

# Microsoft

## Exam Questions DP-700

Implementing Data Engineering Solutions Using Microsoft Fabric (beta)



**NEW QUESTION 1**

HOTSPOT - (Topic 1)

You need to create the product dimension.

How should you complete the Apache Spark SQL code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

```
SELECT ProductID, ProductNumber, ProductName, ModelName, SubCategoryName, CategoryName
FROM ContosoLake.Products p
    ContosoLake.ProductSubCategories s ON p.SubCategoryID = s.SubCategoryID
    ContosoLake.ProductCategories c ON c.CategoryID = s.CategoryID
WHERE
```

The image shows a SQL query editor with three dropdown menus for selecting options. The first dropdown is for the join type between Products and ProductSubCategories, with options: FULL JOIN, INNER JOIN, LEFT ANTI JOIN, LEFT OUTER JOIN, and OUTER JOIN. The second dropdown is for the join type between ProductSubCategories and ProductCategories, with the same options. The third dropdown is for the WHERE clause conditions, with options: CategoryID = 1;, CategoryName is not null;, IsActive = 1;, IsActive is not null;, ProductNumber is not null;, SubCategoryID = 1;, and SubCategoryName is not null;.

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Join between Products and ProductSubCategories: Use an INNER JOIN.

The goal is to include only products that are assigned to a subcategory. An INNER JOIN ensures that only matching records (i.e., products with a valid subcategory) are included.

Join between ProductSubCategories and ProductCategories: Use an INNER JOIN.

Similar to the above logic, we want to include only subcategories assigned to a valid product category. An INNER JOIN ensures this condition is met.

WHERE Clause Condition: IsActive = 1

Only active products (where IsActive equals 1) should be included in the gold layer. This filters out inactive products.

**NEW QUESTION 2**

- (Topic 1)

You need to populate the MAR1 data in the bronze layer.

Which two types of activities should you include in the pipeline? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. ForEach
- B. Copy data
- C. WebHook
- D. Stored procedure

**Answer:** AB

**Explanation:**

MAR1 has seven entities, each accessible via a different API endpoint. A ForEach activity is required to iterate over these endpoints to fetch data from each one. It enables dynamic execution of API calls for each entity.

The Copy data activity is the primary mechanism to extract data from REST APIs and load it into the bronze layer in Delta format. It supports native connectors for REST APIs and Delta, minimizing development effort.

You need to schedule the population of the medallion layers to meet the technical requirements.

What should you do?

- \* A. Schedule a data pipeline that calls other data pipelines.
- \* B. Schedule a notebook.
- \* C. Schedule an Apache Spark job.
- \* D. Schedule multiple data pipelines.

\* Answer: A

The technical requirements specify that:

Medallion layers must be fully populated sequentially (bronze silver gold). Each layer must be populated before the next.

If any step fails, the process must notify the data engineers. Data imports should run simultaneously when possible.

Why Use a Data Pipeline That Calls Other Data Pipelines?

A data pipeline provides a modular and reusable approach to orchestrating the sequential population of medallion layers.

By calling other pipelines, each pipeline can focus on populating a specific layer (bronze, silver, or gold), simplifying development and maintenance.

A parent pipeline can handle:

- Sequential execution of child pipelines.
- Error handling to send email notifications upon failures.
- Parallel execution of tasks where possible (e.g., simultaneous imports into the bronze layer).

**NEW QUESTION 3**

- (Topic 3)

You have an Azure event hub. Each event contains the following fields: BikepointID

Street Neighbourhood

Latitude Longitude No\_Bikes No\_Empty\_Docks

You need to ingest the events. The solution must only retain events that have a Neighbourhood value of Chelsea, and then store the retained events in a Fabric lakehouse.

What should you use?

- A. a KQL queryset
- B. an eventstream
- C. a streaming dataset
- D. Apache Spark Structured Streaming

**Answer: B**

**Explanation:**

An eventstream is the best solution for ingesting data from Azure Event Hub into Fabric, while applying filtering logic such as retaining only the events that have a Neighbourhood value of "Chelsea." Eventstreams in Microsoft Fabric are designed for handling real-time data streams and can apply transformation logic directly on incoming events. In this case, the eventstream can filter events based on the Neighbourhood field before storing the retained events in a Fabric lakehouse. Eventstreams are well-suited for stream processing, such as this case where you need to filter out only specific data (events with a Neighbourhood of "Chelsea") before storing it in the lakehouse.

**NEW QUESTION 4**

- (Topic 3)

You have a Fabric warehouse named DW1 that loads data by using a data pipeline named Pipeline1. Pipeline1 uses a Copy data activity with a dynamic SQL source. Pipeline1 is scheduled to run every 15 minutes.

You discover that Pipeline1 keeps failing.

You need to identify which SQL query was executed when the pipeline failed. What should you do?

- A. From Monitoring hub, select the latest failed run of Pipeline1, and then view the output JSON.
- B. From Monitoring hub, select the latest failed run of Pipeline1, and then view the input JSON.
- C. From Real-time hub, select Fabric events, and then review the details of Microsoft.Fabric.ItemReadFailed.
- D. From Real-time hub, select Fabric events, and then review the details of Microsoft.Fabric.ItemUpdateFailed.
- E. Fabric.ItemUpdateFailed.

**Answer: B**

**Explanation:**

The input JSON contains the configuration details and parameters passed to the Copy data activity during execution, including the dynamically generated SQL query.

Viewing the input JSON for the failed pipeline run provides direct insight into what query was executed at the time of failure.

**NEW QUESTION 5**

HOTSPOT - (Topic 3)

You have three users named User1, User2, and User3.

You have the Fabric workspaces shown in the following table.

Name	Workspace admin
Workspace1	User1
Workspace2	User2

You have a security group named Group1 that contains User1 and User3. The Fabric admin creates the domains shown in the following table.

Name	Domain admin
Domain1	User1
Domain2	User2

User1 creates a new workspace named Workspace3. You add Group1 to the default domain of Domain1. For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

**Answer Area**

Statements	Yes	No
User3 has Viewer role access to Workspace3.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
User3 has Domain contributor access to Domain1.	<input type="checkbox"/>	<input type="checkbox"/>
User2 has Contributor role access to Workspace3.	<input type="checkbox"/>	<input type="checkbox"/>

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

Statements	Yes	No
User3 has Viewer role access to Workspace3.	<input checked="" type="checkbox"/>	<input type="checkbox"/>
User3 has Domain contributor access to Domain1.	<input type="checkbox"/>	<input checked="" type="checkbox"/>
User2 has Contributor role access to Workspace3.	<input type="checkbox"/>	<input checked="" type="checkbox"/>

**NEW QUESTION 6**

- (Topic 3)

You have a Fabric workspace. You have semi-structured data. You need to read the data by using T-SQL, KQL, and Apache Spark. The data will only be written by using Spark. What should you use to store the data?

- A. a lakehouse
- B. an eventhouse
- C. a datamart
- D. a warehouse

**Answer:** A

**Explanation:**

A lakehouse is the best option for storing semi-structured data when you need to read it using T-SQL, KQL, and Apache Spark. A lakehouse combines the flexibility of a data lake (which can handle semi-structured and unstructured data) with the performance features of a data warehouse. It allows data to be written using Apache Spark and can be queried using different technologies such as T-SQL (for SQL-based querying), KQL (Kusto Query Language for querying), and Apache Spark (for distributed processing). This solution is ideal when dealing with semi-structured data and requiring a versatile querying approach.

**NEW QUESTION 7**


HOTSPOT - (Topic 3)

You have a Fabric workspace named Workspace1 that contains a warehouse named Warehouse2. A team of data analysts has Viewer role access to Workspace1. You create a table by running the following statement.

```
CREATE TABLE [warehouse2].[dbo].[CreditCard]
(
    CreditCard varchar(20) NOT NULL
    ,CreditCardType varchar(10) NOT NULL)
GO
```

You need to ensure that the team can view only the first two characters and the last four characters of the Creditcard attribute. How should you complete the statement? To answer, select the appropriate options in the answer area.  
 NOTE: Each correct selection is worth one point.

**Answer Area**



- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Answer Area**

ALTER TABLE dbo.CreditCard  
 COLUMN [CreditCard]  
 WITH (FUNCTION = 'PARTIAL (2, "XXXXXXXXXX", 4)')

ALTER TABLE  
 COLUMN [CreditCard]  
 WITH (FUNCTION = 'PARTIAL (2, "XXXXXXXXXX", 4)')

ALTER TABLE  
 COLUMN [CreditCard]  
 WITH (FUNCTION = 'PARTIAL (2, "XXXXXXXXXX", 4)')

**NEW QUESTION 8**

HOTSPOT - (Topic 3)

You have a Fabric workspace that contains two lakehouses named Lakehouse1 and Lakehouse2. Lakehouse1 contains staging data in a Delta table named Orderlines. Lakehouse2 contains a Type 2 slowly changing dimension (SCD) dimension table named Dim\_Customer.

You need to build a query that will combine data from Orderlines and Dim\_Customer to create a new fact table named Fact\_Orders. The new table must meet the following requirements:

Enable the analysis of customer orders based on historical attributes. Enable the analysis of customer orders based on the current attributes.

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

NOTE: Each correct selection is worth one point.

**Answer Area**

```
SELECT
    OrderLineID order_line_id
    ,OrderDate order_date
    ,c.customer_key
    ,c.customer_id
    ,Quantity order_quantity
    ,unitprice unit_price
    ,taxrate tax_rate
FROM
    Lakehouse1.orderlines o
INNER JOIN
    Lakehouse2.dim_customer c
    ON o.customerid = c.customer_id

AND [ ]
    c.is_current = 1
    o.OrderDate >= c.valid_to_datetime
    o.OrderDate >= c.valid_from_datetime

AND [ ]
    c.is_current = 1
    o.OrderDate <= c.valid_to_datetime
    o.OrderDate <= c.valid_from_datetime
```

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

Answer Area

SELECT

```
OrderLineID order_line_id
,OrderDate order_date
,c.customer_key
,c.customer_id
,Quantity order_quantity
,unitPrice unit_price
,taxRate tax_rate
```

FROM

```
Lakehouse1.orderlines o
```

INNER JOIN

```
Lakehouse2.dim_customer c
ON o.customerid = c.customer_id
```

AND

```
c.is_current = 1
o.OrderDate <= c.valid_to_datetime
o.OrderDate >= c.valid_from_datetime
```

AND

```
c.is_current = 1
o.OrderDate <= c.valid_to_datetime
o.OrderDate <= c.valid_from_datetime
```

NEW QUESTION 9

- (Topic 3)

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 KQL database that contains two tables named Stream and Reference. Stream contains streaming data in the following format.

Column name	Data type
Timestamp	Datetime
GeoLocation	Dynamic
Temperature	Decimal
DeviceId	Int

Reference contains reference data in the following format.

Column name	Data type
DeviceId	Int
DeviceName	String

Both tables contain millions of rows.  
 You have the following KQL queryset.

01 Stream

02 | extend lat = todecimal(GeoLocation.Latitude), long = todecimal(GeoLocation.Longitude)

03 | join kind=inner Reference on DeviceId

04 | project Timestamp, lat, long, Temperature, DeviceName

05 | filter Temperature >= 10

06 | render scatterchart with (kind = map)

You need to reduce how long it takes to run the KQL queryset. Solution: You change the join type to kind=outer.  
 Does this meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:**

An outer join will include unmatched rows from both tables, increasing the dataset size and processing time. It does not improve query performance.

**NEW QUESTION 10**

HOTSPOT - (Topic 3)

You have a Fabric workspace.

You are debugging a statement and discover the following issues: Sometimes, the statement fails to return all the expected rows.

The PurchaseDate output column is NOT in the expected format of mmm dd, yy.

You need to resolve the issues. The solution must ensure that the data types of the results are retained. The results can contain blank cells.

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

NOTE: Each correct selection is worth one point.

## Answer Area

SELECT

item\_id as ItemId

▼

```

,convert(varchar(20), item_name)
,convert(varchar(max), item_name)
try_cast(item_name as varchar(20))
        
```

as ItemName

,item\_description as ItemDescription

▼

```

,convert(varchar, purchase_date, 7)
,convert(varchar, purchase_date, 109)
,convert(varchar, purchase_date, 112)
        
```

as PurchaseDate

FROM

Table1

WHERE

item\_type = @itemtype\_parameter

- A. Mastered
- B. Not Mastered


**Answer:** A


**Explanation:**

## Answer Area

SELECT

item\_id as ItemId

 as ItemName  
 ,convert(varchar(20), item\_name)  
 ,convert(varchar(max), item\_name)  
 ,try\_cast(item\_name as varchar(20))  
 \_\_\_\_\_  
 ,item\_description as ItemDescription

 as PurchaseDate  
 ,convert(varchar, purchase\_date, 7)  
 ,convert(varchar, purchase\_date, 109)  
 ,convert(varchar, purchase\_date, 112)

FROM

Table1

WHERE

item\_type = @itemtype\_parameter

### NEW QUESTION 10

HOTSPOT - (Topic 3)

You have a Fabric workspace that contains a lakehouse named Lakehouse1. Lakehouse1 contains a table named Status\_Target that has the following columns:

- Key
- Status
- LastModified

The data source contains a table named Status\_Source that has the same columns as Status\_Target. Status\_Source is used to populate Status\_Target. In a notebook name Notebook1, you load Status\_Source to a DataFrame named sourceDF and Status\_Target to a DataFrame named targetDF. You need to implement an incremental loading pattern by using Notebook1. The solution must meet the following requirements:

- For all the matching records that have the same value of key, update the value of LastModified in Status\_Target to the value of LastModified in Status\_Source.
- Insert all the records that exist in Status\_Source that do NOT exist in Status\_Target.
- Set the value of Status in Status\_Target to inactive for all the records that were last modified more than seven days ago and that do NOT exist in Status\_Source.

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

NOTE: Each correct selection is worth one point.

Answer Area

```

...
(targetDF
  .merge(sourceDF, "sourceDF.Key" = "targetDF.Key")
    .whenMatchedUpdate(
    .whenMatchedInsert(
    .whenMatchedUpdate(
  ) .whenNotMatchedBySourceInsert(
  .whenNotMatchedBySourceUpdate(
  .whenNotMatchedInsert(
  .whenNotMatchedUpdate(
)
  .whenNotMatchedInsert(
  .whenMatchedInsert(
  .whenMatchedUpdate(
  .whenNotMatchedBySourceInsert(
  .whenNotMatchedBySourceUpdate(
  .whenNotMatchedInsert(
  .whenNotMatchedUpdate(
}
)
  .whenNotMatchedBySourceUpdate(
  .whenMatchedInsert(
  .whenMatchedUpdate(
  .whenNotMatchedBySourceInsert(
  .whenNotMatchedBySourceUpdate(
  .whenNotMatchedInsert(
  .whenNotMatchedUpdate(
)
)
)

```

ent\_date() - INTERVAL '7' DAY)",

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**



Name	Type
Notebook1	Notebook
Notebook2	Notebook
Lakehouse1	Lakehouse
Pipeline1	Data pipeline
Model1	Semantic model

For Model1, the Keep your Direct Lake data up to date option is disabled.

You need to configure the execution of the items to meet the following requirements:

Notebook1 must execute every weekday at 8:00 AM.

Notebook2 must execute when a file is saved to an Azure Blob Storage container. Model1 must refresh when Notebook1 has executed successfully.

How should you orchestrate each item? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

Notebook1:

Notebook2:

Pipeline1:

Model1:

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

## Answer Area

**Notebook1:**

- Add Notebook1 to an Apache Spark job definition.
- Add Notebook1 to Pipeline1.
- From Real-Time hub, configure the execution of Notebook1.

**Notebook2:**

- Add Notebook2 to an Apache Spark job definition.
- Add Notebook2 to Pipeline1.
- From Real-Time hub, configure the execution of Notebook2.

**Pipeline1:**

- Add Pipeline1 to an Apache Spark job definition.
- Configure the execution of Pipeline1 by using a schedule.
- From Real-Time hub, configure the execution of Pipeline1.

**Model1:**

- Add Model1 to Pipeline1.
- From Real-Time hub, configure Model1 to refresh.
- Set Keep your Direct Lake data up to date to On.

### NEW QUESTION 17

- (Topic 3)

You have a Fabric workspace named Workspace1. You plan to integrate Workspace1 with Azure DevOps.

You will use a Fabric deployment pipeline named deployPipeline1 to deploy items from Workspace1 to higher environment workspaces as part of a medallion architecture. You will run deployPipeline1 by using an API call from an Azure DevOps pipeline.

You need to configure API authentication between Azure DevOps and Fabric. Which type of authentication should you use?

- A. service principal
- B. Microsoft Entra username and password
- C. managed private endpoint
- D. workspace identity

**Answer:** A

#### Explanation:

When integrating Azure DevOps with Fabric (Workspace1), using a service principal is the recommended authentication method. A service principal provides a way for applications (such as an Azure DevOps pipeline) to authenticate and interact with resources securely. It allows Azure DevOps to authenticate API calls to Fabric without requiring direct user credentials. This method is ideal for automating tasks such as deploying items through a Fabric deployment pipeline.

### NEW QUESTION 22

- (Topic 3)

You have a Fabric notebook named Notebook1 that has been executing successfully for the last week.

During the last run, Notebook1 executed nine jobs. You need to view the jobs in a timeline chart. What should you use?

- A. Real-Time hub
- B. Monitoring hub
- C. the job history from the application run
- D. Spark History Server
- E. the run series from the details of the application run

**Answer:** E

#### Explanation:

The run series from the details of the application run is the most detailed and relevant feature for visualizing job execution in a timeline format, making it the correct choice for this scenario. It provides an intuitive way to analyze job execution patterns and improve the efficiency of the notebook.

### NEW QUESTION 26

- (Topic 3)

You have a Fabric workspace that contains a warehouse named DW1. DW1 is loaded by using a notebook named Notebook1.

You need to identify which version of Delta was used when Notebook1 was executed. What should you use?

- A. Real-Time hub
- B. OneLake data hub
- C. the Admin monitoring workspace
- D. Fabric Monitor
- E. the Microsoft Fabric Capacity Metrics app

**Answer:** C

#### Explanation:

To identify the version of Delta used when Notebook1 was executed, you should use the Admin monitoring workspace. The Admin monitoring workspace allows you to track and

monitor detailed information about the execution of notebooks and jobs, including the underlying versions of Delta or other technologies used. It provides insights into execution details, including versions and configurations used during job runs, making it the most appropriate choice for identifying the Delta version used during the execution of Notebook1.

**NEW QUESTION 28**

DRAG DROP - (Topic 3)

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

In Warehouse1, you create a table named DimCustomer by running the following statement.

```
CREATE TABLE dbo.DimCustomer (
    CustomerKey VARCHAR(255) NOT NULL,
    Name VARCHAR(255) NOT NULL,
    Email VARCHAR(255) NOT NULL
);
```

You need to set the Customerkey column as a primary key of the DimCustomer table. Which three code segments should you run in sequence? To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.

**Code Segments**

- 0 DROP CONSTRAINT PK\_DimCustomer
- 0 ADD CONSTRAINT PK\_DimCustomer PRIMARY KEY NONCLUSTERED (CustomerKey)
- 0 NOT ENFORCED
- 0 ALTER TABLE dbo.DimCustomer
- 0 ADD CONSTRAINT PK\_DimCustomer PRIMARY KEY CLUSTERED (CustomerKey)
- 0 ENFORCED

**Answer Area**

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

**Code Segments**

- 0 DROP CONSTRAINT PK\_DimCustomer
- 0 ADD CONSTRAINT PK\_DimCustomer PRIMARY KEY NONCLUSTERED (CustomerKey)
- 0 NOT ENFORCED
- 0 ALTER TABLE dbo.DimCustomer
- 0 ADD CONSTRAINT PK\_DimCustomer PRIMARY KEY CLUSTERED (CustomerKey)
- 0 ENFORCED

**Answer Area**

0

0 ADD CONSTRAINT PK\_DimCustomer PRIMARY KEY CLUSTERED (CustomerKey)

0

0 ENFORCED

**NEW QUESTION 32**

- (Topic 3)

You have a Fabric warehouse named DW1 that contains a Type 2 slowly changing dimension (SCD) dimension table named DimCustomer. DimCustomer contains 100 columns and 20 million rows. The columns are of various data types, including int, varchar, date, and varbinary. You need to identify incoming changes to the table and update the records when there is a change. The solution must minimize resource consumption. What should you use to identify changes to attributes?

- A. a direct attributes comparison for the attributes in the source table.
- B. a hash function to compare the attributes in the DimCustomer table.
- C. a direct attributes comparison across the attributes in the DimCustomer table.
- D. a hash function to compare the attributes in the source table.

**Answer: D**

### NEW QUESTION 33

- (Topic 3)

You have a Fabric workspace named Workspace1 that contains a lakehouse named Lakehouse1. Lakehouse1 contains the following tables:

Orders  
Customer Employee  
The Employee table contains Personally Identifiable Information (PII).  
A data engineer is building a workflow that requires writing data to the Customer table, however, the user does NOT have the elevated permissions required to view the contents of the Employee table.

You need to ensure that the data engineer can write data to the Customer table without reading data from the Employee table.

Which three actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. Share Lakehouse1 with the data engineer.
- B. Assign the data engineer the Contributor role for Workspace2.
- C. Assign the data engineer the Viewer role for Workspace2.
- D. Assign the data engineer the Contributor role for Workspace1.
- E. Migrate the Employee table from Lakehouse1 to Lakehouse2.
- F. Create a new workspace named Workspace2 that contains a new lakehouse named Lakehouse2.
- G. Assign the data engineer the Viewer role for Workspace1.

**Answer: ADE**

### Explanation:

To meet the requirements of ensuring that the data engineer can write data to the Customer table without reading data from the Employee table (which contains Personally Identifiable Information, or PII), you can implement the following steps:

? Share Lakehouse1 with the data engineer.

By sharing Lakehouse1 with the data engineer, you provide the necessary access to the data within the lakehouse. However, this access should be controlled through roles and permissions, which will allow writing to the Customer table but prevent reading from the Employee table.

? Assign the data engineer the Contributor role for Workspace1.

Assigning the Contributor role for Workspace1 grants the data engineer the ability to perform actions such as writing to tables (e.g., the Customer table) within the workspace. This role typically allows users to modify and manage data without necessarily granting them access to view all data (e.g., PII data in the Employee table).

? Migrate the Employee table from Lakehouse1 to Lakehouse2.

To prevent the data engineer from accessing the Employee table (which contains PII), you can migrate the Employee table to a separate lakehouse (Lakehouse2) or workspace

(Workspace2). This separation of sensitive data ensures that the data engineer's access is restricted to the Customer table in Lakehouse1, while the Employee table can be managed separately and protected under different access controls.

### NEW QUESTION 35

HOTSPOT - (Topic 3)

You have a Fabric workspace that contains an eventstream named EventStream1. You discover that an EventStream1 transformation fails.

You need to find the following error information: The error details, including the occurrence time The total number of errors

What should you use? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

## Answer Area

To find the error details:

	▼
Data insights	
Data preview	
Details	
Runtime logs	

To find the total number of errors:

	▼
Data insights	
Data preview	
Details	
Runtime logs	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

## Answer Area

To find the error details:

	▼
Data insights	
Data preview	
Details	
Runtime logs	

To find the total number of errors:

	▼
Data insights	
Data preview	
Details	
Runtime logs	

### NEW QUESTION 37

- (Topic 3)

You have a Fabric workspace that contains a semantic model named Model1. You need to dynamically execute and monitor the refresh progress of Model1. What should you use?

- A. dynamic management views in Microsoft SQL Server Management Studio
- B. Monitoring hub

- C. dynamic management views in Azure Data Studio
- D. a semantic link in a notebook

**Answer:** D

**Explanation:**

Semantic models in Microsoft Fabric are part of Power BI datasets and require refreshes to stay updated with the latest data. Dynamically executing and monitoring the refresh progress requires a tool or approach that integrates with Fabric's capabilities for semantic models.

**NEW QUESTION 41**

- (Topic 3)

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 KQL database that contains two tables named Stream and Reference. Stream contains streaming data in the following format.

Column name	Data type
Timestamp	Datetime
GeoLocation	Dynamic
Temperature	Decimal
DeviceId	Int

Reference contains reference data in the following format.

Column name	Data type
DeviceId	Int
DeviceName	String

Both tables contain millions of rows. You have the following KQL queryset.

You need to reduce how long it takes to run the KQL queryset. Solution: You add the make\_list() function to the output columns. Does this meet the goal?

01 Stream

02 | extend lat = todecimal(GeoLocation.Latitude), long = todecimal(GeoLocation.Longitude)

03 | join kind=inner Reference on DeviceId

04 | project Timestamp, lat, long, Temperature, DeviceName

05 | filter Temperature >= 10

06 | render scatterchart with (kind = map)

- A. Yes
- B. No

**Answer:** B

**Explanation:**

Adding an aggregation like make\_list() would require additional processing and memory, which could make the query slower.

**NEW QUESTION 43**

HOTSPOT - (Topic 3)

You have a Fabric warehouse named DW1 that contains four staging tables named ProductCategory, ProductSubcategory, Product, and SalesOrder. ProductCategory, ProductSubcategory, and Product are used often in analytical queries.

You need to implement a star schema for DW1. The solution must minimize development effort.

Which design approach should you use? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

ProductCategory, ProductSubcategory and Product must be:

- Denormalized into a single product dimension table
- Added to the model as individual tables
- Denormalized by being added to the SalesOrder table
- Denormalized into a single product dimension table**

The joining key must be:

- the unique system generated identifier
- The product name and the date
- the unique system generated identifier
- The product category name**

- A. Mastered
- B. Not Mastered

**Answer:** A

**Explanation:**

Answer Area

ProductCategory, ProductSubcategory and Product must be:

- Denormalized into a single product dimension table
- Added to the model as individual tables
- Denormalized by being added to the SalesOrder table
- Denormalized into a single product dimension table**

The joining key must be:

- the unique system generated identifier
- The product name and the date
- the unique system generated identifier
- The product category name**

**NEW QUESTION 45**

- (Topic 3)

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:

```
SELECT BikepointID, Street, Neighbourhood, No_Bikes, No_Empty_Docks, Timestamp
FROM bike_location
WHERE neighbourhood = 'Sands End'
AND no_bikes >= 15
ORDER BY no_bikes
```

Does this meet the goal?

- A. Yes
- B. no

**Answer:** B

**Explanation:**

This code does not meet the goal because this is an SQL-like query and cannot be executed in KQL, which is required for the database.

Correct code should look like:

```
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
```

**NEW QUESTION 47**

HOTSPOT - (Topic 3)

You have a Fabric workspace that contains a warehouse named DW1. DW1 contains the following tables and columns.

Table name	Column name	Description
SalesOrderDetail	ProductID	Contains the product ID of the ordered product
SalesOrderDetail	ModifiedDate	Contains the date of an order
SalesOrderDetail	OrderQty	Contains the order quantity
Product	ProductID	Contains the unique ID of a product
Product	Name	Contains a product name

You need to create an output that presents the summarized values of all the order quantities by year and product. The results must include a summary of the order quantities at the year level for all the products.

How should you complete the code? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

**Answer Area**

```

(SO.ModifiedDate) AS OrderDate
SELECT CAST
SELECT CONVERT
SELECT YEAR
    ,P.Name AS ProductName
    ,SUM(SO.OrderQty) AS OrderQty
FROM [dbo].[SalesOrderDetail] SO
INNER JOIN [dbo].[Product] P
    ON P.ProductID = SO.ProductID
GROUP BY
    CUBE(YEAR(SO.ModifiedDate), P.Name)
    (ROLLUP(CAST(SO.ModifiedDate AS DATE), P.Name), (YEAR(SO.ModifiedDate)))
    ROLLUP(YEAR(SO.ModifiedDate), P.Name)
    YEAR(SO.ModifiedDate), P.Name
ORDER BY OrderDate
  
```

- A. Mastered
- B. Not Mastered

**Answer: A**

**Explanation:**

**Answer Area**

```

(SO.ModifiedDate) AS OrderDate
SELECT CAST
SELECT CONVERT
SELECT YEAR
    ,P.Name AS ProductName
    ,SUM(SO.OrderQty) AS OrderQty
FROM [dbo].[SalesOrderDetail] SO
INNER JOIN [dbo].[Product] P
    ON P.ProductID = SO.ProductID
GROUP BY
    CUBE(YEAR(SO.ModifiedDate), P.Name)
    (ROLLUP(CAST(SO.ModifiedDate AS DATE), P.Name), (YEAR(SO.ModifiedDate)))
    ROLLUP(YEAR(SO.ModifiedDate), P.Name)
    YEAR(SO.ModifiedDate), P.Name
ORDER BY OrderDate
  
```

**NEW QUESTION 49**

- (Topic 3)

You have a Fabric deployment pipeline that uses three workspaces named Dev, Test, and Prod.

You need to deploy an eventhouse as part of the deployment process. What should you use to add the eventhouse to the deployment process?

- A. GitHub Actions
- B. a deployment pipeline
- C. an Azure DevOps pipeline

**Answer: B**

**Explanation:**

A deployment pipeline in Fabric is designed to automate the process of deploying assets (such as reports, datasets, eventhouses, and other objects) between environments like Dev, Test, and Prod. Since you need to deploy an eventhouse as part of the deployment process, a deployment pipeline is the appropriate tool to move this asset through the different stages of your environment.

**NEW QUESTION 50**

- (Topic 3)

You have a Fabric workspace named Workspace1 that contains an Apache Spark job definition named Job1.

You have an Azure SQL database named Source1 that has public internet access disabled.

You need to ensure that Job1 can access the data in Source1. What should you create?

- A. an on-premises data gateway
- B. a managed private endpoint
- C. an integration runtime
- D. a data management gateway

**Answer: B**

**Explanation:**

To allow Job1 in Workspace1 to access an Azure SQL database (Source1) with public internet access disabled, you need to create a managed private endpoint. A managed private endpoint is a secure, private connection that enables services like Fabric (or other Azure services) to access resources such as databases, storage accounts, or other services within a virtual network (VNet) without requiring public internet access. This approach maintains the security and integrity of your data while enabling access to the Azure SQL database.

**NEW QUESTION 54**

- (Topic 3)

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 KQL database that contains two tables named Stream and Reference. Stream contains streaming data in the following format.

Column name	Data type
Timestamp	Datetime
GeoLocation	Dynamic
Temperature	Decimal
DeviceId	Int

Reference contains reference data in the following format.

Column name	Data type
DeviceId	Int
DeviceName	String

Both tables contain millions of rows. You have the following KQL queryset.

You need to reduce how long it takes to run the KQL queryset. Solution: You move the filter to line 02.

```

01 Stream
02 | extend lat = todecimal(GeoLocation.Latitude), long = todecimal(GeoLocation.Longitude)
03 | join kind=inner Reference on DeviceId
04 | project Timestamp, lat, long, Temperature, DeviceName
05 | filter Temperature >= 10
06 | render scatterchart with (kind = map)

```

Does this meet the goal?

- A. Yes
- B. No

**Answer: A**

**Explanation:**

Moving the filter to line 02: Filtering the Stream table before performing the join operation reduces the number of rows that need to be processed during the join. This is an effective optimization technique for queries involving large datasets.

**NEW QUESTION 59**

- (Topic 3)

You have an Azure Data Lake Storage Gen2 account named storage1 and an Amazon S3 bucket named storage2. You have the Delta Parquet files shown in the following table.

Name	Stored in	Size	Description
ProductFile	storage1	50 MB	Contains a list of products and their details
TripsFile	storage2	2 GB	Contains one month's worth of taxi trip data
StoreFile	storage2	25 MB	Contains a list of stores and their addresses

You have a Fabric workspace named Workspace1 that has the cache for shortcuts enabled. Workspace1 contains a lakehouse named Lakehouse1. Lakehouse1 has the following shortcuts:

A shortcut to ProductFile aliased as Products A shortcut to StoreFile aliased as Stores

A shortcut to TripsFile aliased as Trips

The data from which shortcuts will be retrieved from the cache?

- A. Trips and Stores only
- B. Products and Store only
- C. Stores only
- D. Products only
- E. Product
- F. Stores, and Trips

**Answer: B**

**Explanation:**

When the cache for shortcuts is enabled in Fabric, the data retrieval is governed by the caching behavior, which generally retains data for a specific period after it was last accessed. The data from the shortcuts will be retrieved from the cache if the data is stored in locations that support caching. Here's a breakdown based on the data's location: Products: The ProductFile is stored in Azure Data Lake Storage Gen2 (storage1). Since Azure Data Lake is a supported storage system in Fabric and the file is relatively small (50 MB), this data is most likely cached and can be retrieved from the cache.

Stores: The StoreFile is stored in Amazon S3 (storage2), and even though it is stored in a different cloud provider, Fabric can cache data from Amazon S3 if caching is enabled. This data (25 MB) is likely cached and retrievable.

Trips: The TripsFile is stored in Amazon S3 (storage2) and is significantly larger (2 GB) compared to the other files. While Fabric can cache data from Amazon S3, the larger size of the file (2 GB) may exceed typical cache sizes or retention windows, causing this file to likely be retrieved directly from the source instead of the cache.

**NEW QUESTION 64**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **DP-700 Practice Exam Features:**

- \* DP-700 Questions and Answers Updated Frequently
- \* DP-700 Practice Questions Verified by Expert Senior Certified Staff
- \* DP-700 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* DP-700 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The DP-700 Practice Test Here](#)**