Number: DP-600
Passing Score: 800
Time Limit: 120
File Version: 5.0

Exam Code: DP-600

Exam Name: Implementing Analytics Solutions Using Microsoft Fabric

Exam A

QUESTION 1
You have a Fabric tenant that contains a workspace named Workspace^ Workspacel is assigned to a Fabric capacity.
You need to recommend a solution to provide users with the ability to create and publish custom Direct Lake semantic models by using external tools. The solution must follow the principle of least privilege.
Which three actions in the Fabric Admin portal should you include in the recommendation? Each correct answer presents part of the solution.
NOTE: Each correct answer is worth one point.

A. From the Tenant settings, set Allow XMLA Endpoints and Analyze in Excel with on-premises datasets to Enabled
B. From the Tenant settings, set Allow Azure Active Directory guest users to access Microsoft Fabric to Enabled
C. From the Tenant settings, select Users can edit data models in the Power Bl service.
D. From the Capacity settings, set XMLA Endpoint to Read Write
E. From the Tenant settings, set Users can create Fabric items to Enabled
F. From the Tenant settings, enable Publish to Web

Correct Answer: A, C, D
Section:
Explanation:
For users to create and publish custom Direct Lake semantic models using external tools, following the principle of least privilege, the actions to be included are enabling XMLA Endpoints (A), editing data models in Power BI service (C), and setting XMLA Endpoint to Read-Write in the capacity settings (D). Reference = More information can be found in the Admin portal of the Power BI service documentation, detailing tenant and capacity settings.

QUESTION 2
DRAG DROP
You are implementing two dimension tables named Customers and Products in a Fabric warehouse.
You need to use slowly changing dimension (SCO) to manage the versioning of data. The solution must meet the requirements shown in the following table.

Table	Change action
Customers	Create a new version of the row.
Products	Overwrite the existing value in the latest row.

Which type of SCD should you use for each table? To answer, drag the appropriate SCD types to the correct tables. Each SCD type 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:

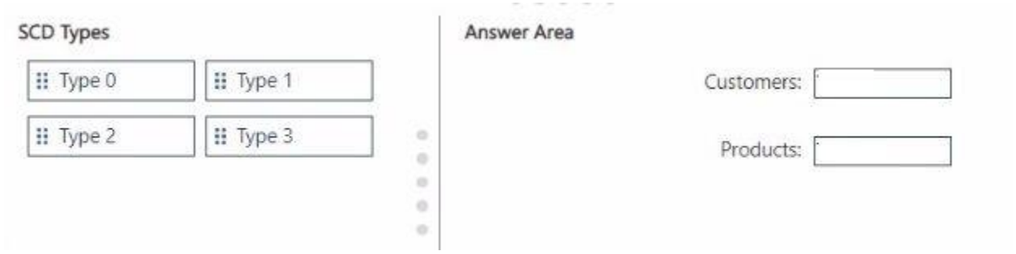

Correct Answer:

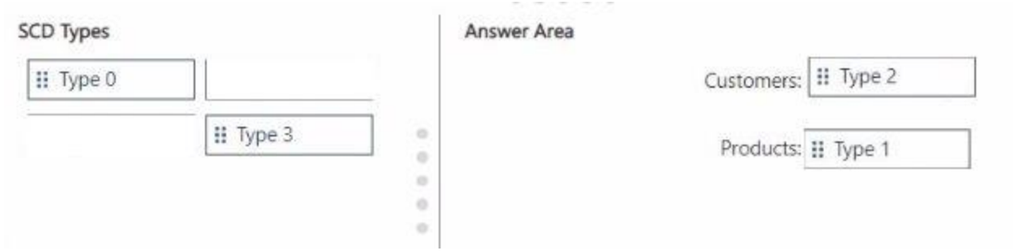

Section:
Explanation:

QUESTION 3
You have a Fabric tenant that contains a lakehouse. You plan to use a visual query to merge two tables.
You need to ensure that the query returns all the rows that are present in both tables. Which type of join should you use?

A. left outer
B. right anti
C. full outer
D. left anti
E. right outer
F. inner

Correct Answer: C
Section:
Explanation:
When you need to return all rows that are present in both tables, you use a full outer join. This type of join combines the results of both left and right outer joins and returns all rows from both tables, with matching rows from both sides where available. If there is no match, the result is NULL on the side of the join where there is no match.

QUESTION 4
You have a Fabric tenant.
You are creating a Fabric Data Factory pipeline.
You have a stored procedure that returns the number of active customers and their average sales for the current month.
You need to add an activity that will execute the stored procedure in a warehouse. The returned values must be available to the downstream activities of the pipeline.
Which type of activity should you add?

A. Stored procedure
B. Get metadata
C. Lookup
D. Copy data

Correct Answer: C
Section:
Explanation:
In a Fabric Data Factory pipeline, to execute a stored procedure and make the returned values available for downstream activities, the Lookup activity is used. This activity can retrieve a dataset from a data store and pass it on for further processing. Here's how you would use the Lookup activity in this context:
Add a Lookup activity to your pipeline.
Configure the Lookup activity to use the stored procedure by providing the necessary SQL statement or stored procedure name.
In the settings, specify that the activity should use the stored procedure mode.

Once the stored procedure executes, the Lookup activity will capture the results and make them available in the pipeline's memory.
Downstream activities can then reference the output of the Lookup activity.

QUESTION 5
DRAG DROP
You are implementing a medallion architecture in a single Fabric workspace.
You have a lakehouse that contains the 8ronze and Silver layers and a warehouse that contains the Gold layer.
You create the items required to populate the layers as shown in the following table.

Layer	Data integration tool
Bronze	Pipelines with Copy activities
Silver	Dataflows
Gold	Stored procedures

You need to ensure that the layers are populated daily in sequential order such that Silver is populated only after Bronze is complete, and Gold is populated only after Silver is complete. The solution must minimize development effort and complexity.
What should you use to execute each set of items? To answer, drag the appropriate options to the correct items. Each option 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:

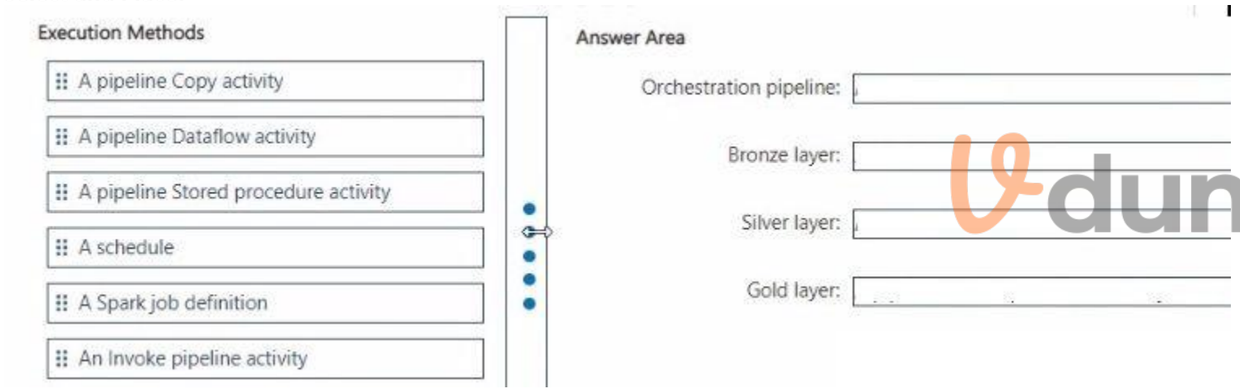

Correct Answer:

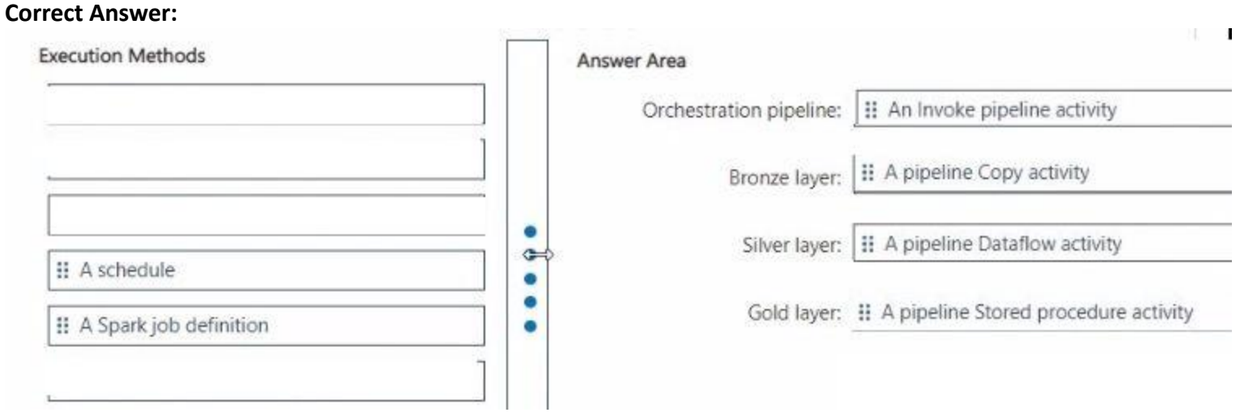

Section:
Explanation:

QUESTION 6
You have a Fabric notebook that has the Python code and output shown in the following exhibit.

```
# Look at a histogram of tips by count by using Matplotlib

ax1 = sampled_taxi_pd_df['tipAmount'].plot(kind='hist', bins=25, facecolor='lightblue'
ax1.set_title('Tip amount distribution')
ax1.set_xlabel('Tip Amount ($)')
ax1.set_ylabel('Counts')
plt.suptitle('')
plt.show()
```

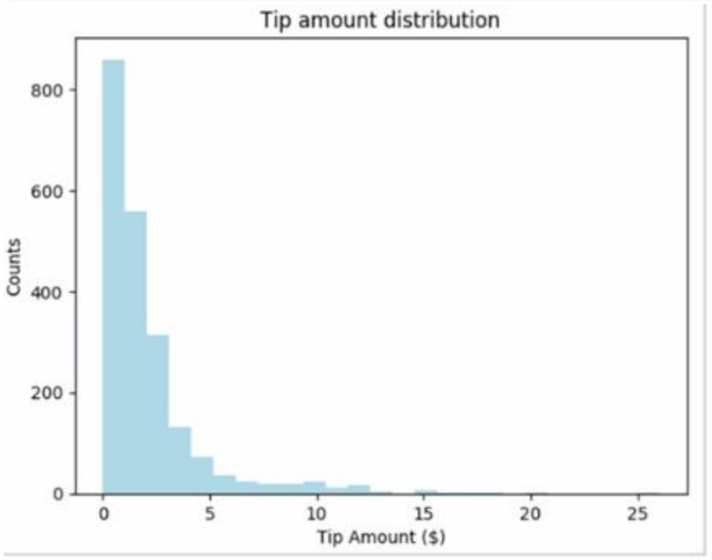

Which type of analytics are you performing?

A. predictive
B. descriptive
C. prescriptive
D. diagnostic

Correct Answer: B
Section:
Explanation:
The Python code and output shown in the exhibit display a histogram, which is a representation of the distribution of data. This kind of analysis is descriptive analytics, which is used to describe or summarize the features of a dataset. Descriptive analytics answers the question of 'what has happened' by providing insight into past data through tools such as mean, median, mode, standard deviation, and graphical representations like histograms.

QUESTION 7
You are creating a semantic model in Microsoft Power Bl Desktop.
You plan to make bulk changes to the model by using the Tabular Model Definition Language (TMDL) extension for Microsoft Visual Studio Code.
You need to save the semantic model to a file.
Which file format should you use?

A. PBIP
B. PBIX
C. PBIT
D. PBIDS

Correct Answer: B
Section:
Explanation:
When saving a semantic model to a file that can be edited using the Tabular Model Scripting Language (TMSL) extension for Visual Studio Code, the PBIX (Power BI Desktop) file format is the correct choice. The PBIX format contains the report, data model, and queries, and is the primary file format for editing in Power BI Desktop. Reference = Microsoft's documentation on Power BI file formats and Visual Studio Code provides further clarification on the usage of PBIX files.

QUESTION 8
HOTSPOT
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:

Answer Area:

Section:
Explanation:

QUESTION 9
You need to provide Power Bl developers with access to the pipeline. The solution must meet the following requirements:
* Ensure that the developers can deploy items to the workspaces for Development and Test.
* Prevent the developers from deploying items to the workspace for Production.
* Follow the principle of least privilege.
Which three levels of access should you assign to the developers? Each correct answer presents part of the solution. NOTE: Each correct answer is worth one point.

A. Build permission to the production semantic models
B. Admin access to the deployment pipeline
C. Viewer access to the Development and Test workspaces
D. Viewer access to the Production workspace

E. Contributor access to the Development and Test workspaces
F. Contributor access to the Production workspace

Correct Answer: B, D, E
Section:
Explanation:
To meet the requirements, developers should have Admin access to the deployment pipeline (B), Contributor access to the Development and Test workspaces (E), and Viewer access to the Production workspace (D). This setup ensures they can perform necessary actions in development and test environments without having the ability to affect production. Reference = The Power BI documentation on workspace access levels and deployment pipelines provides guidelines on assigning appropriate permissions.

QUESTION 10
You have a Fabric workspace that contains a DirectQuery semantic model. The model queries a data source that has 500 million rows.
You have a Microsoft Power Bl report named Report1 that uses the model. Report! contains visuals on multiple pages.
You need to reduce the query execution time for the visuals on all the pages.
What are two features that you can use? Each correct answer presents a complete solution.
NOTE: Each correct answer is worth one point.

A. user-defined aggregations
B. automatic aggregation
C. query caching
D. OneLake integration

Correct Answer: A, B
Section:
Explanation:
User-defined aggregations (A) and query caching (C) are two features that can help reduce query execution time. User-defined aggregations allow precalculation of large datasets, and query caching stores the results of queries temporarily to speed up future queries. Reference = Microsoft Power BI documentation on performance optimization offers in-depth knowledge on these features.

QUESTION 11
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

Correct Answer: C
Section:
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.

QUESTION 12
You need to recommend a solution to prepare the tenant for the PoC.
Which two actions should you recommend performing from the Fabric Admin portal? Each correct answer presents part of the solution.
NOTE: Each correct answer is worth one point.

A. Enable the Users can try Microsoft Fabric paid features option for specific security groups.
B. Enable the Allow Azure Active Directory guest users to access Microsoft Fabric option for specific security groups.
C. Enable the Users can create Fabric items option and exclude specific security groups.
D. Enable the Users can try Microsoft Fabric paid features option for the entire organization.
E. Enable the Users can create Fabric items option for specific security groups.

Correct Answer: A, E
Section:
Explanation:
The PoC is planned to be completed using a Fabric trial capacity, which implies that users involved in the PoC should be able to try paid features. However, this should be limited to specific security groups involved in the PoC to prevent the entire organization from accessing these features before the trial is proven successful (A). The ability for users to create Fabric items should also be enabled for specific security groups to ensure that only the relevant team members participating in the PoC can create items in the Fabric environment (E).

QUESTION 13
You need to ensure the data loading activities in the AnalyticsPOC workspace are executed in the appropriate sequence. The solution must meet the technical requirements.
What should you do?

A. Create a pipeline that has dependencies between activities and schedule the pipeline.
B. Create and schedule a Spark job definition.
C. Create a dataflow that has multiple steps and schedule the dataflow.
D. Create and schedule a Spark notebook.

Correct Answer: A
Section:
Explanation:
To meet the technical requirement that data loading activities must ensure the raw and cleansed data is updated completely before populating the dimensional model, you would need a mechanism that allows for ordered execution. A pipeline in Microsoft Fabric with dependencies set between activities can ensure that activities are executed in a specific sequence. Once set up, the pipeline can be scheduled to run at the required intervals (hourly or daily depending on the data source).

QUESTION 14
You need to implement the date dimension in the data store. The solution must meet the technical requirements.
What are two ways to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct selection is worth one point.

A. Populate the date dimension table by using a dataflow.
B. Populate the date dimension table by using a Stored procedure activity in a pipeline.
C. Populate the date dimension view by using T-SQL.
D. Populate the date dimension table by using a Copy activity in a pipeline.

Correct Answer: A, B
Section:
Explanation:
Both a dataflow (A) and a Stored procedure activity in a pipeline (B) are capable of creating and populating a date dimension table. A dataflow can perform the transformation needed to create the date dimension, and it aligns with the preference for using low-code tools for data ingestion when possible. A Stored procedure could be written to generate the necessary date dimension data and executed within a pipeline, which also adheres to the technical requirements for the PoC.

QUESTION 15
HOTSPOT

You need to design a semantic model for the customer satisfaction report.
Which data source authentication method and mode should you use? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

Hot Area:

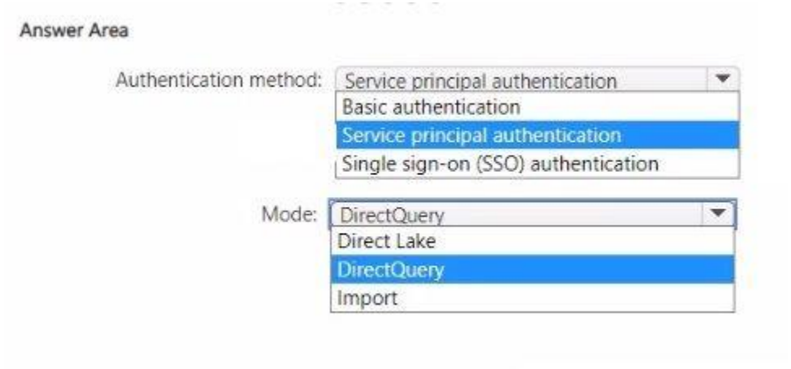

Answer Area:

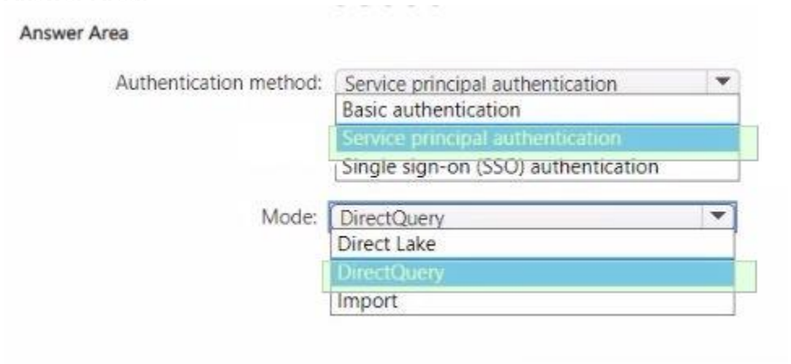

Section:
Explanation:

QUESTION 16
You have a Fabric warehouse that contains a table named Staging.Sales. Staging.Sales contains the following columns.

Name	Data type	Nullable
ProductID	Integer	No
ProductName	Varchar(30)	No
SalesDate	Datetime2(6)	No
WholesalePrice	Decimal(18, 2)	Yes
Amount	Decimal(18, 2)	Yes

You need to write a T-SQL query that will return data for the year 2023 that displays ProductID and ProductName arxl has a summarized Amount that is higher than 10,000. Which query should you use?
A)

```
SELECT ProductID, ProductName, SUM(Amount) AS TotalAmount
FROM Staging.Sales
WHERE DATEPART(YEAR,SaleDate) = '2023'
GROUP BY ProductID, ProductName
HAVING SUM(Amount) > 10000
```

B)

```
SELECT ProductID, ProductName, SUM(Amount) AS TotalAmount
FROM Staging.Sales
GROUP BY ProductID, ProductName
HAVING DATEPART(YEAR,SaleDate) = '2023' AND SUM(Amount) > 10000
```

C)

```
SELECT ProductID, ProductName, SUM(Amount) AS TotalAmount
FROM Staging.Sales
WHERE DATEPART(YEAR,SaleDate) = '2023' AND SUM(Amount) > 10000
```

D)

```
SELECT ProductID, ProductName, SUM(Amount) AS TotalAmount
FROM Staging.Sales
WHERE DATEPART(YEAR,SaleDate) = '2023'
GROUP BY ProductID, ProductName
HAVING TotalAmount > 10000
```

A. Option A
B. Option B
C. Option C
D. Option D

Correct Answer: B
Section:
Explanation:
The correct query to use in order to return data for the year 2023 that displays ProductID, ProductName, and has a summarized Amount greater than 10,000 is Option B. The reason is that it uses the GROUP BY clause to organize the data by ProductID and ProductName and then filters the result using the HAVING clause to only include groups where the sum of Amount is greater than 10,000. Additionally, the DATEPART(YEAR, SaleDate) = '2023' part of the HAVING clause ensures that only records from the year 2023 are included. Reference = For more information, please visit the official documentation on T-SQL queries and the GROUP BY clause at T-SQL GROUP BY.

QUESTION 17
DRAG DROP
You have a Fabric tenant that contains a semantic model. The model contains data about retail stores.
You need to write a DAX query that will be executed by using the XMLA endpoint The query must return a table of stores that have opened since December 1,2023.
How should you complete the DAX expression? To answer, drag the appropriate values to the correct targets. Each value 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:

Correct Answer:

Section:
Explanation:
DAX FILTER Function
DAX SUMMARIZE Function

QUESTION 18
You have a Fabric workspace named Workspace 1 that contains a dataflow named Dataflow1. Dataflow! has a query that returns 2.000 rows. You view the query in Power Query as shown in the following exhibit.

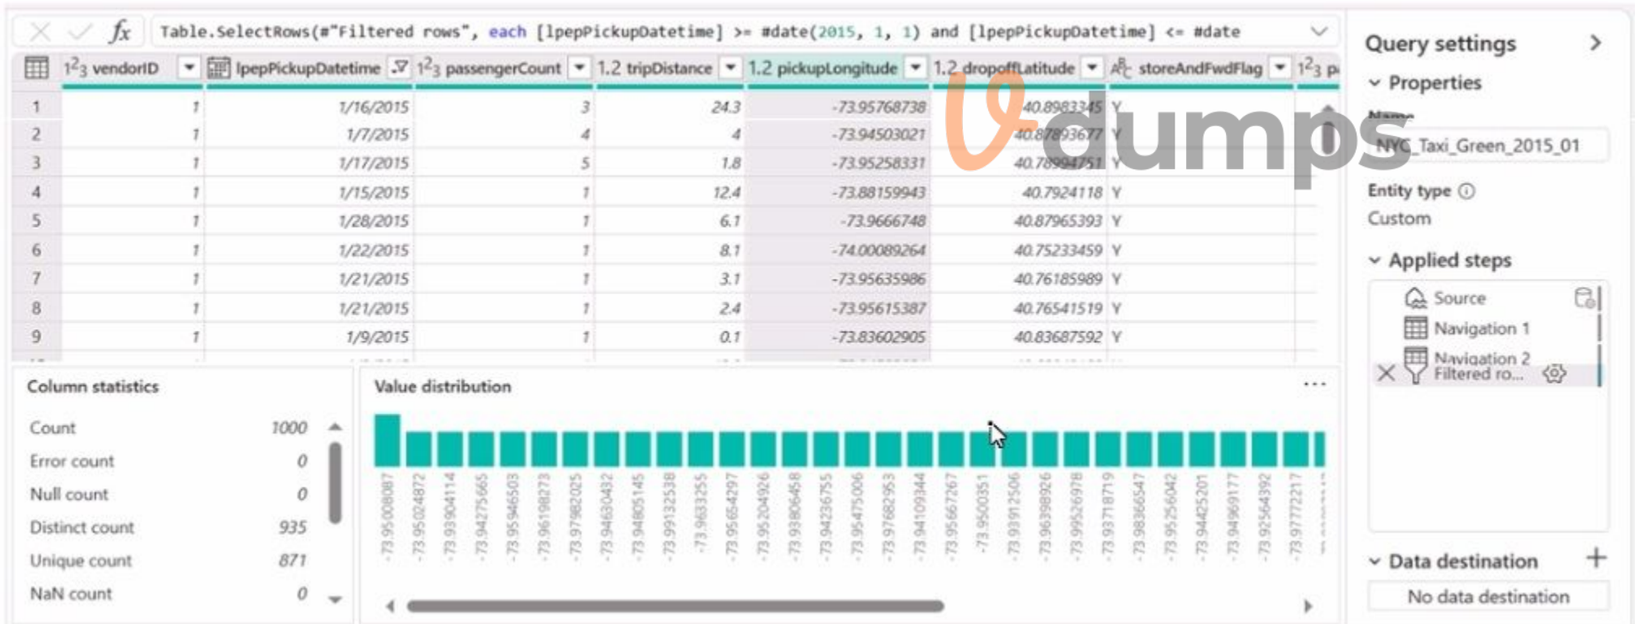

What can you identify about the pickupLongitude column?

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: B

Section:
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. Reference = Microsoft Power BI documentation on data profiling could provide further insights into understanding and interpreting column statistics like these.

QUESTION 19
You have a Fabric tenant named Tenant1 that contains a workspace named WS1. WS1 uses a capacity named C1 and contains a dawset named DS1. You need to ensure read-write access to DS1 is available by using the XMLA endpoint. What should be modified first?

A. the DS1 settings
B. the WS1 settings
C. the C1 settings
D. the Tenant1 settings

Correct Answer: C
Section:
Explanation:
To ensure read-write access to DS1 is available by using the XMLA endpoint, the C1 settings (which refer to the capacity settings) should be modified first. XMLA endpoint configuration is a capacity feature, not specific to individual datasets or workspaces. Reference = The configuration of XMLA endpoints in Power BI capacities is detailed in the Power BI documentation on dataset management.

QUESTION 20
You have a Fabric tenant that contains 30 CSV files in OneLake. The files are updated daily.
You create a Microsoft Power Bl 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: E
Section:
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. Reference = The Power BI documentation on incremental refresh and query folding provides detailed information on this topic.

QUESTION 21
You have a Fabric tenant that uses a Microsoft tower Bl Premium capacity. You need to enable scale-out for a semantic model. What should you do first?

A. At the semantic model level, set Large dataset storage format to Off.
B. At the tenant level, set Create and use Metrics to Enabled.
C. At the semantic model level, set Large dataset storage format to On.
D. At the tenant level, set Data Activator to Enabled.

Correct Answer: C
Section:

Explanation:
To enable scale-out for a semantic model, you should first set Large dataset storage format to On (C) at the semantic model level. This configuration is necessary to handle larger datasets effectively in a scaled-out environment. Reference = Guidance on configuring large dataset storage formats for scale-out is available in the Power BI documentation.

QUESTION 22
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
Section:
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. Reference = The Power BI documentation on DirectQuery provides detailed explanations of how RLS and DAX queries function in this mode.

QUESTION 23
You have a Fabric tenant that contains a complex semantic model. The model is based on a star schema and contains many tables, including a fact table named Sales. You need to create a diagram of the model. The diagram must contain only the Sales table and related tables. What should you use from Microsoft Power Bl Desktop?

A. data categories
B. Data view
C. Model view
D. DAX query view

Correct Answer: C
Section:
Explanation:
To create a diagram that contains only the Sales table and related tables, you should use the Model view (C) in Microsoft Power BI Desktop. This view allows you to visualize and manage the relationships between tables within your semantic model. Reference = Microsoft Power BI Desktop documentation outlines the functionalities available in Model view for managing semantic models.

QUESTION 24
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_COLUMN-iN_SEGMeNTS dynamic management view (DMV).
D. Query the discover_hehory6Rant dynamic management view (DMV).

Correct Answer: B, C
Section:

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. Reference = The Power BI documentation on Vertipaq Analyzer and DMV queries offers detailed guidance on how to use these tools for performance analysis.

QUESTION 25
You have a Fabric tenant that contains a semantic model named Model1. Model1 uses Import mode. Model1 contains a table named Orders. Orders has 100 million rows and the following fields.

Name	Data type	Description
OrderId	Integer	Column imported from the source
OrderDateTime	Date/time	Column imported from the source
Quantity	Integer	Column imported from the source
Price	Decimal	Column imported from the source
TotalSalesAmount	Decimal	Calculated column that multiplies Quantity and Price
TotalQuantity	Integer	Measure

You need to reduce the memory used by Model! and the time it takes to refresh the model. Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct answer is worth one point.

A. Split OrderDateTime into separate date and time columns.
B. Replace TotalQuantity with a calculated column.
C. Convert Quantity into the Text data type.
D. Replace TotalSalesAmount with a measure.

Correct Answer: A, D
Section:
Explanation:
To reduce memory usage and refresh time, splitting the OrderDateTime into separate date and time columns (A) can help optimize the model because date/time data types can be more memory-intensive than separate date and time columns. Moreover, replacing TotalSalesAmount with a measure (D) instead of a calculated column ensures that the calculation is performed at query time, which can reduce the size of the model as the value is not stored but calculated on the fly. Reference = The best practices for optimizing Power BI models are detailed in the Power BI documentation, which recommends using measures for calculations that don't need to be stored and adjusting data types to improve performance.

QUESTION 26
You have a Fabric tenant that contains a semantic model.
You need to prevent report creators from populating visuals by using implicit measures.
What are two tools that you can use to achieve the goal? Each correct answer presents a complete solution.
NOTE: Each correct answer is worth one point.

A. Microsoft Power BI Desktop
B. Tabular Editor
C. Microsoft SQL Server Management Studio (SSMS)
D. DAX Studio

Correct Answer: A, B
Section:
Explanation:
Microsoft Power BI Desktop (A) and Tabular Editor (B) are the tools you can use to prevent report creators from using implicit measures. In Power BI Desktop, you can define explicit measures which can be used in visuals. Tabular Editor allows for advanced model editing, where you can enforce the use of explicit measures. Reference = Guidance on using explicit measures and preventing implicit measures in reports can be found in the Power BI and Tabular Editor official documentation.

QUESTION 27
You have a Fabric tenant that contains a lakehouse named lakehouse1. Lakehouse1 contains a table named Table1.
You are creating a new data pipeline.
You plan to copy external data to Table1. The schema of the external data changes regularly.
You need the copy operation to meet the following requirements:
* Replace Table1 with the schema of the external data.
* Replace all the data in Table1 with the rows in the external data.
You add a Copy data activity to the pipeline. What should you do for the Copy data activity?

A. From the Source tab, add additional columns.
B. From the Destination tab, set Table action to Overwrite.
C. From the Settings tab, select Enable staging
D. From the Source tab, select Enable partition discovery
E. From the Source tab, select Recursively

Correct Answer: B
Section:
Explanation:
For the Copy data activity, from the Destination tab, setting Table action to Overwrite (B) will ensure that Table1 is replaced with the schema and rows of the external data, meeting the requirements of replacing both the schema and data of the destination table. Reference = Information about Copy data activity and table actions in Azure Data Factory, which can be applied to data pipelines in Fabric, is available in the Azure Data Factory documentation.

QUESTION 28
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: B, D
Section:
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. Reference = 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 29
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
Section:
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. Reference = The process for converting and optimizing file formats within a lakehouse is discussed in the lakehouse management documentation.

QUESTION 30
You have a Fabric tenant that contains a lakehouse named lakehouse1. Lakehouse1 contains an unpartitioned table named Table1.
You plan to copy data to Table1 and partition the table based on a date column in the source data.
You create a Copy activity to copy the data to Table1.
You need to specify the partition column in the Destination settings of the Copy activity.
What should you do first?

A. From the Destination tab, set Mode to Append.
B. From the Destination tab, select the partition column,
C. From the Source tab, select Enable partition discovery
D. From the Destination tab, set Mode to Overwrite.

Correct Answer: B
Section:
Explanation:
Before specifying the partition column in the Destination settings of the Copy activity, you should set Mode to Append (A). This will allow the Copy activity to add data to the table while taking the partition column into account. Reference = The configuration options for Copy activities and partitioning in Azure Data Factory, which are applicable to Fabric dataflows, are outlined in the official Azure Data Factory documentation.

QUESTION 31
You have a semantic model named Model 1. Model 1 contains five tables that all use Import mode. Model1 contains a dynamic row-level security (RLS) role named HR. The HR role filters employee data so that HR managers only see the data of the department to which they are assigned.
You publish Model1 to a Fabric tenant and configure RLS role membership. You share the model and related reports to users.
An HR manager reports that the data they see in a report is incomplete.
What should you do to validate the data seen by the HR Manager?

A. Ask the HR manager to open the report in Microsoft Power Bl Desktop.
B. Select Test as role to view the data as the HR role.
C. Select Test as role to view the report as the HR manager,
D. Filter the data in the report to match the intended logic of the filter for the HR department.

Correct Answer: B
Section:
Explanation:
To validate the data seen by the HR manager, you should use the 'Test as role' feature in Power BI service. This allows you to see the data exactly as it would appear for the HR role, considering the dynamic RLS setup. Here is how you would proceed:
Navigate to the Power BI service and locate Model1.
Access the dataset settings for Model1.
Find the security/RLS settings where you configured the roles.
Use the 'Test as role' feature to simulate the report viewing experience as the HR role.

Review the data and the filters applied to ensure that the RLS is functioning correctly.
If discrepancies are found, adjust the RLS expressions or the role membership as needed.

QUESTION 32
You have an Azure Repos Git repository named Repo1 and a Fabric-enabled Microsoft Power Bl Premium capacity. The capacity contains two workspaces named Workspace! and Workspace2. Git integration is enabled at the workspace level.
You plan to use Microsoft Power Bl Desktop and Workspace! to make version-controlled changes to a semantic model stored in Repo1. The changes will be built and deployed lo Workspace2 by using Azure Pipelines.
You need to ensure that report and semantic model definitions are saved as individual text files in a folder hierarchy. The solution must minimize development and maintenance effort.
In which file format should you save the changes?

A. PBIP
B. PBIT
C. PBIX
D. PBIDS

Correct Answer: C
Section:
Explanation:
When working with Power BI Desktop and Git integration for version control, report and semantic model definitions should be saved in the PBIX format. PBIX is the Power BI Desktop file format that contains definitions for reports, data models, and queries, and it can be easily saved and tracked in a version-controlled environment. The solution should minimize development and maintenance effort, and saving in PBIX format allows for the easiest transition from development to deployment, especially when using Azure Pipelines for CI/CD (continuous integration/continuous deployment) practices.

QUESTION 33
You have a Microsoft Fabric tenant that contains a dataflow.
You are exploring a new semantic model.
From Power Query, you need to view column information as shown in the following exhibit.

Which three Data view options should you select? Each correct answer presents part of the solution. NOTE: Each correct answer is worth one point.

A. Enable column profile
B. Show column quality details
C. Show column profile in details pane
D. Enable details pane
E. Show column value distribution

Correct Answer: A, B, E
Section:
Explanation:
To view column information like the one shown in the exhibit in Power Query, you need to select the options that enable profiling and display quality and distribution details. These are: A. Enable column profile - This option turns on profiling for each column, showing statistics such as distinct and unique values. B. Show column quality details - It displays the column quality bar on top of each column showing the percentage of valid, error, and

empty values. E. Show column value distribution - It enables the histogram display of value distribution for each column, which visualizes how often each value occurs.

QUESTION 34
HOTSPOT
You have a Fabric warehouse that contains a table named Sales.Products. Sales.Products contains the following columns.

Name	Data type	Nullable
ProductID	Integer	No
ProductName	Varchar(30)	No
ListPrice	Decimal(18, 2)	No
WholesalePrice	Decimal(18, 2)	Yes
AgentPrice	Decimal(18, 2)	Yes

You need to write a T-SQL query that will return the following columns.

Name	Description
ProductID	Return the ProductID value
HighestSellingPrice	Returns the highest value from ListPrice, WholesalePrice, and AgentPrice
TradePrice	Returns the AgentPrice value if present, otherwise returns the WholesalePrice value if present, otherwise returns the ListPrice value

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

Hot Area:

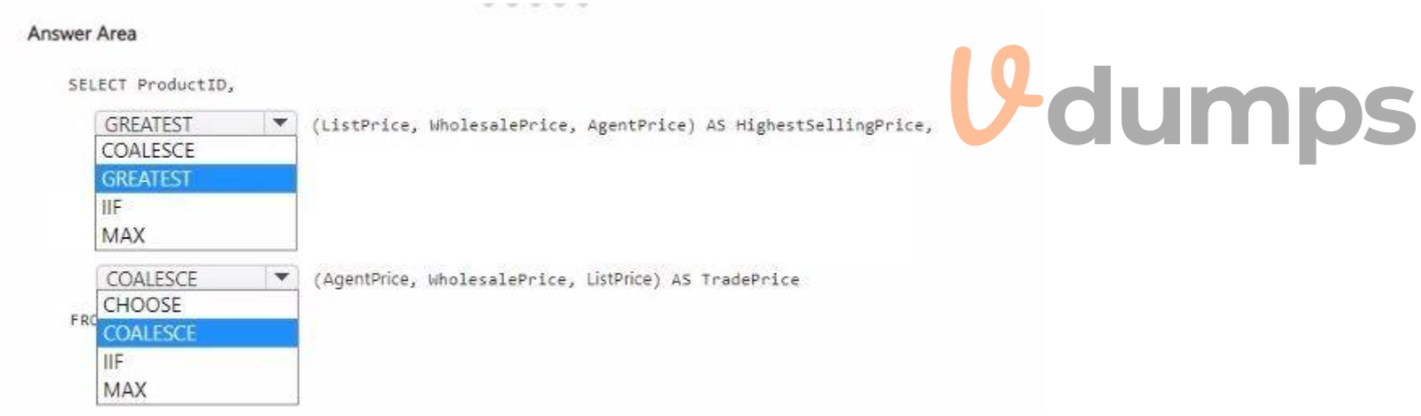

Answer Area:

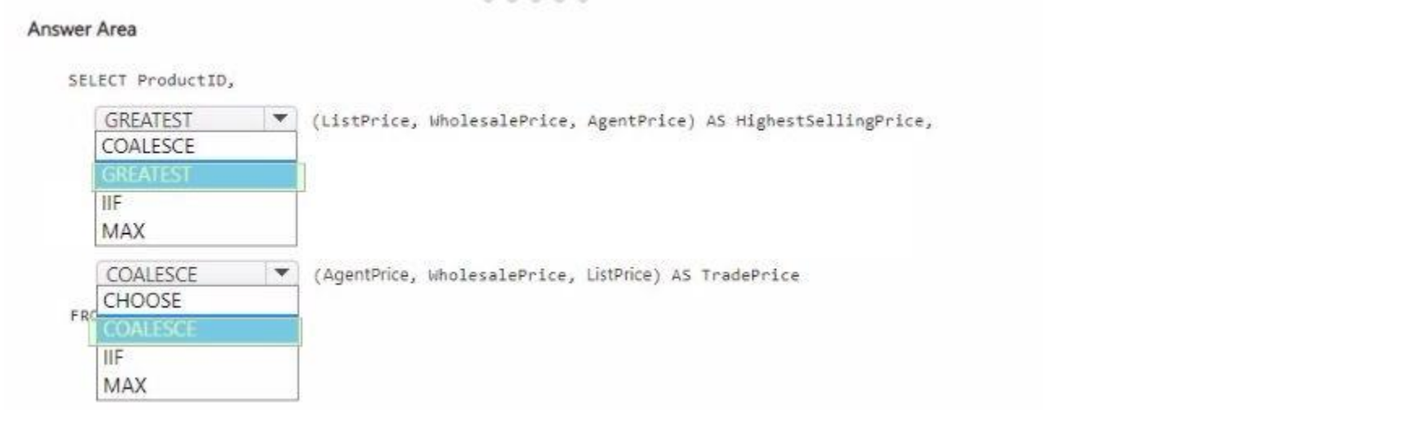

Section:
Explanation:

QUESTION 35
You have a Fabric tenant that contains a Microsoft Power Bl report.
You are exploring a new semantic model.
You need to display the following column statistics:
* Count
* Average
* Null count
* Distinct count
* Standard deviation
Which Power Query function should you run?

A. Table. FuzzyGroup
B. Table.Profile
C. Table.View
D. Table.Schema

Correct Answer: B
Section:
Explanation:
The Table.Profile function in Power Query is used to generate column statistics such as count, average, null count, distinct count, and standard deviation. You can use this function as follows:
Invoke the Power Query Editor.
Apply the Table.Profile function to your table.
The result will be a table where each row represents a column from the original table, and each column in the result represents a different statistic such as those listed in the requirement.

QUESTION 36
You have a Fabric tenant that contains customer churn data stored as Parquet files in OneLake. The data contains details about customer demographics and product usage.
You create a Fabric notebook to read the data into a Spark DataFrame. You then create column charts in the notebook that show the distribution of retained customers as compared to lost customers based on geography, the number of products purchased, age. and customer tenure.
Which type of analytics are you performing?

A. prescriptive
B. diagnostic
C. descriptive
D. predictive

Correct Answer: C
Section:

QUESTION 37
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 tenant that contains a semantic model named Model1.
You discover that the following query performs slowly against Model1.

```
1  EVALUATE
2     FILTER (
3       VALUES ( Customer[Customer Name] ),
4       CALCULATE ( COUNTROWS ( 'Order Item' ) ) > 0
5     )
6     ORDER BY Customer[Customer Name]
```

You need to reduce the execution time of the query.
Solution: You replace line 4 by using the following code:

```
NOT ( CALCULATE ( COUNTROWS ( 'Order Item' ) ) < 0 )
```

Does this meet the goal?

A. Yes
B. No

Correct Answer: B
Section: