

# IT 认证电子书



质 量 更 高 服 务 更 好

半年免费升级服务

<http://www.itrenzheng.com>

**Exam** : **DP-600**

**Title** : **Implementing Analytics  
Solutions Using Microsoft  
Fabric**

**Version** : **DEMO**

## 1. Topic 1, Litware. Inc. Case Study

### Overview

Litware. Inc. is a manufacturing company that has offices throughout North America. The analytics team at Litware contains data engineers, analytics engineers, data analysts, and data scientists.

### Existing Environment

litware has been using a Microsoft Power BI tenant for three years. Litware has NOT enabled any Fabric capacities and features.

### Fabric Environment

Litware has data that must be analyzed as shown in the following table.

| Description                   | Original source                               | Total size |
|-------------------------------|---|------------|
| Customer data                 | Customer relationship management (CRM) system | 50 MB      |
| Product data                  | Customer relationship management (CRM) system | 200 MB     |
| Customer satisfaction surveys | SurveyMonkey                                  | 500 GB     |

The Product data contains a single table and the following columns.

| Name            | Data type |
|-----------------|-----------|
| ProductID       | Integer   |
| ProductName     | String    |
| ProductCategory | String    |
| ListPrice       | Decimal   |

The customer satisfaction data contains the following tables:

- Survey
- Question
- Response

For each survey submitted, the following occurs:

- One row is added to the Survey table.
- One row is added to the Response table for each question in the survey.

The Question table contains the text of each survey question. The third question in each survey response is an overall satisfaction score. Customers can submit a survey after each purchase.

### User Problems

The analytics team has large volumes of data, some of which is semi-structured. The team wants to use Fabric to create a new data store.

Product data is often classified into three pricing groups: high, medium, and low. This logic is implemented in several databases and semantic models, but the logic does NOT always match across

implementations.

### **Planned Changes**

Litware plans to enable Fabric features in the existing tenant. The analytics team will create a new data store as a proof of concept (PoC). The remaining Litware users will only get access to the Fabric features once the PoC is complete. The PoC will be completed by using a Fabric trial capacity.

The following three workspaces will be created:

- AnalyticsPOC: Will contain the data store, semantic models, reports, pipelines, dataflows, and notebooks used to populate the data store
- DataEngPOC: Will contain all the pipelines, dataflows, and notebooks used to populate Onelake
- DataSciPOC: Will contain all the notebooks and reports created by the data scientists

The following will be created in the AnalyticsPOC workspace:

- A data store (type to be decided)
- A custom semantic model
- A default semantic model
- Interactive reports

The data engineers will create data pipelines to load data to OneLake either hourly or daily depending on the data source. The analytics engineers will create processes to ingest transform, and load the data to the data store in the AnalyticsPOC workspace daily. Whenever possible, the data engineers will use low-code tools for data ingestion. The choice of which data cleansing and transformation tools to use will be at the data engineers' discretion.

All the semantic models and reports in the Analytics POC workspace will use the data store as the sole data source.

### **Technical Requirements**

The data store must support the following:

- Read access by using T-SQL or Python
- Semi-structured and unstructured data
- Row-level security (RLS) for users executing T-SQL queries

Files loaded by the data engineers to OneLake will be stored in the Parquet format and will meet Delta Lake specifications.

Data will be loaded without transformation in one area of the AnalyticsPOC data store. The data will then be cleansed, merged, and transformed into a dimensional model.

The data load process must ensure that the raw and cleansed data is updated completely before populating the dimensional model.

The dimensional model must contain a date dimension. There is no existing data source for the date dimension. The Litware fiscal year matches the calendar year. The date dimension must always contain dates from 2010 through the end of the current year.

The product pricing group logic must be maintained by the analytics engineers in a single location. The pricing group data must be made available in the data store for T-SQL queries and in the default semantic model.

The following logic must be used:

- List prices that are less than or equal to 50 are in the low pricing group.
- List prices that are greater than 50 and less than or equal to 1,000 are in the medium pricing group.
- List prices that are greater than 1,000 are in the high pricing group.

### **Security Requirements**

Only Fabric administrators and the analytics team must be able to see the Fabric items created as part of the PoC. Litware identifies the following security requirements for the Fabric items in the AnalyticsPOC workspace:

- Fabric administrators will be the workspace administrators.
- The data engineers must be able to read from and write to the data store. No access must be granted to datasets or reports.
- The analytics engineers must be able to read from, write to, and create schemas in the data store. They also must be able to create and share semantic models with the data analysts and view and modify all reports in the workspace.
- The data scientists must be able to read from the data store, but not write to it. They will access the data by using a Spark notebook.
- The data analysts must have read access to only the dimensional model objects in the data store. They also must have access to create Power BI reports by using the semantic models created by the analytics engineers.
- The date dimension must be available to all users of the data store.
- The principle of least privilege must be followed.

Both the default and custom semantic models must include only tables or views from the dimensional model in the data store.

Litware already has the following Microsoft Entra security groups:

- FabricAdmins: Fabric administrators
- AnalyticsTeam: All the members of the analytics team
- DataAnalysts: The data analysts on the analytics team
- DataScientists: The data scientists on the analytics team
- Data Engineers: The data engineers on the analytics team
- Analytics Engineers: The analytics engineers on the analytics team

### **Report Requirements**

The data analysis must create a customer satisfaction report that meets the following requirements:

- Enables a user to select a product to filter customer survey responses to only those who have purchased that product
- Displays the average overall satisfaction score of all the surveys submitted during the last 12 months up to a selected date
- Shows data as soon as the data is updated in the data store
- Ensures that the report and the semantic model only contain data from the current and previous year
- Ensures that the report respects any table-level security specified in the source data store

- Minimizes the execution time of report queries

### HOTSPOT

You to need assign permissions for the data store in the AnalyticsPOC workspace. The solution must meet the security requirements.

Which additional permissions should you assign when you share the data store? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

#### Answer Area

|                 |   |
|-----------------|---|
| DataEngineers:  | <input type="checkbox"/> Build Reports on the default dataset<br><input type="checkbox"/> Read All Apache Spark<br><input type="checkbox"/> Read All SQL analytics endpoint data            |
| DataAnalysts:   | <input type="checkbox"/> Build Reports on the default dataset<br><input checked="" type="checkbox"/> Read All Apache Spark<br><input type="checkbox"/> Read All SQL analytics endpoint data |
| DataScientists: | <input type="checkbox"/> Build Reports on the default dataset<br><input type="checkbox"/> Read All Apache Spark<br><input checked="" type="checkbox"/> Read All SQL analytics endpoint data |

#### Answer:

#### Answer Area

|                 |  |
|-----------------|--|
| DataEngineers:  | <input checked="" type="checkbox"/> Build Reports on the default dataset<br><input type="checkbox"/> Read All Apache Spark<br><input checked="" type="checkbox"/> Read All SQL analytics endpoint data |
| DataAnalysts:   | <input type="checkbox"/> Build Reports on the default dataset<br><input checked="" type="checkbox"/> Read All Apache Spark<br><input type="checkbox"/> Read All SQL analytics endpoint data            |
| DataScientists: | <input type="checkbox"/> Build Reports on the default dataset<br><input type="checkbox"/> Read All Apache Spark<br><input checked="" type="checkbox"/> Read All SQL analytics endpoint data            |

#### Explanation:

Data Engineers: Read All SQL analytics endpoint data

Data Analysts: Read All Apache Spark

Data Scientists: Read All SQL analytics endpoint data

The permissions for the data store in the Analytics POC workspace should align with the principle of least privilege:

Data Engineers need read and write access but not to datasets or reports.

Data Analysts require read access specifically to the dimensional model objects and the ability to create Power BI reports.

Data Scientists need read access via Spark notebooks. These settings ensure each role has the

necessary permissions to fulfill their responsibilities without exceeding their required access level.

## 2.HOTSPOT

You need to create a DAX measure to calculate the average overall satisfaction score.

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

NOTE: Each correct selection is worth one point.

### Answer Area

```
Rolling 12 Overall Satisfaction =  
VAR NumberOfMonths = 12  
VAR LastCurrentDate = MAX ( 'Date'[Date] )  
VAR Period = DATESINPERIOD ( 'Date'[Date], LastCurrentDate, - NumberOfMonths, MONTH )  
VAR Result =  
    CALCULATE (  
        AVERAGE('Survey'[Response Value]),  
        AVERAGE('Survey'[Response Value]),  
        AVERAGEA('Question'[Question Text]),  
        AVERAGEX(VALUES('Survey'[Customer Key]),  
            NumberOfMonths,  
            LastCurrentDate,  
            NumberOfMonths,  
            Period,  
            'Survey Question'[Question Title] = "Overall Satisfaction"  
        )  
    )  
RETURN  
    Result
```

**Answer:**

**Answer Area**

```
Rolling 12 Overall Satisfaction =
VAR NumberOfMonths = 12
VAR LastCurrentDate = MAX ( 'Date'[Date] )
VAR Period = DATESINPERIOD ( 'Date'[Date], LastCurrentDate, - NumberOfMonths, MONTH )
VAR Result =
```

```
    CALCULATE (
```

The screenshot shows two dropdown menus. The first menu is open, showing the following options:
 

- AVERAGE('Survey'[Response Value])
- AVERAGE('Survey'[Response Value]),**
- AVERAGEA('Question'[Question Text]),
- AVERAGEX(VALUES('Survey'[Customer Key]),

 The second menu is also open, showing the following options:
 

- NumberOfMonths,
- LastCurrentDate,
- NumberOfMonths,**
- Period,

 Red boxes highlight the selected options in both menus.

```
        'Survey Question'[Question Title] = "Overall Satisfaction"
```

```
    )
```

```
RETURN
```

```
    Result
```

**Explanation:**

The measure should use the AVERAGE function to calculate the average value.

It should reference the Response Value column from the 'Survey' table.

The 'Number of months' should be used to define the period for the average calculation.

To calculate the average overall satisfaction score using DAX, you would need to use the AVERAGE function on the response values related to satisfaction questions. The DATESINPERIOD function will help in calculating the rolling average over the last 12 months.

**3.HOTSPOT**

You need to resolve the issue with the pricing group classification.

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.



**Answer Area**

```
CREATE [dbo].[ProductsWithPricingGroup]
AS
SELECT ProductCategory,
Product,
ListPrice,
CASE
WHEN (ListPrice >= 50 AND ListPrice < 1000 ) THEN 'medium'
WHEN (ListPrice > 50 AND ListPrice <= 1000 ) THEN 'medium'
WHEN (ListPrice >= 50 AND ListPrice < 1000 ) THEN 'medium'
WHEN ListPrice BETWEEN 50 AND 1000 ) THEN 'medium'
END AS PricingGroup
FROM dbo.Products
```

**Answer:**

**Answer Area**

```
CREATE [dbo].[ProductsWithPricingGroup]
AS
SELECT ProductCategory,
Product,
ListPrice,
CASE
WHEN (ListPrice >= 50 AND ListPrice < 1000 ) THEN 'medium'
WHEN (ListPrice > 50 AND ListPrice <= 1000 ) THEN 'medium'
WHEN (ListPrice >= 50 AND ListPrice < 1000 ) THEN 'medium'
WHEN ListPrice BETWEEN 50 AND 1000 ) THEN 'medium'
END AS PricingGroup
FROM dbo.Products
```

**Explanation:**

You should use CREATE VIEW to make the pricing group logic available for T-SQL queries. The CASE statement should be used to determine the pricing group based on the list price.

The T-SQL statement should create a view that classifies products into pricing groups based on the list price. The CASE statement is the correct conditional logic to assign each product to the appropriate pricing group. This view will standardize the pricing group logic across different databases and semantic models.

4. What should you recommend using to ingest the customer data into the data store in the AnalyticsPOC workspace?

- A. a stored procedure
- B. a pipeline that contains a KQL activity
- C. a Spark notebook
- D. a dataflow

**Answer:** D

**Explanation:**

For ingesting customer data into the data store in the AnalyticsPOC workspace, a dataflow (D) should be recommended. Dataflows are designed within the Power BI service to ingest, cleanse, transform, and load data into the Power BI environment. They allow for the low-code ingestion and transformation of data as needed by Litware's technical requirements.

Reference = You can learn more about dataflows and their use in Power BI environments in Microsoft's Power BI documentation.

5. Which type of data store should you recommend in the AnalyticsPOC workspace?

- A. a data lake
- B. a warehouse
- C. a lakehouse
- D. an external Hive metaStore

**Answer:** C

**Explanation:**

A lakehouse (C) should be recommended for the AnalyticsPOC workspace. It combines the capabilities of a data warehouse with the flexibility of a data lake. A lakehouse supports semi-structured and unstructured data and allows for T-SQL and Python read access, fulfilling the technical requirements outlined for Litware.

Reference = For further understanding, Microsoft's documentation on the lakehouse architecture provides insights into how it supports various data types and analytical operations.