

100% Money Back
Guarantee

Vendor:Microsoft

Exam Code:70-768

Exam Name:Developing SQL Data Models

Version:Demo

QUESTION 1

A database named DB2 uses the InMemory query mode. Users frequently run the following query:

```
EVALUATE
  FILTER (
    ADDCOLUMNS (
      VALUES ('Date' [Calendar Year]),
      "Sales", CALCULATE (SUM ('Internet Sales' [Sales Amount] ) )
    ),
    [Sales] > 8000000
  )
ORDER BY 'Date' [Calendar Year]
```

You need to configure SQL Server Profiler to determine why the query is performing poorly.

Which three event should you monitor on the SQL Server Profiler trace events configuration page? To answer, select the appropriate options in the answer area.

Hot Area:

Answer area

| Events | |
|-------------------------------------|-------------------------------|
| <input checked="" type="checkbox"/> | Query Processing |
| <input type="checkbox"/> | Calculation Evaluation |
| <input type="checkbox"/> | DAX Query Plan |
| <input type="checkbox"/> | DirectQuery Begin |
| <input type="checkbox"/> | DirectQuery End |
| <input type="checkbox"/> | Query Dimension |
| <input type="checkbox"/> | Query Subcube |
| <input type="checkbox"/> | VertiPaq SE Query Cache Match |
| <input type="checkbox"/> | VertiPaq SE Query Cache Miss |

Correct Answer:

Answer area

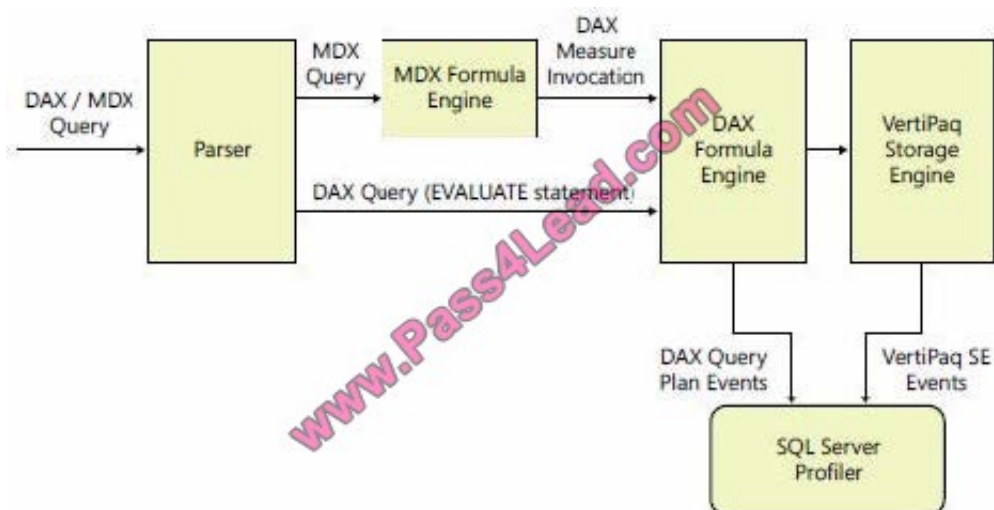
| Events | |
|-------------------------------------|-------------------------------|
| <input checked="" type="checkbox"/> | Query Processing |
| <input type="checkbox"/> | Calculation Evaluation |
| <input type="checkbox"/> | DAX Query Plan |
| <input type="checkbox"/> | DirectQuery Begin |
| <input type="checkbox"/> | DirectQuery End |
| <input type="checkbox"/> | Query Dimension |
| <input type="checkbox"/> | Query Subcube |
| <input type="checkbox"/> | VertiPaq SE Query Cache Match |
| <input type="checkbox"/> | VertiPaq SE Query Cache Miss |

By using SQL Profiler, you can intercept two classes of trace events from Analysis Services, DAX Query Plan and DirectQuery events, both generated by the DirectQuery engine. Here, in this scenario we have a DAX Query.

DAX Query Plan events are generated by the DAX formula.

By using the In-Memory mode, you store a copy of data in the xVelocity (VertiPaq) storage engine.

Figure: This is how a query is executed by using In-Memory mode.



References: Microsoft SQL Server 2012 Analysis Services, The BISM Tabular Model, Microsoft Press (July 2012), page 331 From Scenario: Users report that the query takes a long time to complete.

QUESTION 2

You are writing a MDX query to retrieve data from a Microsoft SQL Server Analysis Services (SSAS) cube named Channel Sales. The cube defines two measures named Sales and Cost. The cube also defines a Date dimension and a Product dimension.

You need to retrieve profit values for a year named CY2016.

How should you complete the MDX statement? To answer, drag the appropriate MDX segment to the correct locations. Each MDX segment 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.

Select and Place:

MDX segments

- WITH MEMBER [Measures].[Profit] AS ([Sales]-[Cost])
- WITH SET [Measures].[Profit] AS ([Sales]-[Cost])
- WHERE ([Date].[Year].[CY2016])
- WHERE ([Date].[Year] = [CY2016])

Answer Area

```
MDX segment
```

```
SELECT
```

```
    {[Measures].[Profit]} ON COLUMNS,
```

```
    [Product].[Category].[Category].MEMBERS ON ROWS
```

```
FROM [Channel Sales]
```

```
MDX segment
```

Correct Answer:

MDX segments

-
- WITH SET [Measures].[Profit] AS ([Sales]-[Cost])
-
- WHERE ([Date].[Year] = [CY2016])

Answer Area

```
WITH MEMBER [Measures].[Profit] AS ([Sales]-[Cost])
```

```
SELECT
```

```
    {[Measures].[Profit]} ON COLUMNS,
```

```
    [Product].[Category].[Category].MEMBERS ON ROWS
```

```
FROM [Channel Sales]
```

```
WHERE ([Date].[Year].[CY2016])
```

Box 1:WITH MEMBER

QUESTION 3

You need to resolve the issues that the users report.

Which processing options should you use? To answer, drag the appropriate processing option to the correct location or locations. Each processing option may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Select and Place:

Processing options

- Process Clear
- Process Update
- Process Index
- Process Default
- Process Data
- Process Full

Answer Area

Data availability during cube processing

- Maximum data availability
- Less than maximum data availability
- Least data availability

Processing option

- Processing option
- Processing option
- Processing option

Correct Answer:

Processing options

- Process Clear
-
- Process Index
-
- Process Data
-

Answer Area

Data availability during cube processing

- Maximum data availability
- Less than maximum data availability
- Least data availability

Processing option

- Process Full
- Process Default
- Process Update

Box1: Process Full:

When Process Full is executed against an object that has already been processed, Analysis Services drops all data in the object, and then processes the object. This kind of processing is required when a structural change has been made to

an object, for example, when an attribute hierarchy is added, deleted, or renamed.

Box 2: Process Default

Detects the process state of database objects, and performs processing necessary to deliver unprocessed or partially processed objects to a fully processed state. If you change a data binding, Process Default will do a Process Full on the affected object.

Box 3:

Not Process Update: Forces a re-read of data and an update of dimension attributes. Flexible aggregations and indexes on related partitions will be dropped.

QUESTION 4

You need to create the cube processing job and the dimension processing job. Which processing task should you use for each job? To answer, drag the appropriate processing tasks to the correct locations. Each processing task 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.

Select and Place:

| Processing tasks | Answer Area | | | | | | |
|----------------------------------|--|-----|-----------------|-----------------------------|-----------------|----------------------------------|-----------------|
| Process Clear | <table border="1"><thead><tr><th>Job</th><th>Processing task</th></tr></thead><tbody><tr><td>Incremental cube processing</td><td>Processing task</td></tr><tr><td>Incremental dimension processing</td><td>Processing task</td></tr></tbody></table> | Job | Processing task | Incremental cube processing | Processing task | Incremental dimension processing | Processing task |
| Job | Processing task | | | | | | |
| Incremental cube processing | Processing task | | | | | | |
| Incremental dimension processing | Processing task | | | | | | |
| Process Update | | | | | | | |
| Process Index | | | | | | | |
| Process Add | | | | | | | |
| Process Data | | | | | | | |
| Process Structure | | | | | | | |

Correct Answer:

| Processing tasks | Answer Area | | | | | | |
|----------------------------------|--|-----|-----------------|-----------------------------|--------------|----------------------------------|----------------|
| Process Clear | <table border="1"><thead><tr><th>Job</th><th>Processing task</th></tr></thead><tbody><tr><td>Incremental cube processing</td><td>Process Data</td></tr><tr><td>Incremental dimension processing</td><td>Process Update</td></tr></tbody></table> | Job | Processing task | Incremental cube processing | Process Data | Incremental dimension processing | Process Update |
| Job | Processing task | | | | | | |
| Incremental cube processing | Process Data | | | | | | |
| Incremental dimension processing | Process Update | | | | | | |
| Process Update | | | | | | | |
| Process Index | | | | | | | |
| Process Add | | | | | | | |
| Process Data | | | | | | | |
| Process Structure | | | | | | | |

Box 1: ProcessData Processes data only without building aggregations or indexes. If there is data in the partitions, it will be dropped before re-populating the partition with source data. Box 2: Process Update Forces a re-read of data and an update of dimension attributes. Flexible aggregations and indexes on related partitions will be dropped. References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensionalmodels/processing-options-and-settings-analysis-services>

QUESTION 5

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains

a unique solution. Determine whether the solution meets the stated goals.

You have a Microsoft SQL Server Analysis Services (SSAS) multidimensional database that stores customer and order data for customers in the United States only. The database contains the following objects:

| Type | Name | Content |
|-----------|---|---|
| Measure | Reseller Average Unit Price | the average unit price of sales |
| Dimension | Geography | the location of resellers |
| Hierarchy | Geography.State-Province | the state or province where the reseller is located |
| Member | Geography.State-Province.&[WA]&[US], Geography.State-Province.&[GA]&[US] | a specific state and country/region |

You must create a KPI named Large Sales Target that uses the Traffic Light indicator to display status. The KPI must contain:

| Expression type | Description |
|-----------------|--|
| Value | the reseller average unit price |
| Goal | the average reseller average unit price for US states other than Colorado (CO) |
| Status | a green indicator if the value is at least 10 percent above the goal, a red indicator if the value is 15 percent or more below the goal, and a yellow indicator for other values |
| Trend | the value for trend is always 0 |

You need to create the KPI.

Solution: You set the value of the Status expression to:

```

Case
  When KpiValue("Large Sales Target")/KpiGoal("Large Sales Target") >= 1.1
    Then 1
  When KpiValue("Large Sales Target")/KpiGoal("Large Sales Target") < 1.1
    And
      KpiValue("Large Sales Target")/KpiGoal("Large Sales Target") > .85
    Then 0
  Else-1
End
  
```

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

QUESTION 6

You are developing a SQL Server Analysis Services (SSAS) tabular project.

A column named City must be added to the table named Customer. The column will be used in the definition of a hierarchy. The City column exists in the Geography table that is related to the Customer table.

You need to add the City column to the Customer table.

How should you write the calculation?

- A. City:= LOOKUP(Geography[City],Geography[GeographyKey],[GeographyKey])
- B. City:= LOOKUPVALUE(Geography[City],Geography[GeographyKey],[GeographyKey]) C
. =RELATED(Geography[City])
- C. =RELATED(Geography.City)
- D. =VALUES(Geography[City])
- E. City:=VALUES(Geography[City])

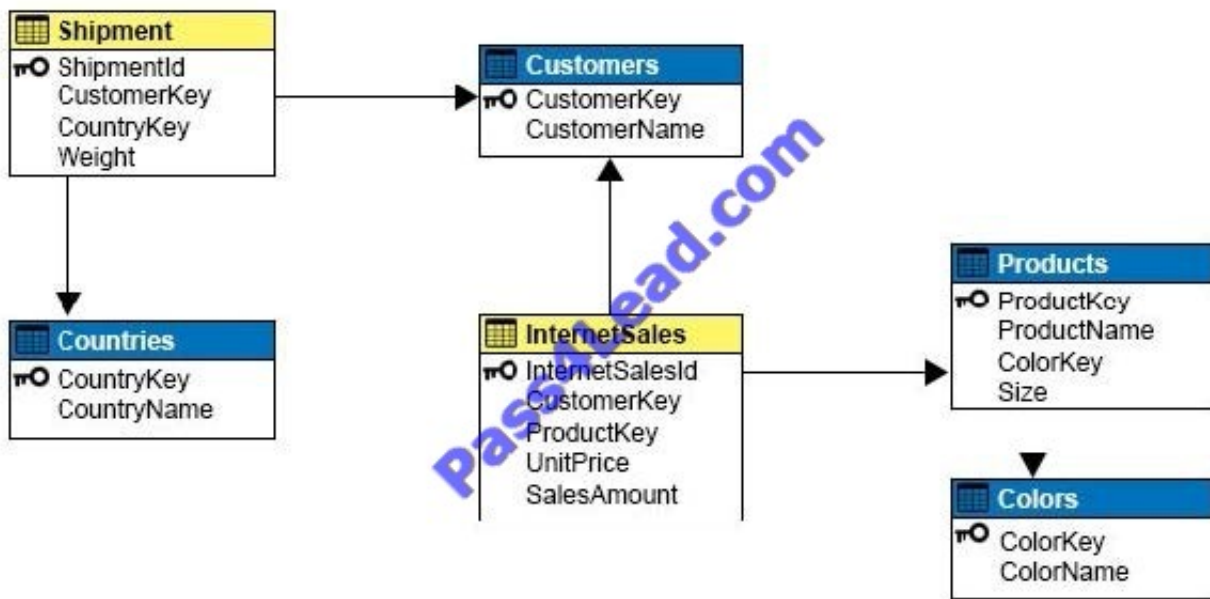
Correct Answer: C

Explanation: * RELATED Function Returns a related value from another table.

QUESTION 7

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You have a Microsoft SQL Server Analysis Services (SSAS) instance that is configured to use multidimensional mode. You create the following cube:



Users need to be able to analyze sales by product and color.

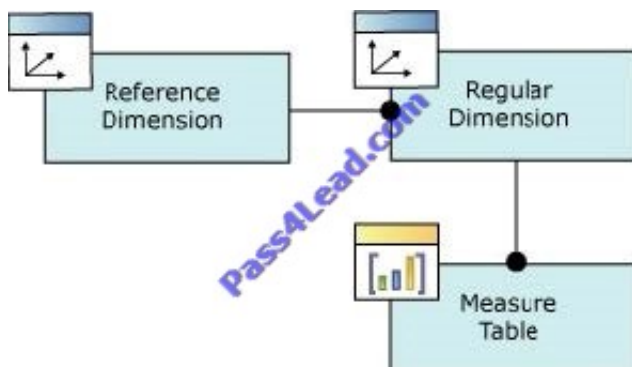
You need to create the dimension.

Which relationship type should you use between the InternetSales table and the new dimension?

- A. no relationship
- B. regular
- C. fact
- D. referenced
- E. many-to-many
- F. data mining

Correct Answer: D

A reference dimension relationship between a cube dimension and a measure group exists when the key column for the dimension is joined indirectly to the fact table through a key in another dimension table, as shown in the following illustration.



A reference dimension relationship represents the relationship between dimension tables and a fact table in a snowflake schema design. When dimension tables are connected in a snowflake schema, you can define a single dimension using columns from multiple tables, or you can define separate dimensions based on the separate dimension tables and then define a link between them using the reference dimension relationship setting. The following figure shows one fact table named InternetSales, and two dimension tables called Customer and Geography, in a snowflake schema.



You can create two dimensions related to the InternetSales measure group: a dimension based on the Customer table, and a dimension based on the Geography table. You can then relate the Geography dimension to the InternetSales measure group using a reference dimension relationship using the Customer dimension.

QUESTION 8

You need to configure the project option settings to minimize deployment time for the CustomerAnalysis data model. What should you do? To answer, select the appropriate setting from each list in the answer area.

Hot Area:

Answer Area

| Location | Setting |
|--------------------------|---|
| Processing option | <div style="border: 1px solid gray; padding: 2px;">▼</div> <div style="border: 1px solid gray; padding: 2px;">Default</div> <div style="border: 1px solid gray; padding: 2px;">Do not process</div> <div style="border: 1px solid gray; padding: 2px;">Full</div> |
| Transactional deployment | <div style="border: 1px solid gray; padding: 2px;">▼</div> <div style="border: 1px solid gray; padding: 2px;">False</div> <div style="border: 1px solid gray; padding: 2px;">True</div> |

Correct Answer:

Answer Area

| Location | Setting | | | | |
|--------------------------|--|---|---------|----------------|------|
| Processing option | <table border="1"><tr><td>▼</td></tr><tr><td>Default</td></tr><tr><td>Do not process</td></tr><tr><td>Full</td></tr></table> | ▼ | Default | Do not process | Full |
| ▼ | | | | | |
| Default | | | | | |
| Do not process | | | | | |
| Full | | | | | |
| Transactional deployment | <table border="1"><tr><td>▼</td></tr><tr><td>False</td></tr><tr><td>True</td></tr></table> | ▼ | False | True | |
| ▼ | | | | | |
| False | | | | | |
| True | | | | | |

Scenario:

Box 1, Processing option:Default

Process Default detects the process state of database objects, and performs processing necessary to deliver unprocessed or partially processed objects to a fully processed state. If you change a data binding, Process Default will do a

Process Full on the affected object.

Note: Processing Method This setting controls whether the deployed objects are processed after deployment and the type of processing that will be performed. There are three processing options:

Default processing (default)

Full processing

None

Box 2, Transactional deployment: False

If this option is False, Analysis Services deploys the metadata changes in a single transaction, and deploys each processing command in its own transaction.

Scenario: The CustomerAnalysis data model will contain a large amount of data and needs to be shared with other developers even if a deployment fails. Each time you deploy a change during development, processing takes a long time.

References:<https://docs.microsoft.com/en-us/sql/analysis-services/multidimensionalmodels/deployment-script-files-specifying-processing-options>

QUESTION 9

You are responsible for installing new database server instances.

You must install Microsoft SQL Server Analysis Services (SSAS) to support deployment of the following projects. You develop both projects by using SQL Server Data Tools.

You need to install the appropriate services to support both projects.

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

- A. Install one tabular instance of SSAS and enable the Data Mining Extensions.
- B. Install one multidimensional instance of SSAS.
- C. Install one tabular instance of SSAS.
- D. Install a multidimensional instance and a Power Pivot instance of SSAS on the same server.
- E. Install two separate tabular instances of SSAS.

Correct Answer: BC

Analysis Services can be installed in one of three server modes: Multidimensional and Data Mining (default), Power Pivot for SharePoint, and Tabular. References: <https://docs.microsoft.com/en-us/sql/analysis-services/comparing-tabular-and-multidimensional-solutions-ssas>

QUESTION 10

You are deploying a multidimensional Microsoft SQL Server Analysis Services (SSAS) project. You add two new role-playing dimensions named Picker and Salesperson to the cube. Both of the cube dimensions are based upon the underlying dimension named Employee in the data source view.

Users report that they are unable to differentiate the Salesperson attributes from the Picker attributes.

You need to ensure that the Salesperson and Picker attributes in each dimension use unique names.

In the table below, identify an option that you would use as part of the process to alter the names of the attributes for each of the dimensions.

NOTE: Make only one selection in each column.

Hot Area:

Answer Area

| Option | Dimension Picker | Dimension Salesperson |
|---|--------------------------|--------------------------|
| Create a second data source view. | <input type="checkbox"/> | <input type="checkbox"/> |
| Rename the Employee dimension. | <input type="checkbox"/> | <input type="checkbox"/> |
| Create a new named query for both dimensions. | <input type="checkbox"/> | <input type="checkbox"/> |

Correct Answer:

Answer Area

| Option | Dimension Picker | Dimension Salesperson |
|---|-------------------------------------|-------------------------------------|
| Create a second data source view. | <input type="checkbox"/> | <input type="checkbox"/> |
| Rename the Employee dimension. | <input type="checkbox"/> | <input type="checkbox"/> |
| Create a new named query for both dimensions. | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> |

A named query is a SQL expression represented as a table. In a named query, you can specify an SQL expression to select rows and columns returned from one or more tables in one or more data sources. A named query is like any other

table in a data source view (DSV) with rows and relationships, except that the named query is based on an expression.

A named query lets you extend the relational schema of existing tables in DSV without modifying the underlying data source.

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensionalmodels/define-named-queries-in-a-data-source-view-analysis-services>

QUESTION 11

You are developing a SQL Server Analysis Services (SSAS) tabular project.

You need to grant the minimum permissions necessary to enable users to query data in a data model.

Which role permission should you use?

- A. Explorer
- B. Process
- C. Browser
- D. Administrator
- E. Select
- F. Read

Correct Answer: F

QUESTION 12

You are developing a SQL Server Analysis Services (SSAS) tabular project. The model includes a table named Sales. The Sales table includes a single date column.

The Sales table must meet the following requirements:

Queries must be able to return all rows.

Must be able to support four different processing schedules for different date ranges.

Date ranges must not include any overlapping data.

You need to implement a solution that meets the requirements.

What should you do?

- A. Create four partitions for the Sales table. Create four roles. Use the same row filter queries for each role and partition.
- B. Convert the Sales table into four smaller tables by using row filter queries. Use one perspective for all four tables.
- C. Create four partitions for the Sales table. Use row filter queries for each partition.
- D. Convert the Sales table into four smaller tables by using row filter queries. Use one perspective for each of the four tables.

Correct Answer: C

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate


More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



| | | |
|---|---|--|
|  <p>One Year Free Update Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p> |  <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p> |  <p>Security & Privacy We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p> |
|---|---|--|

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.