

Microsoft.DP-600.v2026-10-12.q87

Exam Code:	DP-600
Exam Name:	Implementing Analytics Solutions Using Microsoft Fabric
Certification Provider:	Microsoft
Free Question Number:	87
Version:	v2026-10-12
# of views:	106
# of Questions views:	870
https://www.freepdf dumps.com/Microsoft.DP-600.v2026-10-12.q87.html	

NEW QUESTION: 1

You need to ensure that Contoso can use version control to meet the data analytics requirements and the general requirements. What should you do?

- A. Store all the semantic models and reports in Data Lake Gen2 storage.
- B. Modify the settings of the Research workspaces to use a GitHub repository.
- C. Store all the semantic models and reports in Microsoft OneDrive.
- D. Modify the settings of the Research division workspaces to use an Azure Repos repository.

Answer: D (LEAVE A REPLY)

You need to ensure that Contoso can use version control to meet the data analytics requirements and the general requirements.

Requirement: "All the semantic models and reports for the Research division must use version control that supports branching." Fabric supports Git integration (Azure Repos or GitHub) for semantic models and reports.

Storing in Data Lake Gen2 or OneDrive does not provide version control with branching.

The correct action is to integrate workspaces with a Git repository.

answer: D. Modify the settings of the Research division workspaces to use an Azure Repos repository.

NEW QUESTION: 2

Which syntax should you use in a notebook to access the Research division data for Productline1?

- A. `spark.sql("SELECT * FROM Lakehouse1.ResearchProduct ")`
- B. `spark.read.format("delta").load("Files/ResearchProduct ")`
- C. `spark.sql("SELECT * FROM Lakehouse1.productline1.ResearchProduct ")`
- D. `external_table('Tables/ResearchProduct)`

Answer: A (LEAVE A REPLY)

Comprehensive Detailed Explanation

The question asks: Which syntax should you use in a Fabric notebook to access the Research division data for Productline1?

Key Background from the Case

In Productline1ws, a lakehouse named Lakehouse1 is created.

In Lakehouse1, a shortcut is created to storage1, named ResearchProduct.

Storage1 contains the Research division data for Productline1 in Delta format.

Requirement: All data in lakehouses must be presented as managed tables in Lakehouse explorer.

Analyzing the Syntax Options

Option A:

```
spark.sql( " SELECT * FROM Lakehouse1.ResearchProduct " )
```

This syntax directly queries the ResearchProduct shortcut within Lakehouse1 using Spark SQL.

Since the shortcut points to Delta data, Spark can directly query it.

This is the correct way to retrieve Productline1 data from Lakehouse1.

Option B:

```
spark.sql( " SELECT * FROM Lakehouse1.productline1.ResearchProduct " )
```

This introduces an extra schema-like path (productline1) that is not part of the shortcut name.

Incorrect, because the shortcut was created as ResearchProduct inside Lakehouse1, not under another schema.

Option C:

```
external_table( ' Tables/ResearchProduct ' )
```

external_table is not the correct way to access a Lakehouse shortcut.

Shortcuts in Lakehouses appear as tables and can be queried using Spark SQL directly.

Option D:

```
spark.sql( " SELECT * FROM Lakehouse1.productline1.ResearchProduct " )
```

Same issue as Option B, includes a schema path that does not exist.

Correct Choice

Since the shortcut to ResearchProduct was created inside Lakehouse1, and Spark SQL can query it directly, the correct syntax is:

```
spark.sql( " SELECT * FROM Lakehouse1.ResearchProduct " )
```

That matches Option A.

References

Microsoft Fabric Lakehouse - Shortcuts

Query data in a lakehouse using Spark SQL

Delta format support in Microsoft Fabric

NEW QUESTION: 3

You have a Fabric semantic model named Modell that contains a table named Sales. Sales contains one million rows and uses Import storage mode. Model1 includes 50 DAX formulas that reference the Sales table and two DAX measures named Sales Quantity and Sales Quantity LY and 15 other measures.

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
Changing the storage mode of the Sales table to Dual will improve the performance of the formulas.	☐	☐
Creating a variable in the Sales Quantity LY measure to store the result of the Sales Quantity measure will improve the overall performance of Model1.	☐	☐
Using the DIVIDE() DAX function in the measures will improve the overall performance of Model1.	☐	☐

Answer:

Answer Area

Statements	Yes	No
Changing the storage mode of the Sales table to Dual will improve the performance of the formulas.	☐	☒
Creating a variable in the Sales Quantity LY measure to store the result of the Sales Quantity measure will improve the overall performance of Model1.	☐	☒
Using the DIVIDE() DAX function in the measures will improve the overall performance of Model1.	☐	☒

Explanation:

Answer Area

Statements	Yes	No
Changing the storage mode of the Sales table to Dual will improve the performance of the formulas.	☐	☒
Creating a variable in the Sales Quantity LY measure to store the result of the Sales Quantity measure will improve the overall performance of Model1.	☐	☒
Using the DIVIDE() DAX function in the measures will improve the overall performance of Model1.	☐	☒

NEW QUESTION: 4

You have a Fabric tenant that contains a lakehouse named Lakehouse1

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

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

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

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

What should you do? To answer, drag the appropriate actions to the correct requirements. Each action may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Actions	Answer Area
Set the autoCompact table setting.	Remove the files: Combine the files:
Set the optimizeWrite table setting.	
Run the VACUUM command on a schedule.	
Set the autoCompact SparkSession setting.	
Run the OPTIMIZE command on a schedule.	
Set the parallelDelete SparkSession setting.	

Answer:

Actions	Answer Area
Set the autoCompact table setting.	Remove the files: Combine the files:
Set the optimizeWrite table setting.	
Run the VACUUM command on a schedule.	
Set the autoCompact SparkSession setting.	
Run the OPTIMIZE command on a schedule.	
Set the parallelDelete SparkSession setting.	

Explanation:

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

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

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

NEW QUESTION: 5

You have a Fabric tenant that contains a semantic model named Model1.

Model1 contains a table named TemperatureUnit. TemperatureUnit contains a column named Unit. Unit contains manually added options for Celsius and Fahrenheit.

Model1 contains the following measures:

- A measure that contains the following expression.
[TemperatureMetric] = AVERAGE(Weather[Tempt])
- A dynamic measure that contains the following expression.
[TemperatureDisplay] =
VAR Unit = SELECTEDVALUE(TemperatureUnit[Unit], "Celsius")
VAR CelsiusValue = [TemperatureMetric]
RETURN
SWITCH(
Unit,
"Fahrenheit", (CelsiusValue * 9/5) + 32,
CelsiusValue
)

You need to add a dynamic format string to the [TemperatureDisplay] measure to enable users to switch display values between metric (Celsius) and imperial (Fahrenheit).

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

NOTE: Each correct selection is worth one point.

Answer Area

TempFormatted =

SWITCH(

CALCULATE

LOOKUPVALUE

SELECTEDVALUE

TemperatureMetric[Unit]

Temperature[Unit]

Unit[TemperatureUnit]

, "Celsius"),

Answer:

Answer Area

TempFormatted =

SWITCH(

CALCULATE

LOOKUPVALUE

SELECTEDVALUE

TemperatureMetric[Unit]

Temperature[Unit]

Unit[TemperatureUnit]

, "Celsius"),

)

Explanation:

Answer Area

TempFormatted =

SWITCH(
 SELECTEDVALUE (Temperature[Unit] , "Celsius"),
 "Celsius", "0.0 °C",
 "Fahrenheit", "0.0 °F",
 "0.0"
)

NEW QUESTION: 6

You have a Fabric tenant that contains three users named User1, User2, and User3. The tenant contains a security group named Group 1. User1 and User3 are members of Group " 1. The tenant contains the workspaces shown in the following table.

Name	Admin
Workspace1	User1
Workspace2	User2

The tenant contains the domains shown in the following table.

Name	Admin
Domain1	User1
Domain2	User2

User1 creates a new workspace named Workspace3.
You assign Domain! as the default domain of Group1.
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

Microsoft

Statements

User2 is assigned the Contributor role for Workspace3.

User3 is assigned the Viewer role for Workspace3.

User3 is assigned the Contributor role for Workspace1.

Yes

No

Answer:

Answer Area

Statements	Yes	No
User2 is assigned the Contributor role for Workspace3.	☐	☒
User3 is assigned the Viewer role for Workspace3.	☐	☒
User3 is assigned the Contributor role for Workspace1.	☐	☒

Explanation:

Statements	Yes	No
User2 is assigned the Contributor role for Workspace3.	☐	☒
User3 is assigned the Viewer role for Workspace3.	☐	☒
User3 is assigned the Contributor role for Workspace1.	☐	☒

NEW QUESTION: 7

You have a Microsoft Power BI semantic model that contains a measure named TotalSalesAmount.

TotalSalesAmount returns a sales revenue amount that is translated into a selected currency.

You need to ensure that the value returned by TotalSalesAmount is formatted to use the correct currency symbol.

What should you include in the solution?

- A. A dynamic format string
- B. The WINDOW DAX function
- C. A linguistic schema
- D. A field parameter

Answer: ([SHOW ANSWER](#))

Scenario:

Semantic model contains a measure: TotalSalesAmount.

This measure returns sales revenue translated into a selected currency.

Requirement: Format the result using the correct currency symbol dynamically.

Analysis:

Dynamic format string in DAX is specifically designed for applying formats (like currency symbols) dynamically based on slicers or calculation context.

WINDOW function # used for windowing calculations, not formatting.

Linguistic schema # used for Q&A natural language, not for formatting.

Field parameter # used to swap fields/measures dynamically, not for formatting.

NEW QUESTION: 8

You have a Fabric tenant that contains a workspace named Workspace1 and a user named User1. Workspace1 contains a warehouse named DW1.

You share DW1 with User1 and assign User1 the default permissions for DW1.

What can User1 do?

- A.** Build reports by using the default dataset.
- B.** Read the underlying Parquet files from OneLake.
- C.** Connect to DW1 via the TDS (Tabular Data Stream) endpoint.
- D.** Read data from the tables in DW1.

Answer: A (LEAVE A REPLY)

Comprehensive Detailed Explanation

Step 1: Default permissions when sharing a warehouse in Fabric

When you share a Fabric warehouse (DW1) with a user and assign them default permissions, the user:

Gets Build permission by default.

Build permission allows the user to:

Use the default dataset automatically created for the warehouse.

Build reports and dashboards against that dataset.

It does not grant direct query permissions to the underlying tables or files.

Step 2: Analyze the options

A). Build reports by using the default dataset. #

Correct. Sharing a warehouse gives Build permission on the warehouse's default dataset.

User1 can create Power BI reports but cannot directly query the warehouse via SQL.

B). Read the underlying Parquet files from OneLake.

Incorrect. Access to underlying OneLake files requires Direct Lake storage access or shortcut permissions, not provided by default sharing.

C). Connect to DW1 via the TDS endpoint.

Incorrect. The TDS endpoint is supported by semantic models, not warehouses.

D). Read data from the tables in DW1.

Incorrect. Default permissions do not allow table-level SQL access. They allow using the dataset for reporting.

Step 3: Correct Answer

The only capability User1 has by default is:

A). Build reports by using the default dataset.

References

Fabric warehouses - permissions

Build permission in Power BI

NEW QUESTION: 9

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.

Answer: A,E (LEAVE A REPLY)

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).

NEW QUESTION: 10

You have Fabric tenant that contains a semantic model named Mode1. Model! contains the following tables.

Name	Description
FactPackageEvents	Contains two columns named PackageID and CarrierID
DimPackage	Contains a column named PackageID
DimCarrier	Contains a column named CarrierID

You are designing a Microsoft Power BI report named Report1 that will deliver near-real-time (NRT) shipping data. Report1 will contain data about many shipped packages. You need to ensure that each table is in the appropriate storage mode to support Report1. The solution must support near-realtime data delivery while maintaining high query performance.

Which storage mode should you select for each table? To answer, drag the appropriate storage modes to the correct tables. Each mode 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.

Storage modes

DirectQuery

Dual

Hybrid

Import

Answer Area

FactPackageEvents:

DimPackage:

DimCarrier:

Answer:

Storage modes

DirectQuery

Dual

Hybrid

Import

Answer Area

FactPackageEvents:

Hybrid

DimPackage:

Dual

DimCarrier:

Dual

Explanation:

Storage modes

☐ DirectQuery

☐ Dual

☐ Hybrid

☐ Import

Answer Area

FactPackageEvents:

DimPackage:

DimCarrier:

NEW QUESTION: 11

You have a Fabric tenant that contains a semantic model named Modell. Model! contains a fact table with millions of rows of shipment data.

You need to enable partitioning to improve the query performance and manageability of Modell. Which two tools can you use to achieve the goal? Each correct answer presents a complete solution. NOTE: Each correct answer is worth one point.

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

Answer: A,B ([LEAVE A REPLY](#))

NEW QUESTION: 12

You have a Microsoft Power BI project that contains a file named definition.pbir. definition.pbir contains the following JSON.

```
{
  "version": "1.0",
  "datasetReference": {
    "byPath": {
      "path": "Sales.Dataset"
    }
    "byConnection": null
  }
}
```

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
definition.pbir is in the PBIR-Legacy format.	☐	☐
The semantic model referenced by definition.pbir is located in the Power BI service.	☐	☐
When the related report is opened, Power BI Desktop will open the semantic model in full edit mode.	☐	☐

Answer:

Answer Area

Statements	Yes	No
definition.pbir is in the PBIR-Legacy format.	☐	☒
The semantic model referenced by definition.pbir is located in the Power BI service.	☐	☒
When the related report is opened, Power BI Desktop will open the semantic model in full edit mode.	☒	☐

Explanation:

Answer Area

Statements	Yes	No
definition.pbir is in the PBIR-Legacy format.	☐	☒
The semantic model referenced by definition.pbir is located in the Power BI service.	☐	☒
When the related report is opened, Power BI Desktop will open the semantic model in full edit mode.	☒	☐

We are analyzing the JSON for definition.pbir in a Power BI Project:

```
{
  "version": "1.0",
  "datasetReference": {
    "byPath": {
      "path": "../Sales.Dataset"
    },
    "byConnection": null
  }
}
```

Statement 1:

" definition.pbir is in the PBIR-Legacy format. "

No # The legacy PBIR format references datasets by byConnection (pointing to the Power BI service).

Here, it uses byPath, which is the new project format, not legacy.

Statement 2:

" The semantic model referenced by definition.pbir is located in the Power BI service. " No # The JSON shows byPath: " ../Sales.Dataset " , meaning the semantic model is referenced locally within the project folder (Sales.Dataset), not in the Power BI service.

Statement 3:

" When the related report is opened, Power BI Desktop will open the semantic model in full edit mode. " Yes # Since the semantic model is referenced by path (local PBIP project files), Power BI Desktop will open the model in full edit mode. If it had been byConnection, it would open in live connect mode to the service instead.

Final Answer:

PBIR-Legacy format # No

Semantic model in Power BI service # No

Opens in full edit mode # Yes

References:

Power BI Project (PBIP) structure

PBIR formats: byPath vs byConnection

NEW QUESTION: 13

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. Script
- B. Copy data
- C. Switch
- D. Append variable

Answer: A (LEAVE A REPLY)

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.

References: The functionality and use of Lookup activity within Azure Data Factory is documented in Microsoft's official documentation for Azure Data Factory, under the section for pipeline activities.

NEW QUESTION: 14

You have a Fabric tenant that contains a data warehouse.

You need to load rows into a large Type 2 slowly changing dimension (SCD). The solution must minimize resource usage. Which T-SQL statement should you use?

- A. update and insert
- B. MERGE
- C. TRUNCATE TABLE and INSERT
- D. CREATE TABLE AS SELECT

Answer: ([SHOW ANSWER](#))

For Type 2 Slowly Changing Dimensions (SCD):

MERGE is the standard T-SQL approach, allowing conditional UPDATE for existing rows and INSERT for new history records, in a single statement.

UPDATE + INSERT would require multiple steps and higher resource usage.

TRUNCATE + INSERT would destroy historical data, violating SCD rules.

CREATE TABLE AS SELECT is not suitable for incremental Type 2 handling.

Correct answer: B .

Reference: Implement SCD Type 2 with MERGE

NEW QUESTION: 15

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

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

- * Support the use of dataflows to load and append data to the data store.
- * Ensure that Delta tables are V-Order optimized and compacted automatically.

Which type of data store should you use?

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

Answer: A,C ([LEAVE A REPLY](#))

A lakehouse (A) is the type of data store you should use. It supports dataflows to load and append data and ensures that Delta tables are Z-Order optimized and compacted automatically.

References = The capabilities of a lakehouse and its support for Delta tables are described in the lakehouse and Delta table documentation.

NEW QUESTION: 16

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]),  
        AVERAGEA('Question'[Question Text]),  
        AVERAGEX(VALUES('Survey'[Customer Key]),  
            [NumberOfMonths],  
            [LastCurrentDate],  
            [NumberOfMonths],  
            [Period],  
            'Survey Question'[Question Title] = "Overall Satisfaction"  
    )  
RETURN  
    Result
```



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]),  
        AVERAGEA('Question'[Question Text]),  
        AVERAGEX(VALUES('Survey'[Customer Key]),  
            [NumberOfMonths],  
            [LastCurrentDate],  
            [NumberOfMonths],  
            [Period],  
            'Survey Question'[Question Title] = "Overall Satisfaction"  
    )  
RETURN  
    Result
```



Microsoft

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 (  
    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
```

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.

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
```

Valid DP-600 Dumps shared by Actual4test.com for Helping Passing DP-600 Exam! Actual4test.com now offer the **newest DP-600 exam dumps**, the Actual4test.com DP-600 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-600 dumps with Test Engine here:

https://www.actual4test.com/DP-600_examcollection.html (212 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 17

You have a Fabric tenant that contains two workspaces named Workspace1 and Workspace2 and a user named User1.

You need to ensure that User1 can perform the following tasks:

Create a new domain.

Create two subdomains named subdomain1 and subdomain2.

Assign Workspace1 to subdomain1.

Assign Workspace2 to subdomain2.

The solution must follow the principle of least privilege.

Which role should you assign to User1?

- A.** Domain admin
- B.** Domain contributor
- C.** Workspace Admin
- D.** Fabric admin

Answer: (SHOW ANSWER)

User1 must be able to:

Create a new domain.

Create two subdomains.

Assign Workspace1 to subdomain1.

Assign Workspace2 to subdomain2.

Key Role Definitions in Fabric:

Domain admin: Can manage domain settings, add/remove workspaces, manage contributors.

Domain contributor: Can contribute to an existing domain but cannot create a domain.

Workspace admin: Only controls permissions/settings inside a workspace, not domains.

Fabric admin: Tenant-wide admin, broader than required.

Requirement: Follow least privilege.

To create a new domain # Requires Domain admin.

To assign workspaces to subdomains # Domain admin rights are sufficient.

Fabric admin is excessive and breaks least privilege.

NEW QUESTION: 18

You have a Fabric tenant that contains a warehouse.

A user discovers that a report that usually takes two minutes to render has been running for 45 minutes and has still not rendered.

You need to identify what is preventing the report query from completing.

Which dynamic management view (DMV) should you use?

- A.** sys.dm-exec_requests
- B.** sys.dm_pdw_exec_requests

C. sys.dm._exec._connections

D. sys.dm._exec._sessions

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 19

You have a Fabric tenant that contains a semantic model named Model1. Model1 contains a fact table named FactSales. FactSales has a relationship to a dimension table named DimDate on a column named OrderDate.

To FactSales, you add a second relationship between FactSales and DimDate based on a column named ShippedDate.

You need to filter by using ShippedDate in a new report. The solution must NOT affect the function of existing reports.

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

NOTE: Each correct answer is worth one point.

A. Disable the ShippedDate relationship between the tables.

B. Use a relatedtable DAX function.

C. Enable the ShippedDate relationship.

D. Use a userrelationship DAX function.

E. Use a crossfilter DAX function.

Answer: D,E ([LEAVE A REPLY](#))

NEW QUESTION: 20

You have two Microsoft Power BI queries named Employee and Retired Roles.

You need to merge the Employee query with the Retired Roles query. The solution must ensure that rows in the Employee query that match the Retired Roles query are removed.

Which column and Join Kind should you use in Power Query Editor? To answer, select the appropriate options in the answer area.

NOTE: Each correct answer is worth one point

Answer Area

Merge

Select a table and matching columns to create a merged table.

Employee

ID	Division	Title Detail	Role	Department	Department Number	Division Number
198	Services	Marketing	Vice President	Marketing	22	14
299	Services	Business Development	Vice President	Sales	17	14
269	Services	Human Resources	Vice President	Human Resources	23	14
246	Services	HSE & Risk Management	Vice President	Safety	16	14

Retired Roles

Division	Role
Services	Technical Analyst
Services	Technical Support Analyst
Services	Senior Engineer
Services	Chief Technical Officer
Services	Applied Technology Specialist

Join Kind

Left Outer (all from first, matching from second)

Left Outer (all from first, matching from second)

Right Outer (all from second, matching from first)

Full Outer (all rows from both)

Inner (only matching rows)

Left Anti (rows only in first)

Right Anti (rows only in second)

to continue

OK

Cancel



Answer:

Select a table and matching columns to create a merged table.

Employee

ID	Division	Title Detail	Role	Department	Department Number	Division Number
298	Services	Marketing	Vice President	Marketing	22	14
299	Services	Business Development	Vice President	Sales	17	14
169	Services	Human Resources	Vice President	Human Resources	13	14
146	Services	HSE & Risk Management	Vice President	Safety	16	14

Retired Roles

Division	2	Role	1
Services		Technical Analyst	
Services		Technical Support Analyst	
Services		Senior Engineer	
Services		Chief Technical Officer	
Services		Applied Technology Specialist	

Join Kind

Left Outer (all from first, matching from second)

- Left Outer (all from first, matching from second)
- Right Outer (all from second, matching from first)
- Full Outer (all rows from both)
- Inner (only matching rows)
- Left Anti (rows only in first)
- Right Anti (rows only in second)

to continue

OK

Cancel

Explanation:

Answer Area

Merge

Select a table and matching columns to create a merged table.

Employee

ID	Division	Title Detail	Role	Department	Department Number	Division Number
198	Services	Marketing	Vice President	Marketing	22	14
299	Services	Business Development	Vice President	Sales	17	14
169	Services	Human Resources	Vice President	Human Resources	13	14
146	Services	HSE & Risk Management	Vice President	Safety	16	14

Retired Roles

Division	Role
Services	Technical Analyst
Services	Technical Support Analyst
Services	Senior Engineer
Services	Chief Technical Officer
Services	Applied Technology Specialist

Join Kind

- Left Outer (all from first, matching from second)
- Left Outer (all from first, matching from second)
- Right Outer (all from second, matching from first)
- Full Outer (all rows from both)
- Inner (only matching rows) ☒
- Left Anti (rows only in first)
- Right Anti (rows only in second)

to continue.

OK

Cancel

NEW QUESTION: 21

You have a Fabric tenant named Tenant1 that contains a lakehouse named Lakehouse1.

You need to add data to Lakehouse1 from a CSV file in an Azure Storage account outside of Fabric. The solution must minimize development effort.

What should you use to add the data?

- A. a Dataflow Gen2
- B. a shortcut
- C. a copy job
- D. a pipeline

Answer: (SHOW ANSWER)

Comprehensive Detailed Explanation

Requirement

Tenant: Fabric tenant Tenant1

Data target: Lakehouse1

Data source: CSV file in an Azure Storage account (outside of Fabric)

Requirement: Minimize development effort

Step 1: Understand options

A). Dataflow Gen2

Used for ETL (transformations, cleansing, ingestion).

More effort because you need to build Power Query transformations.

Not minimal effort for simply exposing a file.

B). Shortcut

Fabric supports creating shortcuts to external data sources (OneLake, ADLS Gen2, Azure Blob, Dataverse, etc.).

Shortcut directly exposes external files/tables inside the lakehouse without copying or moving data.

This is the lowest-effort solution to access external CSV.

C). Copy job

Pipelines have a copy activity to physically copy data into Fabric.

Requires setup and execution. More development effort than a shortcut.

D). Pipeline

Used for orchestrating complex ETL workflows.

Overhead for this scenario.

Step 2: Correct Answer

Since the goal is simply to access a CSV in external storage with minimal development effort, the right solution is:

B). a shortcut

References

Shortcuts in Microsoft Fabric OneLake

Dataflow Gen2 overview

NEW QUESTION: 22

You have a Fabric tenant that contains a workspace named Workspace1. Workspace1 contains a single semantic model that has two Microsoft Power BI reports.

You have a Microsoft 365 subscription that contains a data loss prevention (DLP) policy named DLP1.

You need to apply DLP1 to the items in Workspace1.

What should you do?

A. Apply a master data endorsement to the semantic model.

B. Apply a certified endorsement to the semantic model.

C. Create a workspace identity.

D. Apply sensitivity labels to the semantic model and reports.

Answer: D (LEAVE A REPLY)

Step 1: Understand the scenario

Workspace1 contains a semantic model and two Power BI reports.

A Microsoft 365 Data Loss Prevention (DLP) policy (DLP1) already exists.

Requirement: Apply DLP1 to the workspace items.

In Microsoft Fabric and Power BI, DLP policies are enforced through Microsoft Purview Information Protection (MIP) sensitivity labels.

Step 2: Evaluate each option

A). Apply a master data endorsement to the semantic model.

Endorsements (promoted, certified) help users find and trust datasets but do not enforce DLP policies.

Not correct.

B). Apply a certified endorsement to the semantic model.

Certification improves discoverability and trust, but again, it does not trigger DLP policy enforcement.

Not correct.

C). Create a workspace identity.

Workspace identity (service principal) is used for connecting to external resources with managed identity.

It does not relate to DLP policy application.

Not correct.

D). Apply sensitivity labels to the semantic model and reports.

Sensitivity labels (from Microsoft Purview) integrate with Microsoft 365 DLP policies.

When sensitivity labels are applied to Power BI datasets, reports, dashboards, and workspaces, DLP1 can enforce protections such as preventing sharing or exporting sensitive data.

This is the correct answer.

Step 3: Correct Action

To apply DLP1 to Workspace1 items, you must:

Configure sensitivity labels in Microsoft Purview.

Apply those sensitivity labels to the semantic model and reports inside Workspace1.

The existing DLP1 policy will then govern data handling for those items.

References

Sensitivity labels in Power BI

Apply sensitivity labels to Power BI items

DLP policies in Microsoft Purview

NEW QUESTION: 23

You have a Fabric tenant that contains a lakehouse named Lakehouse1. Lakehouse1 contains a Delta table that has one million Parquet files.

You need to remove files that were NOT referenced by the table during the past 30 days. The solution must ensure that the transaction log remains consistent, and the ACID properties of the table are maintained. What should you do?

A. Run the optimize command and specify the v-order parameter.

B. Run the OPTIMIZE command and specify the z-order parameter.

C. Run the vacuum command.

D. From OneLake file explorer, delete the files.

Answer: C ([LEAVE A REPLY](#))

VACUUM in Delta Lake removes unreferenced files older than a retention period (default 7 days, but configurable, e.g., 30 days).

It ensures that the transaction log remains consistent and ACID properties are maintained.

OPTIMIZE with V-Order or Z-Order improves file organization and query performance, but it does not remove unused files.

Deleting files manually from OneLake would break the transaction log consistency.


Correct answer: C. Run the VACUUM command.

Reference: VACUUM in Delta Lake

NEW QUESTION: 24

You are creating a report and a semantic model in Microsoft Power BI Desktop.

The Value measure has the expression shown in the following exhibit.



Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

Answer Area

A [answer choice] was added to the Value measure.

dynamic format string
calculation group
data category
dynamic format string
synonym

The Value measure can return values formatted as [answer choice].

percentages or whole numbers
percentages only
percentages or decimals
whole numbers only
percentages or whole numbers

Answer:

Answer Area

A [answer choice] was added to the Value measure.

The Value measure can return values formatted as [answer choice].

Explanation:

Answer Area

A [answer choice] was added to the Value measure.

The Value measure can return values formatted as [answer choice].

Comprehensive Detailed Explanation

Step 1: Reviewing the provided DAX expression

```
SWITCH(  
    SELECTEDVALUE(Metric[Metric]),  
    "# of Customers ", "# ",  
    "Gross Margin % ", " 0.00% ",  
    "#,### "  
)
```

This is applied as a dynamic format string.

It changes the display formatting of the Value measure depending on the selected metric.

Step 2: Identify the feature

The screenshot shows Format = Dynamic in the Power BI Modeling tab.

This confirms that a dynamic format string has been applied.

So, for the first blank # dynamic format string.

Step 3: Identify the output formats

If the metric = # of Customers # format is " # " # whole numbers.

If the metric = Gross Margin % # format is " 0.00% " # percentages.

Fallback # " #,### " # also whole numbers.

Thus, the measure supports both whole numbers and percentages.

So, for the second blank # percentages or whole numbers.

References

Dynamic format strings in Power BI

SWITCH function in DAX

NEW QUESTION: 25

You have a Fabric tenant that contains a semantic model named model1. The two largest columns in model1 are shown in the following table.

Name	Cardinality	Compressed size	Description
TransactionKey	2,500,350,000	160 GB	Contains an incrementing surrogate key for a fact table named SalesTransaction
SaleDateTime	6,850,273	120 GB	Contains the datetime (to the nearest second) of when a sale occurred

You need to optimize model 1. The solution must meet the following requirements:

- * Reduce the model size.
- * Increase refresh performance when using Import mode.
- * Ensure that the datetime value for each sales transaction is available in the model.

What should you do on each column? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

TransactionKey:

SaleDateTime:

Answer:



Explanation:



We are optimizing model1 in Fabric with two very large columns:

TransactionKey

Cardinality: 2.5 billion

Size: 160 GB

Description: surrogate key for SalesTransaction fact table.

SaleDateTime

Cardinality: 6.8 million

Size: 120 GB

Description: datetime (to the second) of when a sale occurred.

Goal:

Reduce model size.

Improve refresh performance (Import mode).

Ensure datetime values remain available.

Analysis per column

1. TransactionKey (very high cardinality)

Surrogate keys with extremely high cardinality do not add analytical value in a semantic model because they are not used in measures or grouping.

They take a lot of memory and reduce refresh performance.

Best practice: remove the column from the model.

2. SaleDateTime

Needed to analyze transactions by date and time.

However, storing the full datetime at second-level cardinality causes huge memory usage.

To optimize: split the column into Date and Time columns.

Reduces cardinality significantly.

Preserves full datetime information (can recombine if needed).

Improves compression and refresh speed.

Final Answer:

TransactionKey: Remove the column.

SaleDateTime: Split the column.

References:

Microsoft Fabric optimization guidance - reduce cardinality

Best practices for datetime optimization in semantic models

Answer Selection:

TransactionKey # Remove the column

SaleDateTime # Split the column

NEW QUESTION: 26

You have a Fabric tenant that contains a workspace named Workspace1.

You plan to deploy a semantic model named Model1 by using the XMLA endpoint.

You need to optimize the deployment of Model1. The solution must minimize how long it takes to deploy Model1.

What should you do in Workspace1?

- A. Select Small semantic model storage format.
- B. Select Users can edit data models in the Power BI service.
- C. Set Enable Cache for Shortcuts to On.
- D. Select Large semantic model storage format.

Answer: ([SHOW ANSWER](#))

When using XMLA endpoint for deployment or processing, the semantic model storage format impacts deployment times.

Small storage format allows faster metadata operations and shorter deployment times.

Large semantic model storage format is only needed for extremely large models (> 10 GB dataset limit), and it increases deployment overhead.

The other options (B: service editing, C: cache for shortcuts) are unrelated to XMLA endpoint deployment performance.

Reference:

Power BI semantic model storage format

NEW QUESTION: 27

You have a Fabric tenant that contains a lakehouse named Lakehouse1. Lakehouse1 contains a table named Nyctaxi_raw. Nyctaxi_raw contains the following columns.

Name	Data type
pickupDateTime	Timestamp
passengerCount	Integer
fareAmount	Double
paymentType	String
tipAmount	Double

You create a Fabric notebook and attach it to lakehouse1.

You need to use PySpark code to transform the data. The solution must meet the following requirements:

- * Add a column named pickupDate that will contain only the date portion of pickupDateTime.
- * Filter the DataFrame to include only rows where fareAmount is a positive number that is less than 100.

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

```
df = spark.read.format("delta").load("Tables/nyctaxi_raw")
df2 = df.withColumn("pickupDate", df['tpepPickupDateTime'].cast('date')).filter("fareAmount > 0 AND fareAmount < 100")
```

df.withColumn
df.columns
df.select
df.withColumn
df.withColumnsRenamed

.filter("fareAmount > 0 AND fareAmount < 100")
.filter("fareAmount > 0 AND fareAmount < 100")
.filter(col("fareAmount").contains("1..100"))
.when(df.fareAmount > 0 AND fareAmount < 100)
.where(df.fareAmount.isin([1,100]))

.cast('date')
.alias('date')
.cast('date')
.cast('pickupDate')
.getField('date')

Answer:

Answer Area

```
df = spark.read.format("delta").load("Tables/nyctaxi_raw")
df2 = df.withColumn("pickupDate", df['tpepPickupDateTime'].cast('date')).filter("fareAmount > 0 AND fareAmount < 100")
```

df.withColumn
df.columns
df.select
df.withColumn
df.withColumnsRenamed

.filter("fareAmount > 0 AND fareAmount < 100")
.filter("fareAmount > 0 AND fareAmount < 100")
.filter(col("fareAmount").contains("1..100"))
.when(df.fareAmount > 0 AND fareAmount < 100)
.where(df.fareAmount.isin([1,100]))

.cast('date')
.alias('date')
.cast('date')
.cast('pickupDate')
.getField('date')

Explanation:

Answer Area

```
df = spark.read.format("delta").load("Tables/nyctaxi_raw")
df2 = df.withColumn("pickupDate", df['tpepPickupDateTime'].cast('date')).filter("fareAmount > 0 AND fareAmount < 100")
```

df.withColumn ("pickupDate", df['tpepPickupDateTime'].cast('date'))
.filter("fareAmount > 0 AND fareAmount < 100")

Microsoft

Add the pickupDate column: `.withColumn("pickupDate", df["tpepPickupDateTime"].cast("date"))` Filter the DataFrame: `.filter("fareAmount > 0 AND fareAmount < 100")` In PySpark, you can add a new column to a DataFrame using the `.withColumn` method, where the first argument is the new column name and the second argument is the expression to generate the content of the new column. Here, we use the `.cast("date")` function to extract only the date part from a timestamp. To filter the DataFrame, you use the `.filter` method with a condition that selects rows where fareAmount is greater than 0 and less than 100, thus ensuring only positive values less than 100 are included.

NEW QUESTION: 28

You have a Fabric warehouse named Warehouse1 that contains a table named TaxiTrips. TaxiTrips contains the following columns.

Name	Data type	Description
Company	Varchar(100)	Name of each taxi company.
passengerCount	Integer	Number of passengers per trip.
tripDistance	Decimal(10,4)	Distance traveled in miles during a trip.
pickupDate	Datetime2	Date and time a passenger was picked up.
FareAmount	Decimal(10,4)	Cost of each trip.

You need to create a query that shows the top three taxi companies based on the total miles traveled.

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

Microsoft

SELECT

FETCH NEXT 3 ROWS ONLY

Limit (3)

Top

company AS TaxiCompany,

sum(passengerCount) AS TotalPassengers,

sum(tripDistance) TotalDistance

FROM [Warehouse1].[dbo].[TaxiTrips]

GROUP BY company

ORDER BY tripDistance

DESC;

GROUP BY company

GROUP BY sum(tripDistance)

ORDER BY tripDistance

DESC;

Answer:

SELECT

FETCH NEXT 3 ROWS ONLY

Limit (3)

Top

company AS TaxiCompany,

sum(passengerCount) AS TotalPassengers,

sum(tripDistance) TotalDistance

FROM [Warehouse1].[dbo].[TaxiTrips]

GROUP BY company

ORDER BY tripDistance

DESC;

GROUP BY company

GROUP BY sum(tripDistance)

ORDER BY tripDistance

DESC;

Microsoft

Explanation:

Answer Area

Microsoft

SELECT

Top (3)

company AS TaxiCompany,

sum(passengerCount) AS TotalPassengers,

sum(tripDistance) TotalDistance

FROM [Warehouse1].[dbo].[TaxiTrips]

GROUP BY company

ORDER BY sum(tripDistance)

DESC;

freepdfdumps.com

NEW QUESTION: 29

You have a Fabric tenant that contains a Microsoft Power BI report named Report 1.

Report1 is slow to render. You suspect that an inefficient DAX query is being executed.

You need to identify the slowest DAX query, and then review how long the query spends in the formula engine as compared to the storage engine.

Which five actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

View the Server Timings tab.

From Performance analyzer, capture a recording.

Enable Query Timings and Server Timings. Run the query.

View the Query Timings tab.

Sort the Duration (ms) column in descending order by DAX query time.

Copy the first query to DAX Studio.

Answer Area



Answer:

Actions

View the Server Timings tab.

From Performance analyzer, capture a recording.

Enable Query Timings and Server Timings. Run the query.

View the Query Timings tab.

Sort the Duration (ms) column in descending order by DAX query time.

Copy the first query to DAX Studio.

Answer Area

From Performance analyzer, capture a recording.

View the Server Timings tab.

Enable Query Timings and Server Timings. Run the query.

View the Query Timings tab.

Sort the Duration (ms) column in descending order by DAX query time.

Explanation:

To identify the slowest DAX query and analyze the time it spends in the formula engine compared to the storage engine, you should perform the following actions in sequence:

- * From Performance analyzer, capture a recording.
- * View the Server Timings tab.
- * Enable Query Timings and Server Timings. Run the query.
- * View the Query Timings tab.
- * Sort the Duration (ms) column in descending order by DAX query time.

NEW QUESTION: 30

You are analyzing the data in a Fabric notebook.

You have a Spark DataFrame assigned to a variable named df.

You need to use the Chart view in the notebook to explore the data manually.

Which function should you run to make the data available in the Chart view?

A. displayHTML

B. show

C. write

D. display

Answer: D ([LEAVE A REPLY](#))

The display function is the correct choice to make the data available in the Chart view within a Fabric notebook. This function is used to visualize Spark DataFrames in various formats including charts and graphs directly within the notebook environment. References = Further explanation of the display function can be found in the official documentation on Azure Synapse Analytics notebooks.

NEW QUESTION: 31

You have a query in Power Query Editor that contains two columns named Order_Date and Shipping_Date.

You need to create a column that will calculate the number of days between Order_Date and Shipping_Date for each row.

Which Power Query function should you use?

A. Duration.From

B. Date.AddDays

C. Duration.Days

D. DateTime.LocalNow

Answer: ([SHOW ANSWER](#)**)**

Comprehensive Detailed Explanation

Requirement

We have two columns in Power Query Editor:

Order_Date

Shipping_Date

We need to calculate the number of days between the two dates.

Step 1: How date differences work in Power Query (M language)

Subtracting two date or datetime values in Power Query returns a duration value.

Example:

[Shipping_Date] - [Order_Date]

This produces a duration such as 3.00:00:00 (3 days).

Step 2: Converting duration into days

To get the number of days as an integer, we use Duration.Days().

Example:

Duration.Days([Shipping_Date] - [Order_Date])

This returns 3 if the shipping date is 3 days after the order date.

Step 3: Evaluate other options

A). Duration.From # Converts a value to a duration type but does not extract the number of days.

B). Date.AddDays # Adds days to a given date, not used for calculating the difference.

- C). Duration.Days # Correct, extracts total days from a duration.
- D). DateTime.LocalNow # Returns the current system datetime, unrelated to this requirement.

References

Duration.Days function (Power Query M)

Working with dates in Power Query

Valid DP-600 Dumps shared by Actual4test.com for Helping Passing DP-600 Exam! Actual4test.com now offer the **newest DP-600 exam dumps**, the Actual4test.com DP-600 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-600 dumps with Test Engine here:

https://www.actual4test.com/DP-600_examcollection.html (212 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 32

You have a Fabric tenant that contains a lakehouse.

You are using a Fabric notebook to save a large DataFrame by using the following code.

```
df.write.partitionBy("year","month","day").mode("overwrite").parquet("Files/SalesOrder")
```

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
The results will form a hierarchy of folders for each partition key.	☐	☐
The resulting file partitions can be read in parallel across multiple nodes.	☐	☐
The resulting file partitions will use file compression.	☐	☐

Answer:

Answer Area

Statements	Yes	No
The results will form a hierarchy of folders for each partition key.	☒	☐
The resulting file partitions can be read in parallel across multiple nodes.	☒	☐
The resulting file partitions will use file compression.	☐	☒

Explanation:

The results will form a hierarchy of folders for each partition key. - Yes The resulting file partitions can be read in parallel across multiple nodes. - Yes The resulting file partitions will use file compression. - No Partitioning data by columns such as year, month, and day, as shown in the DataFrame write operation, organizes the output into a directory hierarchy that reflects the partitioning structure. This organization can improve the performance of read operations, as queries that filter by the partitioned columns can scan only the relevant directories.

Moreover, partitioning facilitates parallelism because each partition can be processed independently across different nodes in a distributed system like Spark. However, the code snippet provided does not explicitly specify that file compression should be used, so we cannot assume that the output will be compressed without additional context.

References =

DataFrame write partitionBy

Apache Spark optimization with partitioning

NEW QUESTION: 33

You are designing a Fabric semantic model named Model1 for a financial analytics workspace. Model1 includes a DAX measure named Measure1 that returns total sales.

You have a table named Currency that contains a column named Format. The Format column defines currency formatting strings for each record as follows: \$8,880.00, ta,880.00, or yv.sve.ee.

You have a report that uses the Currency column to identify which symbol and format string to display. In the report Measure1 must appear with two decimal places and the appropriate currency symbol based on the selected value in the Currency column.

You create the following DAX formula.

FORMAT = SELECTEDVALUE(' Currency ' [Format])

You need to ensure that Measure1 appears correctly in visuals. The solution must minimize logic duplication.

What should you do?

- A. Copy the DAX expression from the Format column and profile the expression into the format string expression of Measure.
- B. For Model1, set the data category of Measure to Country/Region.
- C. For Model1, set the format type of Measure1 to Currency.
- D. Create a dynamic format string for Measure1 that references the FORMAT measure.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 34

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

Answer: A,B,D ([LEAVE A REPLY](#))

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.

References: These features and their descriptions are typically found in the Power Query documentation, under the section for data profiling and quality features.

NEW QUESTION: 35

You have a Fabric workspace named Workspace1 and a user named User1. You need to ensure that User1 can perform the following tasks in Workspace1:

- * Create dashboards.
- * Publish semantic models.
- * Create and publish reports.

The solution must follow the principle of least privilege. Which workspace role should you assign to User1?

- A.** Member
- B.** Admin
- C.** Viewer
- D.** Contributor

Answer: ([SHOW ANSWER](#))

NEW QUESTION: 36

You have a Fabric tenant that contains a semantic model named Model1. Model1 contains a fact table named FactSales. FactSales has a relationship to a dimension table named DimDate on a column named OrderDate.

To FactSales, you add a second relationship between FactSales and DimDate based on a column named ShippedDate.

You need to filter by using ShippedDate in a new report. The solution must NOT affect the function of existing reports.

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

NOTE: Each correct answer is worth one point.

- A.** Enable the ShippedDate relationship.
- B.** Use a crossfilter DAX function.
- C.** Disable the ShippedDate relationship between the tables.
- D.** Use a userrelationship DAX function.
- E.** Use a relatedtable DAX function.

Answer: B,D ([LEAVE A REPLY](#))

NEW QUESTION: 37

You have a Fabric workspace named Workspace1 and an Azure SQL database.

You plan to create a dataflow that will read data from the database, and then transform the data by performing an inner join. You need to ignore spaces in the values when performing the inner join. The solution must minimize development effort. What should you do?

- A.** Append the queries by using fuzzy matching.
- B.** Merge the queries by using fuzzy matching.
- C.** Append the queries by using a lookup table.
- D.** Merge the queries by using a lookup table.

Answer: B (LEAVE A REPLY)

You need to perform an inner join and ignore spaces in the values.

In Power Query, to handle joins where values may contain spaces or minor mismatches, you use fuzzy matching in the Merge Queries operation.

Append queries combines rows, not columns, so it does not apply here.

Lookup table is not required since the transformation can be handled natively in Power Query with fuzzy matching.

Correct solution: Merge the queries by using fuzzy matching.

Reference: Fuzzy merge in Power Query

NEW QUESTION: 38

You have a Fabric warehouse that contains a table named SalesOrderDetail. SalesOrderDetail contains three columns named OrderQty, ProductID and SalesOrderID. SalesOrderDetail contains one row per combination of SalesOrderID and ProductID.

You need to calculate the proportion of the total quantity of each sales order represented by each product within the sales order.

Which T-SQL statement should you run?

```
SELECT SalesOrderID, ProductID, OrderQty
      ,CAST(1. * OrderQty / SUM(OrderQty) OVER(ORDER BY ProductID) * 100
      AS DECIMAL(5,2)) AS PercentbyProductID
FROM Sales.SalesOrderDetail;
```

A.

```
SELECT SalesOrderID, ProductID, OrderQty
      ,CAST(1. * OrderQty / SUM(OrderQty) OVER(PARTITION BY ProductID) * 100
      AS DECIMAL(5,2)) AS PercentbyProductID
FROM Sales.SalesOrderDetail;
```

B.

```
SELECT SalesOrderID, ProductID, OrderQty
      ,CAST(1. * OrderQty / SUM(OrderQty) OVER(PARTITION BY SalesOrderID) * 100
      AS DECIMAL(5,2)) AS PercentbyProductID
FROM Sales.SalesOrderDetail;
```

C.

```
SELECT SalesOrderID, ProductID, OrderQty
      ,CAST(1. * OrderQty / SUM(OrderQty) OVER(ORDER BY SalesOrderID) * 100
      AS DECIMAL(5,2)) AS PercentbyProductID
FROM Sales.SalesOrderDetail;
```

D.

Answer: C (LEAVE A REPLY)

Comprehensive Detailed Explanation

We need to calculate the proportion of the total quantity of each sales order represented by each product within the sales order.

Step 1: Analyze the requirement

Table: SalesOrderDetail

Columns: OrderQty, ProductID, SalesOrderID

Each row = one product in a sales order.

Requirement: For each SalesOrderID, calculate what percentage each product contributes to the total order quantity.

This means we must:

Calculate the total order quantity per SalesOrderID.

Divide each product's OrderQty by that total.

Express it as a percentage.

Step 2: Evaluate the options

A). Uses OVER(ORDER BY ProductID)

Wrong: ORDER BY does not segment data, it only defines sequence. We need grouping by sales order.

B). Uses OVER(PARTITION BY ProductID)

Wrong: This would calculate percentages per product across all sales orders, not per sales order.

C). Uses OVER(PARTITION BY SalesOrderID)

Correct: This computes the total OrderQty per SalesOrderID and divides each product's quantity by that total.

Exactly what the requirement asks.

D). Uses OVER(ORDER BY SalesOrderID)

Wrong: Again, ORDER BY just sequences rows, does not group them.

Step 3: Correct Query

```
SELECT
SalesOrderID,
ProductID,
OrderQty,
CAST(1. * OrderQty / SUM(OrderQty)
OVER(PARTITION BY SalesOrderID) * 100 AS DECIMAL(5,2))
AS PercentByProductID
FROM Sales.SalesOrderDetail;
```

Why Option C is Correct

PARTITION BY SalesOrderID ensures the denominator is the total quantity for the current order.

This produces the correct percentage breakdown of each product inside its sales order.

References

T-SQL Window Functions

Aggregate Functions with PARTITION BY

Microsoft Fabric Warehouse T-SQL Support

NEW QUESTION: 39

You have a Fabric tenant that contains a semantic model. The model contains 15 tables.

You need to programmatically change each column that ends in the word Key to meet the following requirements:

- * Hide the column.
- * Set Nullable to False.
- * Set Summarize By to None
- * Set Available in MDX to False.
- * Mark the column as a key column.

What should you use?

- A.** Microsoft Power BI Desktop
- B.** Tabular Editor
- C.** ALM Toolkit
- D.** DAX Studio

Answer: B ([LEAVE A REPLY](#))

Tabular Editor is an advanced tool for editing Tabular models outside of Power BI Desktop that allows you to script out changes and apply them across multiple columns or tables. To accomplish the task programmatically, you would:

- * Open the model in Tabular Editor.
- * Create an Advanced Script using C# to iterate over all tables and their respective columns.
- * Within the script, check if the column name ends with 'Key'.
- * For columns that meet the condition, set the properties accordingly: IsHidden = true, IsNullable = false, SummarizeBy = None, IsAvailableInMDX = false.
- * Additionally, mark the column as a key column.
- * Save the changes and deploy them back to the Fabric tenant.

References: The ability to batch-edit properties using scripts in Tabular Editor is well-documented in the tool's official documentation and user community resources.

NEW QUESTION: 40

You have an Azure SQL database named DB1 and a Fabric workspace named Workspace1. Workspace 1 contains a lakehouse named LH1 and a Dataflow Gen2 named Dataflow1. Dataflow1 includes a query named Query1 that loads data from DB1, applies transformations to the data, and then filters the data.

You discover that Query1 loads all the data before applying the transformations.

You need to ensure that Query1 uses query folding.

What should you do?

- A.** Remove unused columns from Query1.
- B.** Enable incremental refresh for Query1.
- C.** Move the filter steps before the transformation steps in Query1.
- D.** Replace DB1 with a Microsoft SQL Server database.

Answer: C ([LEAVE A REPLY](#))

NEW QUESTION: 41

You have a Fabric tenant that contains a warehouse.

Several times a day, the performance of all warehouse queries degrades. You suspect that Fabric is throttling the compute used by the warehouse.

What should you use to identify whether throttling is occurring?

- A.** the Capacity settings
- B.** the Monitoring hub
- C.** dynamic management views (DMVs)
- D.** the Microsoft Fabric Capacity Metrics app

Answer: (SHOW ANSWER)

To identify whether throttling is occurring, you should use the Monitoring hub (B). This provides a centralized place where you can monitor and manage the health, performance, and reliability of your data estate, and see if the compute resources are being throttled. References = The use of the Monitoring hub for performance management and troubleshooting is detailed in the Azure Synapse Analytics documentation.

NEW QUESTION: 42

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

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

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

Solution: You use the following PySpark expression:

df.summary()
Does this meet the goal?
A. Yes
B. No

Answer: A (LEAVE A REPLY)

Yes, the df.summary() method does meet the goal. This method is used to compute specified statistics for numeric and string columns. By default, it provides statistics such as count, mean, stddev, min, and max.
References = The PySpark API documentation details the summary() function and the statistics it provides.

NEW QUESTION: 43

You have a Fabric tenant that contains a data warehouse named DW1. DW1 contains a table named DimCustomer. DimCustomer contains the fields shown in the following table.

ID	CustomerKey	CustomerAltKey	Title	FirstName	LastName	AddressLine1	City
1	29546	christopher1@contoso.com	Mr.	Christopher	Beck	93-2501, Blackfriars Road	London
2	29587	john12@contoso.com	Mr.	John	Brooks	99-3 Forest Works	London
3	29612	richard2@contoso.com	Mr.	Richard	Byham	Knaves Beech Ind	High Wycombe

You need to identify duplicate email addresses in DimCustomer. The solution must return a maximum of 1,000 records.
Which four T-SQL statements should you run in sequence? To answer, move the appropriate statements from the list of statements to the answer area and arrange them in the correct order.

Statements

Microsoft

freepdfdumps.com

LIMIT 1000

SELECT TOP(1000) CustomerAltKey, COUNT(*)

WHERE COUNT(*) > 1

SELECT CustomerAltKey, COUNT(*)

FROM DimCustomer

GROUP BY CustomerAltKey

HAVING COUNT(*) > 1

Answer Area

Answer:

Statements

LIMIT 1000

SELECT TOP(1000) CustomerAltKey, COUNT(*)

WHERE COUNT(*) > 1

SELECT CustomerAltKey, COUNT(*)

FROM DimCustomer

GROUP BY CustomerAltKey

HAVING COUNT(*) > 1

Answer Area

SELECT CustomerAltKey, COUNT(*)

FROM DimCustomer

GROUP BY CustomerAltKey

HAVING COUNT(*) > 1

Explanation:

Statements

LIMIT 1000

SELECT TOP(1000) CustomerAltKey, COUNT(*)

WHERE COUNT(*) > 1

Answer Area

1 SELECT CustomerAltKey, COUNT(*)

2 FROM DimCustomer

3 GROUP BY CustomerAltKey

4 HAVING COUNT(*) > 1

Step 1 - Select the required fields and count

We want to group by CustomerAltKey and count occurrences.

SELECT CustomerAltKey, COUNT(*)

Step 2 - From the table

FROM DimCustomer

Step 3 - Grouping by email

GROUP BY CustomerAltKey

Step 4 - Filtering duplicates (only values with count > 1)

HAVING COUNT(*) > 1

Step 5 - Limiting the result to 1,000 rows

LIMIT 1000

(or in T-SQL, SELECT TOP(1000) ...). Since both options are shown, we use LIMIT 1000 because Fabric Warehouse supports Synapse SQL T-SQL + ANSI SQL style.

Correct Sequence:

SELECT CustomerAltKey, COUNT(*)

FROM DimCustomer

GROUP BY CustomerAltKey

HAVING COUNT(*) > 1
LIMIT 1000

NEW QUESTION: 44

You are creating a dataflow in Fabric to ingest data from an Azure SQL database by using a T-SQL statement.

You need to ensure that any foldable Power Query transformation steps are processed by the Microsoft SQL Server engine.

How should you complete the code? 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.

Values

- EnableFolding
- NativeQuery
- Optimize
- Record
- StopFolding
- Table
- Value

Answer Area

```
let
    Source = Sql.Databases(
        "server.database.windows.net"
    ),
    Database = Source[[Name = "db"]][Data],
    Query = [ ] . [ ] (
        Database,
        " SELECT * FROM customer WHERE country IN ('USA', 'UK')",
        null,
        [ [ ] = true]
    )
in
    Query
```

Microsoft

Answer:

Values

- EnableFolding
- NativeQuery
- Optimize
- Record
- StopFolding
- Table
- Value

Answer Area

```
let
    Source = Sql.Databases(
        "server.database.windows.net"
    ),
    Database = Source[[Name = "db"]][Data],
    Query = [ Table ] . [ NativeQuery ] (
        Database,
        " SELECT * FROM customer WHERE country IN ('USA', 'UK')",
        null,
        [ [ EnableFolding ] = true]
    )
in
    Query
```

Microsoft

Explanation:

You should complete the code as follows:

- * Table
- * NativeQuery
- * EnableFolding

In Power Query, using Table before the SQL statement ensures that the result of the SQL query is treated as a table. NativeQuery allows a native database query to be passed through from Power Query to the source database. The EnableFolding option ensures that any subsequent transformations that can be folded will be sent back and executed at the source database (Microsoft SQL Server engine in this case).

NEW QUESTION: 45

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. full outer
- C. right anti
- D. right outer
- E. left anti
- F. inner

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 46

You have a Microsoft Power BI Premium Per User (PPU) workspace that contains a semantic model.

You have an Azure App Service app named App1 that modifies row-level security (RLS) for the model by using the XMLA endpoint. App1 requires users to sign in by using their Microsoft Entra credentials to access the XMLA endpoint. You need to configure App1 to use a service account to access the model. What should you do first?

- A. Add a managed identity to the workspace.
- B. Add a managed identity to App1.
- C. Upgrade the workspace to Premium capacity.
- D. Modify the XMLA Endpoint setting.

Answer: ([SHOW ANSWER](#))

Power BI Premium Per User (PPU) already supports XMLA read/write . You don't need to upgrade to Premium capacity.

The workspace itself does not get a managed identity; rather, App1 (Azure App Service) can be configured with a managed identity .

Once App1 has a managed identity, you grant it access to the Power BI model in the workspace via XMLA endpoint.

Modifying the XMLA endpoint setting is not sufficient; authentication must come from the service (App1).

Correct answer: B .

Reference: Service principals and managed identities for Power BI XMLA

Valid DP-600 Dumps shared by Actual4test.com for Helping Passing DP-600 Exam! Actual4test.com now offer the **newest DP-600 exam dumps**, the Actual4test.com DP-600 exam questions have been updated and answers have been corrected get the **newest** Actual4test.com DP-600 dumps with Test Engine here:

https://www.actual4test.com/DP-600_examcollection.html (212 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 47

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: following code:

```
NOT ( ISEMPTY ( CALCULATETABLE ( 'Order Item' ) ) )
```

Does this meet the goal?

A. Yes

B. No

Answer: A ([LEAVE A REPLY](#))

ISEMPTY(CALCULATETABLE(' Order Item ')) checks if there are no related rows in ' Order Item ' for a customer. Wrapping this in NOT returns true if the customer has at least one order. This avoids calculating row counts and instead just checks existence, which is more efficient for performance. This reduces execution time and meets the goal.

NEW QUESTION: 48

You are the administrator of a Fabric workspace that contains a lakehouse named Lakehouse1. Lakehouse1 contains the following tables:

* Table1: A Delta table created by using a shortcut

* Table2: An external table created by using Spark

* Table3: A managed table

You plan to connect to Lakehouse1 by using its SQL endpoint. What will you be able to do after connecting to Lakehouse1?

A. Update the data in Table1.

B. Update the data Table3.

C. ReadTable2.

D. ReadTable3.

Answer: D ([LEAVE A REPLY](#))

NEW QUESTION: 49

You have a Fabric workspace named Workspace1 that contains a lakehouse named Lakehouse1 and a user named User1. You need to ensure that User1 can read all the data in Lakehouse1. The solution must follow the principle of least privilege. Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.


A. Share Lakehouse1 by using item permissions.

- B. Select Read all Apache Spark
- C. Assign the Contributor role to User1.
- D. Assign the Viewer role to User1.
- E. Assign the Member role to User1.

Answer: A,B (LEAVE A REPLY)

NEW QUESTION: 50

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



What can you identify about the pickupLongitude column?

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

Answer: A (LEAVE A REPLY)

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

NEW QUESTION: 51

You have a Fabric tenant that contains a workspace named Workspace! Workspace1 uses the Pro license mode and contains a semantic model named Model1. You have an Azure DevOps organization. You need to enable version control for Workspace1. The solution must ensure that Model 1 is added to the repository.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
☐ Create a deployment pipeline.	
☐ Configure branch policies in Azure DevOps.	
☐ Assign Workspace1 to a Fabric capacity.	
☐ Connect Workspace1 to a Git provider.	
☐ Sync Workspace1 with the repository.	

Answer:

Actions	Answer Area
☐ Create a deployment pipeline.	☐ Assign Workspace1 to a Fabric capacity.
☐ Configure branch policies in Azure DevOps.	
☐ Assign Workspace1 to a Fabric capacity.	☐ Connect Workspace1 to a Git provider.
☐ Connect Workspace1 to a Git provider.	☐ Sync Workspace1 with the repository.
☐ Sync Workspace1 with the repository.	

Explanation:

Actions	Answer Area
☐ Create a deployment pipeline.	1 ☐ Assign Workspace1 to a Fabric capacity.
☐ Configure branch policies in Azure DevOps.	2 ☐ Connect Workspace1 to a Git provider.
	3 ☐ Sync Workspace1 with the repository.

We need to enable version control for a Fabric workspace (Workspace1) and ensure that Model1 (a semantic model) is added to an Azure DevOps repository.

Key facts:

Workspace1 is in Pro mode. To enable Git integration, the workspace must be assigned to a Fabric capacity.

Git integration steps in Fabric:

Assign the workspace to a Fabric capacity.

Connect the workspace to a Git provider (e.g., Azure DevOps or GitHub).

Sync the workspace with the repository (to push artifacts like semantic models).

Branch policies and deployment pipelines are not needed for simply enabling version control.

Correct Sequence:

Assign Workspace1 to a Fabric capacity.

Connect Workspace1 to a Git provider.

Sync Workspace1 with the repository.

Final Answer:

Step 1 # Assign Workspace1 to a Fabric capacity

Step 2 # Connect Workspace1 to a Git provider

Step 3 # Sync Workspace1 with the repository

References:

Microsoft Fabric Git integration

Enable Git in Fabric workspace

This ensures Model1 will be version-controlled in the linked DevOps repository.


NEW QUESTION: 52

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.

Answer Area



Answer:



Explanation:

Answer Area



GRANT ALTER ON SCHEMA::schemaA TO User1;

The ALTER permission allows a user to modify the schema of an object, and granting ALTER on a schema will allow the user to perform operations like TRUNCATE TABLE on any object within that schema. It is the correct permission to grant to User1 for truncating tables in schemaA.

References =

GRANT Schema Permissions

Permissions That Can Be Granted on a Schema

NEW QUESTION: 53

You create a semantic model by using Microsoft Power BI Desktop. The model contains one security role named SalesRegionManager and the following tables:

- * Sales
- * SalesRegion
- * Sales Address

You need to modify the model to ensure that users assigned the SalesRegionManager role cannot see a column named Address in Sales Address.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

⋮ Open the model in Power BI Desktop.

⋮ Set Object Level Security to **Default** for SalesRegionManager.

⋮ Set the Hidden property to **True**.

⋮ Open the model in Tabular Editor.

⋮ Select the **Address** column in SalesAddress.

⋮ Set Object Level Security to **None** for SalesRegionManager.

Answer Area



Answer:

Actions

- Open the model in Power BI Desktop.
- Set Object Level Security to **Default** for SalesRegionManager.
- Set the Hidden property to **True**.
- Open the model in Tabular Editor.
- Select the **Address** column in SalesAddress.
- Set Object Level Security to **None** for SalesRegionManager.

Answer Area

- Open the model in Tabular Editor.
- Select the **Address** column in SalesAddress.
- Set Object Level Security to **None** for SalesRegionManager.



Explanation:

Actions	Answer Area
Open the model in Power BI Desktop.	1 Open the model in Tabular Editor.
Set Object Level Security to Default for SalesRegionManager.	2 Select the Address column in SalesAddress.
Set the Hidden property to True .	3 Set Object Level Security to None for SalesRegionManager.

To ensure that users assigned the SalesRegionManager role cannot see the Address column in the SalesAddress table, follow these steps in sequence:

Open the model in Tabular Editor.

Select the Address column in SalesAddress.

Set Object Level Security to None for SalesRegionManager.

NEW QUESTION: 54

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

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

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

Solution: You use the following PySpark expression:

df.show()

Does this meet the goal?

A. Yes

B. No

Answer: B (LEAVE A REPLY)

The df.show() method also does not meet the goal. It is used to show the contents of the DataFrame, not to compute statistical functions. References = The usage of the show() function is documented in the PySpark API documentation.

NEW QUESTION: 55

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.

Answer: C,D (LEAVE A REPLY)

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

NEW QUESTION: 56

You need to migrate the Research division data for Productline2. The solution must meet the data preparation requirements. 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

```
df=spark.read.format('csv').options(header="true",inferSchema="true").load("fs://storage1.dfs.core.windows.net/files/productline2")

df.write.mode("overwrite").format(" ").save(" ")
```

delta

csv

delta

parquet

Tables/productline2

productline2

Tables/productline2

Tables/research/productline2

Answer:



Explanation:



Comprehensive Detailed Explanation

From the case study:

Research division data for Productline2 is currently in CSV format in storage2.

Requirement: "All the Research division data in the lakehouses must be presented as managed tables in Lakehouse explorer." In Fabric lakehouses, managed tables are stored in Delta format inside the Tables folder.

Step 1: Reading the source

```
df = spark.read.format( " csv " ) \
options(header= " true " , inferSchema= " true " ) \
load( " abfss://storage1.dfs.core.windows.net/files/productline2 " )
```

This correctly ingests the CSV source from ADLS Gen2.

Step 2: Writing to Lakehouse as a managed table

You must write the data in Delta format to ensure it is queryable and managed within the lakehouse.

The correct path is under Tables/, because this is where Fabric automatically manages Lakehouse managed tables.

The target table should be named productline2, so the correct path is:

```
df.write.mode( " overwrite " ).format( " delta " ).save( " Tables/productline2 " )
```

Why not other options?
CSV or Parquet formats would not create a managed Lakehouse table; they would just create files.

Writing to productline2 directly (without Tables/) would store unmanaged files in the Lakehouse Files area, not managed tables.

Writing to Tables/research/productline2 adds an unnecessary subdirectory and is not the standard structure for managed tables.

Correct Final Code

```
df.write.mode( " overwrite " ).format( " delta " ).save( " Tables/productline2 " )
```

References Managed tables in Microsoft Fabric Lakehouse Delta Lake support in Fabric Spark write to Lakehouse

NEW QUESTION: 57

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

Answer: ([SHOW ANSWER](#))

The charts in the Fabric notebook are summarizing customer churn data (retained vs lost customers) across dimensions such as geography, products purchased, age, and tenure. This is describing "what has happened" in the data.

Descriptive analytics summarizes and visualizes historical data.

Diagnostic analytics explains "why" events occurred.

Predictive analytics forecasts future outcomes.

Prescriptive analytics recommends actions.

Since only distributions are being shown, this is descriptive analytics.

References:

Microsoft Learn - Types of Analytics in Fabric

Descriptive Analytics concepts

NEW QUESTION: 58

You have a Fabric tenant that contains a warehouse named DW1 and a lakehouse named LH1. DW1 contains a table named Sales.Product. LH1 contains a table named Sales.Orders.

You plan to schedule an automated process that will create a new point-in-time (PIT) table named Sales.

ProductOrder in DW1. Sales.ProductOrder will be built by using the results of a query that will join Sales.

Product and Sales.Orders.

You need to ensure that the types of columns in Sales. ProductOrder match the column types in the source tables. The solution must minimize the number of operations required to create the new table.

Which operation should you use?

- A. CREATE TABLE AS SELECT (CTAS)
- B. INSERT INTO
- C. CREATE MATERIALIZED VIEW AS SELECT
- D. CREATE TABLE AS CLONE OF

Answer: ([SHOW ANSWER](#))

CTAS creates a new table from the results of a query, automatically inheriting the column types of the source tables. It is the simplest and most efficient way to create a point-in-time (PIT) table.

INSERT INTO requires a table to already exist.

CREATE MATERIALIZED VIEW keeps results up to date but is heavier and not required for PIT.

CLONE creates a shallow copy of an existing table, but we are joining two tables, so clone is not applicable.

Correct: A.

Reference: CTAS in Fabric data warehouse

NEW QUESTION: 59

You have Fabric tenant that contains four workspaces named Development, Test, QA, and Production. All the workspaces are in Premium Per User (PPU) license mode.

You plan to use a release pipeline to support the development lifecycle from Development to Production.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions	Answer Area
Move each workspace to Pro license mode.	
Create the QA stage.	
Create a deployment pipeline.	
Assign workspaces.	
Create a deployment rule.	

Answer:

Actions	Answer Area
Move each workspace to Pro license mode.	Create a deployment pipeline.
Create the QA stage.	Create the QA stage.
Create a deployment pipeline.	
Assign workspaces.	Assign workspaces.
Create a deployment rule.	

Explanation:

Actions	Answer Area
Create a deployment rule.	1 Create a deployment pipeline.
Move each workspace to Pro license mode.	2 Create the QA stage.
	3 Assign workspaces.

Scenario Recap

Fabric tenant with workspaces: Development, Test, QA, Production.

All workspaces are already in Premium Per User (PPU) mode (which is required for deployment pipelines).

Task: Set up a release pipeline to move through Dev # Test # QA # Production.

Step 1: Create a deployment pipeline

The first step is always to create a new deployment pipeline in the Fabric/Power BI service. This establishes the structure for Dev # Test # Prod (and optionally QA).

Step 2: Create the QA stage

By default, pipelines have 3 stages (Dev, Test, Prod).

Since we need a QA stage in addition, we must explicitly add it as a stage to the pipeline.

Step 3: Assign workspaces

Once the pipeline structure is ready (with QA included), we assign each workspace to the corresponding stage:

Development # Development

Test # Test

QA # QA

Production # Production

Why Not the Other Actions?

Move each workspace to Pro license mode # Incorrect, because deployment pipelines require Premium capacity or PPU, and the workspaces already are in PPU mode.

Create a deployment rule # Deployment rules (e.g., parameter or data source rules) are optional for environment-specific changes, but not part of the initial required setup sequence.

References

Deployment pipelines overview

Assigning workspaces to deployment pipelines

NEW QUESTION: 60

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.

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 61

You have a Microsoft Power BI semantic model that is used by your company ' s finance and sales departments for monthly reporting. The semantic model contains two explicit measures named sales Amount and Orders and a measure named KPI Title that returns text for dynamic visual titles.

You need to implement a calculation group named Time Intelligence that contains a column named Time Calculation. The solution must meet the following requirements

- * Apply time-intelligence patterns to both explicit measures by using SELECTEDMEASURE().
- * Format only the yoy% calculation item as a percentage.

What should you configure ' To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point

Answer Area

Measure expression pattern for Time Intelligence:

CALCULATE(SELECTEDMEASURE(), DATESYTD('Date'[Date]))

FORMAT(SELECTEDMEASURE(), "General Number")

SUMX(VALUES('Date'[Date]), SELECTEDMEASURE())

To format the YOY% calculation item:

Wrap the YOY% expression with FORMAT.

Set the static format string of the measure to #,##0.00%.

Set the dynamic format string of the calculation item to #,##0.00%.

Answer:

Answer Area

Measure expression pattern for Time Intelligence:

Microsoft

freepdfdumps.com

Calculate (SELECTEDMEASURE(), DATESYTD('Date'[Date]))

Format (SELECTEDMEASURE(), "General Number")

SUMX (VALUES('Date'[Date]), SELECTEDMEASURE())

To format the YOY% calculation item:

Wrap the YOY% expression with FORMAT.

Set the static format string of the measure to #,##0.00%.

Set the dynamic format string of the calculation item to #,##0.00%.

Explanation:

Answer Area

Measure expression pattern for Time Intelligence:

Microsoft

Calculate (SELECTEDMEASURE(), DATESYTD('Date'[Date]))

To format the YOY% calculation item:

Set the dynamic format string of the calculation item to #,##0.00%.

Valid DP-600 Dumps shared by Actual4test.com for Helping Passing DP-600 Exam! Actual4test.com now offer the **newest DP-600 exam dumps**, the Actual4test.com DP-600 exam questions have been updated and answers have been corrected get the **newest** Actual4test.com DP-600 dumps with Test Engine here:

https://www.actual4test.com/DP-600_examcollection.html (212 Q&As Dumps, **30%OFF Special Discount: Freepdfdumps**)

NEW QUESTION: 62

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

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

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

Answer Area

Microsoft

Statements	Yes	No
A replica of dbo.Sales is created in the test schema by copying the metadata only.	☐	☐
Additional schema changes to dbo.FactSales will also apply to test.FactSales.	☐	☐
Additional data changes to dbo.FactSales will also apply to test.FactSales.	☐	☐

Answer:

Answer Area

Statements	Yes	No
A replica of dbo.Sales is created in the test schema by copying the metadata only.	☐	☒
Additional schema changes to dbo.FactSales will also apply to test.FactSales.	☐	☒
Additional data changes to dbo.FactSales will also apply to test.FactSales.	☒	☐

Explanation:

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

References =

CREATE TABLE AS SELECT (CTAS) in SQL Data Warehouse

NEW QUESTION: 63

You have a Fabric tenant that contains a workspace named Workspace_DEV. Workspace_DEV contains the semantic models shown in the following table.

Name	Scheduled refresh policy
Model1	Configured
Model2	Not configured

Workspace_DEV contains the dataflows shown in the following table.

Name	Type	Scheduled refresh policy
DF1	Dataflow Gen1	Configured
DF2	Dataflow Gen2	Not configured

Answer Area

Statements	Yes	No
DF1 will be deployed to Workspace_TEST.	☐	☐
Data from Model1 will be deployed to Workspace_TEST.	☐	☐
The scheduled refresh policy for Model1 will be deployed to Workspace_TEST.	☐	☐

Answer:

Statements	Yes	No
DF1 will be deployed to Workspace_TEST.	☒	☐
Data from Model1 will be deployed to Workspace_TEST.	☐	☒
The scheduled refresh policy for Model1 will be deployed to Workspace_TEST.	☐	☒

Explanation:

Answer Area	
Statements	
DF1 will be deployed to Workspace_TEST.	☒ Yes ☐ No
Data from Model1 will be deployed to Workspace_TEST.	☐ Yes ☒ No
The scheduled refresh policy for Model1 will be deployed to Workspace_TEST.	☐ Yes ☒ No

DF1 will be deployed to Workspace_TEST: Yes

Data from Model1 will be deployed to Workspace_TEST: No

The scheduled refresh policy for Model1 will be deployed to Workspace_TEST: No DF1 is a Dataflow Gen1 with a configured scheduled refresh policy, so it will be deployed to Workspace_TEST.

Data from Model1 will not be deployed as Model1 's scheduled refresh policy is not configured, indicating it is not set up for deployment.

The scheduled refresh policy for Model1 will not be deployed since it is listed as " Not configured. "

NEW QUESTION: 64

You are creating a semantic model in Microsoft Power BI 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

Answer: A ([LEAVE A REPLY](#))

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. References = Microsoft's documentation on Power BI file formats and Visual Studio Code provides further clarification on the usage of PBIX files.

NEW QUESTION: 65

You have a Fabric tenant that contains a data pipeline.

You need to ensure that the pipeline runs every four hours on Mondays and Fridays.

To what should you set Repeat for the schedule?

- A. Daily
- B. By the minute
- C. Weekly
- D. Hourly

Answer: ([SHOW ANSWER](#)**)**

You should set Repeat for the schedule to Weekly (C). This allows you to specify the pipeline to run on specific days of the week, in this case, every four hours on Mondays and Fridays.

References = Scheduling options for data pipelines are available in the Azure Data Factory documentation, which includes details on configuring recurring triggers.

NEW QUESTION: 66

You have source data in a CSV file that has the following fields:

- * SalesTransactionID
- * SaleDate
- * CustomerCode
- * CustomerName
- * CustomerAddress
- * ProductCode
- * ProductName
- * Quantity
- * UnitPrice

You plan to implement a star schema for the tables in WH1. The dimension tables in WH1 will implement Type 2 slowly changing dimension (SCD) logic.

You need to design the tables that will be used for sales transaction analysis and load the source data.

Which type of target table should you specify for the CustomerName, CustomerCode, and SaleDate fields?

To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

CustomerCode:

CustomerName:

SaleDate:



The screenshot shows the 'Answer Area' with three dropdown menus. The first two, 'CustomerCode' and 'CustomerName', have 'Dimension' selected. The third, 'SaleDate', has 'Junk dimension' selected. A watermark 'freepdfdumps.com' is visible across the middle of the image.

Answer:



Explanation:



We are designing a star schema for sales transactions. In a star schema:

Dimension tables contain descriptive attributes (e.g., Customer, Product, Date).

Fact tables contain measurable, quantitative data (e.g., SalesTransactionID, Quantity, UnitPrice).

Factless fact tables capture events without numeric measures.

Junk dimensions store miscellaneous low-cardinality attributes.

The requirement says dimension tables will use Type 2 Slowly Changing Dimensions (SCD) # so Customer and Date attributes belong to Dimension tables.

Field-by-field classification

CustomerCode

A business key (natural key) for the customer dimension.

Belongs in the Customer Dimension table.

answer: Dimension.

CustomerName

A descriptive attribute of the customer.

Belongs in the Customer Dimension table.

answer: Dimension.

SaleDate

A time attribute used for analysis.

Belongs in the Date Dimension table.

answer: Dimension.

Final Answer:

CustomerCode # Dimension

CustomerName # Dimension

SaleDate # Dimension

References:

Star schema design for Power BI / Fabric

Slowly Changing Dimensions in data warehousing

This follows Kimball's star schema best practices.

NEW QUESTION: 67

You have a Fabric workspace that uses the default Spark starter pool and runtime version 1,2.

You plan to read a CSV file named Sales.raw.csv in a lakehouse, select columns, and save the data as a Delta table to the managed area of the lakehouse. Sales_raw.csv contains 12 columns.

You have the following code.

```
from pyspark.sql.functions import year

(spark
 .read
 .format("csv")
 .option("header", 'true')
 .load("Files/sales_raw.csv")
 .select('SalesOrderNumber', 'OrderDate', 'CustomerName', 'UnitPrice')
 .withColumn("Year", year("OrderDate"))
 .write
 .partitionBy('Year')
 .saveAsTable("sales")
)
```

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
The Spark engine will read only the 'SalesOrderNumber', 'OrderDate', 'CustomerName', 'UnitPrice' columns from Sales_raw.csv.	☐	☐
Removing the partition will reduce the execution time of the query.	☐	☐
Adding inferSchema='true' to the options will increase the execution time of the query.	☐	☐

Answer:

Answer Area

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

Explanation:

Answer Area

Statements

The Spark engine will read only the 'SalesOrderNumber', 'OrderDate', 'CustomerName', 'UnitPrice' columns from Sales_raw.csv. ☒ Yes ☐ No

Removing the partition will reduce the execution time of the query. ☐ Yes ☒ No

Adding inferSchema='true' to the options will increase the execution time of the query. ☒ Yes ☐ No

The Spark engine will read only the 'SalesOrderNumber', 'OrderDate', 'CustomerName', 'UnitPrice' columns from Sales_raw.csv. - Yes Removing the partition will reduce the execution time of the query. - No Adding inferSchema='true' to the options will increase the execution time of the query. - Yes The code specifies the selection of certain columns, which means only those columns will be read into the DataFrame. Partitions in Spark are a way to optimize the execution of queries by organizing the data into parts that can be processed in parallel. Removing the partition could potentially increase the execution time because Spark would no longer be able to process the data in parallel efficiently. The inferSchema option allows Spark to automatically detect the column data types, which can increase the execution time of the initial read operation because it requires Spark to read through the data to infer the schema.

NEW QUESTION: 68

You have a Fabric tenant that contains a Microsoft Power BI 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

Answer: B (LEAVE A REPLY)

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.

References: The use of Table.Profile is part of Power Query M function documentation where it explains how to gather column statistics for a given table.

NEW QUESTION: 69

You have a Fabric tenant

You plan to create a data pipeline named Pipeline1. Pipeline1 will include two activities that will execute in sequence. You need to ensure that a failure of the first activity will NOT block the second activity. Which conditional path should you configure between the first activity and the second activity?

- A. Upon Completion
- B. Upon Success
- C. Upon Skip
- D. Upon Failure

Answer: A (LEAVE A REPLY)

Upon Success: downstream runs only if the first activity succeeds.

Upon Failure: downstream runs only if the first activity fails.

Upon Skip: downstream runs if the activity was skipped.

Upon Completion: downstream runs regardless of whether the first activity succeeded or failed.

Since we want the second activity to run even if the first fails, the correct answer is Upon Completion.

Reference: Pipeline activity dependencies in Fabric Data Factory

NEW QUESTION: 70

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.

Answer Area

SELECT ProductID,

GREATEST

COALESCE

GREATEST

IIF

MAX

(ListPrice, WholesalePrice, AgentPrice) AS HighestSellingPrice,

COALESCE

CHOOSE

COALESCE

IIF

MAX

(AgentPrice, WholesalePrice, ListPrice) AS TradePrice

FROM

Answer:



Explanation:

For the HighestSellingPrice, you should use the GREATEST function to find the highest value from the given price columns. However, T-SQL does not have a GREATEST function as found in some other SQL dialects, so you would typically use a CASE statement or an IIF statement with nested MAX functions. Since neither of those are provided in the options, you should select MAX as a placeholder to indicate the function that would be used to find the highest value if combining multiple MAX functions or a similar logic was available.

For the TradePrice, you should use the COALESCE function, which returns the first non-null value in a list.

The COALESCE function is the correct choice as it will return AgentPrice if it 's not null; if AgentPrice is null, it will check WholesalePrice, and if that is also null, it will return ListPrice.

The complete code with the correct SQL functions would look like this:

```
SELECT ProductID,
```

```
MAX(ListPrice, WholesalePrice, AgentPrice) AS HighestSellingPrice, -- MAX is used as a placeholder  
COALESCE(AgentPrice, WholesalePrice, ListPrice) AS TradePrice FROM  
Sales.Products
```

Select MAX for HighestSellingPrice and COALESCE for TradePrice in the answer area.

NEW QUESTION: 71

You have a Fabric tenant that contains a lakehouse named Lakehouse1.

You need to prevent new tables added to Lakehouse1 from being added automatically to the default semantic model of the lakehouse.

What should you configure? (5)

- A. the semantic model settings
- B. the Lakehouse1 settings
- C. the workspace settings
- D. the SQL analytics endpoint settings


Answer: (SHOW ANSWER)

To prevent new tables added to Lakehouse1 from being automatically added to the default semantic model, you should configure the semantic model settings. There should be an option within the settings of the semantic model to include or exclude new tables by default. By adjusting these settings, you can control the automatic inclusion of new tables.

References: The management of semantic models and their settings would be covered under the documentation for the semantic layer or modeling features of the Fabric tenant's lakehouse solution.

NEW QUESTION: 72

You have a Microsoft Power BI report and a semantic model that uses Direct Lake mode. From Power BI Desktop, you open Performance analyzer as shown in the following exhibit.



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

Answer Area

The Direct Lake fallback behavior is set to [answer choice].

The query for the table visual is executed by using [answer choice].

Answer:

Answer Area

The Direct Lake fallback behavior is set to [answer choice].

The query for the table visual is executed by using [answer choice].

Explanation:

The Direct Lake fallback behavior is set to: DirectQueryOnly

The query for the table visual is executed by using: DirectQuery

In the context of Microsoft Power BI, when using DirectQuery in Direct Lake mode, there is no caching of data and all queries are sent directly to the underlying data source. The Performance Analyzer tool shows the time taken for different operations, and from the options provided, it indicates that DirectQuery mode is being used for the visuals, which is consistent with the Direct Lake setting. DirectQueryOnly as the fallback behavior ensures that only DirectQuery will be used without reverting to import mode.

NEW QUESTION: 73

You have a Microsoft Power BI semantic model.

You need to identify any surrogate key columns in the model that have the Summarize By property set to a value other than to None. The solution must minimize effort.

What should you use?

- A. DAX Formatter in DAX Studio
- B. Model view in Microsoft Power BI Desktop
- C. Model explorer in Microsoft Power BI Desktop
- D. Best Practice Analyzer in Tabular Editor

Answer: ([SHOW ANSWER](#))

To identify surrogate key columns with the "Summarize By" property set to a value other than "None," the Best Practice Analyzer in Tabular Editor is the most efficient tool. The Best Practice Analyzer can analyze the entire model and provide a report on all columns that do not meet a specified best practice, such as having the "Summarize By" property set correctly for surrogate key columns. Here's how you would proceed:

- * Open your Power BI model in Tabular Editor.
- * Go to the Advanced Scripting window.
- * Write or use an existing script that checks the "Summarize By" property of each column.
- * Execute the script to get a report on the surrogate key columns that do not have their "Summarize By" property set to "None".
- * You can then review and adjust the properties of the columns directly within the Tabular Editor.

References: The functionality of the Best Practice Analyzer in Tabular Editor is documented in the community and learning resources for Power BI.

NEW QUESTION: 74

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 BI report named Report1 that uses the model. Report1 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

Answer: ([SHOW ANSWER](#))

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. References = Microsoft Power BI documentation on performance optimization offers in-depth knowledge on these features.


NEW QUESTION: 75

You have a Fabric tenant that contains lakehouse named Lakehouse1. Lakehouse1 contains a Delta table with eight columns. You receive new data that contains the same eight columns and two additional columns.


You create a Spark DataFrame and assign the DataFrame to a variable named df. The DataFrame contains the new data. You need to add the new data to the Delta table to meet the following requirements:

- * Keep all the existing rows.
- * Ensure that all the new data is added to the table.

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



Answer:



Explanation:



To add new data to the Delta table while meeting the specified requirements:

You should use the append mode to ensure that all new data is added to the table without affecting the existing rows.

You should set the mergeSchema option to true to allow the schema of the Delta table to be updated with the new columns found in the DataFrame.

The completed code would look like this:

```
df.write.format( " delta " ).mode( " append " )
option( " mergeSchema " , " true " )
saveAsTable( " Lakehouse1.TableName " )
```

NEW QUESTION: 76

You are analyzing customer purchases in a Fabric notebook by using PySpark. You have the following DataFrames:

- **transactions:** Contains five columns named `transaction_id`, `customer_id`, `product_id`, `amount`, and `date` and has 10 million rows, with each row representing a transaction
- **customers:** Contains customer details in 1,000 rows and three columns named `customer_id`, `name`, and `country`

You need to join the DataFrames on the `customer_id` column. The solution must minimize data shuffling. You write the following code.

```
from pyspark.sql import functions as F

results =
```

Which code should you run to populate the results DataFrame?

- A. `transactions.join(F.broadcast(customers), transactions.customer_id == customers.customer_id)`
- B. `transactions.join(customers, transactions.customer_id == customers.customer_id).distinct()`
- C. `transactions.join(customers, transactions.customer_id == customers.customer_id)`
- D. `transactions.crossJoin(customers).where(transactions.customer_id == customers.customer_id)`

Answer: A (LEAVE A REPLY)

The correct code to populate the results DataFrame with minimal data shuffling is Option A. Using the broadcast function in PySpark is a way to minimize data movement by broadcasting the smaller DataFrame (customers) to each node in the cluster. This is ideal when one DataFrame is much smaller than the other, as in this case with customers. References = You can refer to the official Apache Spark documentation for more details on joins and the broadcast hint.

Valid DP-600 Dumps shared by Actual4test.com for Helping Passing DP-600 Exam! Actual4test.com now offer the **newest DP-600 exam dumps**, the Actual4test.com DP-600 exam questions have been updated and answers have been corrected get the **newest** Actual4test.com DP-600 dumps with Test Engine here:

https://www.actual4test.com/DP-600_examcollection.html (212 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)

NEW QUESTION: 77

You have a Fabric workspace named Workspace1 that contains a data flow named Dataflow1. Dataflow1 contains a query that returns the data shown in the following exhibit.



You need to transform the date columns into attribute-value pairs, where columns become rows.

You select the VendorID column.

Which transformation should you select from the context menu of the VendorID column?

- A. Group by
- B. Unpivot columns
- C. Unpivot other columns
- D. Split column
- E. Remove other columns

Answer: (SHOW ANSWER)

The transformation you should select from the context menu of the VendorID column to transform the date columns into attribute-value pairs, where columns become rows, is Unpivot columns (B). This transformation will turn the selected columns into rows with two new columns, one for the attribute (the original column names) and one for the value (the data from the cells). References = Techniques for unpivoting columns are covered in the Power Query documentation, which explains how to use the transformation in data modeling.

NEW QUESTION: 78

You have a Fabric lakehouse named Lakehouse1 that contains the following data.

Table name	Column name	Data type
dbo.publicholidays	countryOrRegion	Char(10)
dbo.publicholidays	holidayName	Varchar(6000)
dbo.publicholidays	date	Date
dbo.sales	OrderDate	Date
dbo.sales	Quantity	Float
dbo.sales	UnitPrice	Float

You need build a T-SQL statement that will return the total sales amount by OrderDate only for the days that are holidays in Australia. The total sales amount must sum the quantity multiplied by the price on each row in the dbo.sales table.

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 s.OrderDate,

s.Quantity * S.UnitPrice
Sum(s.Quantity * s.UnitPrice)
Sum(s.quantity) * s.UnitPrice
Sum(s.Quantity) * Sum(s.UnitPrice)
GROUP BY s.OrderDate

AS TotalSalesAmt

JOIN [Lakehouse1].[dbo].[publicholidays] AS ph ON s.OrderDate = ph.date



Microsoft

Cross
Full outer
Inner
Right outer

JOIN [Lakehouse1].[dbo].[publicholidays] AS ph ON s.OrderDate = ph.date
region = 'Australia'

Answer:

Answer Area

SELECT s.OrderDate,

s.Quantity * S.UnitPrice
Sum(s.Quantity * s.UnitPrice)
Sum(s.quantity) * s.UnitPrice
Sum(s.Quantity) * Sum(s.UnitPrice)
GROUP BY s.OrderDate

AS TotalSalesAmt

JOIN [Lakehouse1].[dbo].[publicholidays] AS ph ON s.OrderDate = ph.date




Microsoft

Cross
Full outer
Inner
Right outer

JOIN [Lakehouse1].[dbo].[publicholidays] AS ph ON s.OrderDate = ph.date
region = 'Australia'

Explanation:



Microsoft

SELECT s.OrderDate,

s.Quantity * S.UnitPrice

AS TotalSalesAmt

FROM [Lakehouse1].[dbo].[sales] AS s

INNER

JOIN [Lakehouse1].[dbo].[publicholidays] AS ph ON s.OrderDate = ph.date

WHERE ph.countryregion = 'Australia'

GROUP BY s.OrderDate

Step 1: Requirement

We need to calculate total sales amount by OrderDate.

Total Sales Amount = Quantity × UnitPrice per row.

Then SUM over all rows per OrderDate.

Only include rows where the OrderDate matches a public holiday in Australia.

Step 2: Calculation Logic

If we select s.Quantity * s.UnitPrice, it only computes row-level values.

To calculate the total sales amount per OrderDate, we need:

SUM(s.Quantity * s.UnitPrice) AS TotalSalesAmt

Thus, the correct calculation: SUM(s.Quantity * s.UnitPrice)

Step 3: Join Type

We want to return only those sales that occurred on public holidays in Australia.

If we used a LEFT JOIN, we'd include all sales regardless of holiday.

A FULL OUTER JOIN would include non-matching rows, not required.

A CROSS JOIN would duplicate results unnecessarily.

A RIGHT OUTER JOIN would include holidays with no sales, which is not needed.

The correct choice is an INNER JOIN on matching dates.

Step 4: Final Query

SELECT

s.OrderDate,

SUM(s.Quantity * s.UnitPrice) AS TotalSalesAmt

FROM dbo.sales AS s

INNER JOIN dbo.publicholidays AS ph

ON s.OrderDate = ph.Date

WHERE ph.countryOrRegion = ' Australia '

GROUP BY s.OrderDate;

Why This is Correct

SUM(s.Quantity * s.UnitPrice) # ensures sales are aggregated correctly.

INNER JOIN # ensures only rows where OrderDate matches a holiday are included.

WHERE ph.countryOrRegion = ' Australia ' # filters holidays to Australia only.

References

SUM (Transact-SQL)

JOIN operations in T-SQL

NEW QUESTION: 79

You plan to deploy Microsoft Power BI items by using Fabric deployment pipelines. You have a deployment pipeline that contains three stages named Development, Test, and Production.

A workspace is assigned to each stage.

You need to provide Power BI 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.

Ensure that developers can view items in 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. Contributor access to the Development and Test workspaces
- B. Viewer access to the Development and Test workspaces
- C. Admin access to the deployment pipeline
- D. Contributor access to the Production workspace
- E. Viewer access to the Production workspace
- F. Build permission to the production semantic models

Answer: A,C,E ([LEAVE A REPLY](#))

NEW QUESTION: 80

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

You create a Microsoft Power BI semantic model named Modell that uses the CSV files as a data source. You configure incremental refresh for Model 1 and publish the model to a Premium capacity in the Fabric tenant.

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

What is a possible cause of the failure?

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

Answer: D ([LEAVE A REPLY](#))

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

NEW QUESTION: 81

You have a KQL database that contains a table named Readings.

You need to query Readings and return the results shown in the following table.

City	Area	MeterReading	Datetime	PrevMeterReading	PrevDatetime
Copenhagen	Area1	2700	2022-04-20 10:00:00	None	None
Copenhagen	Area2	2720	2022-04-20 11:00:00	2700	2022-04-20 10:00:00
Copenhagen	Area1	2750	2022-04-20 12:00:00	2720	2022-04-20 11:00:00
Copenhagen	Area2	2780	2022-04-20 13:00:00	2750	2022-04-20 12:00:00

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

NOTE: Each correct selection is worth one point.

Answer Area

Readings

| filter City == "Copenhagen"

| sort by Datetime

| extend

| evaluate

| extend

| lookup

| project

| summarize

PrevMeterReading = prev(MeterReading), PrevTimestamp = prev(Datetime)

| project

| evaluate

| extend

| lookup

| project

| summarize

City, Area, MeterReading, Datetime, PrevMeterReading, PrevDatetime

Answer:
Answer Area

Readings

| filter City == "Copenhagen"

| sort by Datetime

| extend

| evaluate

| extend

| lookup

| project

| summarize

PrevMeterReading = prev(MeterReading), PrevTimestamp = prev(Datetime)

| project

| evaluate

| extend

| lookup

| project

| summarize

City, Area, MeterReading, Datetime, PrevMeterReading, PrevDatetime

Explanation:

Answer Area

Readings

| filter City == "Copenhagen"

| sort by Datetime

| extend

PrevMeterReading = prev(MeterReading), PrevTimestamp = prev(Datetime)

| project

City, Area, MeterReading, Datetime, PrevMeterReading, PrevDatetime

Comprehensive Detailed Explanation

We need to query a KQL database table named Readings to generate output that includes both the current reading and the previous reading values and datetime for each row.

Step 1: Understanding Requirements

The output table shows:

City, Area, MeterReading, Datetime

PrevMeterReading, PrevDatetime

The prev() function in KQL is used to access the value of a column from the previous row (based on order).

Step 2: Query Construction

We start with filtering:

Readings

| where City == " Copenhagen "

| sort by Datetime

To add columns that calculate the previous reading and previous datetime, we use extend:

| extend PrevMeterReading = prev(MeterReading), PrevDatetime = prev(Datetime) extend creates new columns.

prev() pulls the previous row's values.

Finally, we only want specific columns in the output. For this, we use project:

| project City, Area, MeterReading, Datetime, PrevMeterReading, PrevDatetime

| where City == " Copenhagen "

| sort by Datetime

| extend PrevMeterReading = prev(MeterReading), PrevDatetime = prev(Datetime)

| project City, Area, MeterReading, Datetime, PrevMeterReading, PrevDatetime

prev() gives access to the previous row.

project selects only the required columns.

This matches exactly the table shown in the question.

References

Kusto Query Language - extend operator

Kusto Query Language - project operator

prev() function in KQL

NEW QUESTION: 82

You have the following KQL query named Query1.

```
SigninLogs
| summarize FailedLogons = countif(ResultType != 0) by UserPrincipalName, OfficeLocation
| project UserPrincipalName, FailedLogons, OfficeLocation
| where OfficeLocation == 'Seattle'
```

You need to improve the performance of Query1.

The solution must minimize resource use How should you restructure Query1?

A.

```
SigninLogs
| where OfficeLocation == 'Seattle'
| summarize FailedLogons = countif(ResultType != 0) by UserPrincipalName, OfficeLocation
| project UserPrincipalName, FailedLogons, OfficeLocation
```

B.

```
SigninLogs
| summarize FailedLogons = countif(ResultType != 0) by UserPrincipalName, OfficeLocation
| where OfficeLocation == 'Seattle'
| project UserPrincipalName, FailedLogons, OfficeLocation
```

C.

```
SigninLogs
| project UserPrincipalName, FailedLogons, OfficeLocation
| where OfficeLocation == 'Seattle'
| summarize FailedLogons = countif(ResultType != 0) by UserPrincipalName, OfficeLocation
```

D.

```
SigninLogs
| project UserPrincipalName, FailedLogons, OfficeLocation
| summarize FailedLogons = countif(ResultType != 0) by UserPrincipalName, OfficeLocation
| where OfficeLocation == 'Seattle'
```

Answer: A ([LEAVE A REPLY](#))

NEW QUESTION: 83

You have a Fabric workspace that contains a warehouse named Warehouse1. Warehouse1 contains the following data.

Table name	Column name	Data type
Employee	EmployeeID	Int
Employee	EmployeeName	Varchar(128)
Employee	EmployeePosition	Varchar(64)
Contract	EmployeeID	Int
Contract	ContractType	Varchar(64)
Contract	StartDate	Datetime2
Contract	EndDate	Datetime2

You need to create a T-SQL statement that will denormalize the tables and include the ContractType and StartDate attributes in the results. The solution must meet the following requirements:

- * Include attributes from matching rows in the Contract table.
- * Ensure that all the rows from the Employee table are preserved.
- * Return the total number of employees per contract type for all the contract types that have more than two employees.

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

WITH result AS(

SELECT

e.EmployeeID

, e.EmployeeName

, e.EmployeePosition

, c.ContractType

FROM Employee AS e

LEFT OUTER JOIN

CROSS JOIN

INNER JOIN

LEFT OUTER JOIN

RIGHT OUTER JOIN

Contract AS c on c.EmployeeID = e.EmployeeID)

AS TotalEmployees, ContractType

FROM result

GROUP BY ContractType

HAVING

CONTAINS

HAVING

LIMIT

WHERE

COUNT(DISTINCT EmployeeID) > 2



Answer:

```

WITH result AS(
SELECT
    e.EmployeeID
    , e.EmployeeName
    , e.EmployeePosition
    , c.ContractType
FROM Employee AS e
LEFT OUTER JOIN Contract AS c on c.EmployeeID = e.EmployeeID
) AS TotalEmployees, ContractType
FROM result
GROUP BY ContractType
HAVING COUNT(DISTINCT EmployeeID) > 2

```

Explanation:

Answer Area

```

WITH result AS(
SELECT
    e.EmployeeID
    , e.EmployeeName
    , e.EmployeePosition
    , c.ContractType
FROM Employee AS e
LEFT OUTER JOIN Contract AS c on c.EmployeeID = e.EmployeeID
) AS TotalEmployees, ContractType
FROM result
GROUP BY ContractType
HAVING COUNT(DISTINCT EmployeeID) > 2

```

Comprehensive Detailed Explanation

We are tasked with writing a T-SQL query to denormalize the tables Employee and Contract in a Fabric warehouse.

Requirements Breakdown

Include attributes from matching rows in the Contract table

This means we must join Employee and Contract on EmployeeID.

Ensure that all rows from the Employee table are preserved

This requires a LEFT OUTER JOIN from Employee to Contract. If an employee has no contract, we still want the employee record included.

Return the total number of employees per contract type for all contract types that have more than two employees This means we need to:

Group by ContractType

Count distinct EmployeeID

Filter groups where that count > 2.

To filter aggregated results, we use HAVING, not WHERE.

Completed Query

WITH result AS (

SELECT

e.EmployeeID,

e.EmployeeName,

e.EmployeePosition,

c.ContractType,

c.StartDate

FROM Employee AS e

LEFT OUTER JOIN Contract AS c

ON c.EmployeeID = e.EmployeeID

)

SELECT

ContractType,

COUNT(DISTINCT EmployeeID) AS TotalEmployees

FROM result

GROUP BY ContractType

HAVING COUNT(DISTINCT EmployeeID) > 2;

Why This is Correct

LEFT OUTER JOIN ensures all employees are included, regardless of contracts.

HAVING allows filtering aggregated groups by employee counts.

The query returns ContractType, StartDate (as requested in denormalization), and aggregated employee counts.

References

T-SQL JOINS in Microsoft Fabric

Aggregate functions and HAVING clause

Denormalization best practices in Fabric warehouses

NEW QUESTION: 84

You have a Fabric workspace named Workspace1 that contains a KQL database. The database contains a table named Table1. Workspace1 is allocated to an F 128 capacity. Table1 contains 10 billion rows and 25 columns named d through c25. Column cl contains identifiers that range from 1 to 15 and are equally distributed.

You have the following query.

```
Table1
| project c1, c2, c3
| where c1 == 1
```

You modify the query as shown.

```
Table1
| where c1 == 1
| project c1, c2, c3
```


How will the execution time of the query be affected?

- A. The execution time will longer.
- B. The execution time will remain the same.
- C. The query will fail
- D. The execution time will be shorter.

Answer: B ([LEAVE A REPLY](#))

NEW QUESTION: 85

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



Which type of analytics are you performing?

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

Answer: B ([LEAVE A REPLY](#))

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.

References: Descriptive analytics and the use of histograms as a way to visualize data distribution are basic concepts in data analysis, often covered in introductory analytics and Python programming resources.

NEW QUESTION: 86

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

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

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

Solution: You use the following PySpark expression:

```
df.explain()
```

Does this meet the goal?

- A. Yes
- B. No

Answer: ([SHOW ANSWER](#))

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

NEW QUESTION: 87

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 BI Desktop?

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

Answer: C ([LEAVE A REPLY](#))

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. References = Microsoft Power BI Desktop documentation outlines the functionalities available in Model view for managing semantic models.

Valid DP-600 Dumps shared by Actual4test.com for Helping Passing DP-600 Exam! Actual4test.com now offer the **newest DP-600 exam dumps**, the Actual4test.com DP-600 exam **questions have been updated** and **answers have been corrected** get the **newest** Actual4test.com DP-600 dumps with Test Engine here:
https://www.actual4test.com/DP-600_examcollection.html (212 Q&As Dumps, **30%OFF** Special Discount: **Freepdfdumps**)