

70-768.exam.19q

Number: 70-768
Passing Score: 800
Time Limit: 120 min



Website: <https://vceplus.com>
VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>
Facebook: <https://www.facebook.com/VCE.For.All.VN/>
Twitter : https://twitter.com/VCE_Plus

<https://vceplus.com/>

70-768

Developing SQL Data Models

Exam A

QUESTION 1

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:



<https://vceplus.com/>

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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 2

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:

```
AVG({
    COUSIN(
        [Geography].[State-Province].&[CO]&[US],
        [Geography].[State-Province].&[CO]
    )
})
[Measures].[Reseller Average Unit Price])
```

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:



QUESTION 3

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

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 4

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.

A company has an e-commerce website. When a customer places an order, information about the transaction is inserted into tables in a Microsoft SQL Server relational database named OLTP1. The company has a SQL Server Analysis Services (SSAS) instance that is configured to use Tabular mode. SSAS uses data from OLTP1 to populate a data model.

Sales analysts build reports based on the SSAS model. Reports must be able to access data as soon as it is available in the relational database.

You need to configure and deploy an Analysis Services project to the Analysis Services instance that allows near real-time data source access.

Solution: In the Deployment Option property for the report, you set the Query Mode to InMemory.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

With InMemory the queries can use the cache only.

References: [https://msdn.microsoft.com/en-us/library/hh230898\(v=sql.120\).aspx](https://msdn.microsoft.com/en-us/library/hh230898(v=sql.120).aspx)

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.

A company has an e-commerce website. When a customer places an order, information about the transaction is inserted into tables in a Microsoft SQL Server relational database named OLTP1. The company has a SQL Server Analysis Services (SSAS) instance that is configured to use Tabular mode. SSAS uses data from OLTP1 to populate a data model.

Sales analysts build reports based on the SSAS model. Reports must be able to access data as soon as it is available in the relational database.

You need to configure and deploy an Analysis Services project to the Analysis Services instance that allows near real-time data source access.



Solution: In the Deployment Option property for the report, you set the Query Mode to DirectQuery with InMemory.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

With DirectQuerywithInMemory mode the queries use the relational data source by default, unless otherwise specified in the connection string from the client.

References:[https://msdn.microsoft.com/en-us/library/hh230898\(v=sql.120\).aspx](https://msdn.microsoft.com/en-us/library/hh230898(v=sql.120).aspx)

QUESTION 6

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 an existing multidimensional cube that provides sales analysis. The users can slice by date, product, location, customer, and employee.

The management team plans to evaluate sales employee performance relative to sales targets. You identify the following metrics for employees:

- Ninety percent or greater relative to sales target values is considered on target.
- Between 75 percent and 90 percent is considered slightly off target. ▪

Below 75 percent is considered off target.

You need to implement the KPI based on the Status expression.

Solution: You design the following solution:


```
Case
  WHEN ([Measures].[Total Including Tax]) / (SUM([Date].[Calendar Year].Cur-
rentMember.Lag(1), [Measures].[Total Including Tax])) > 0.9
  THEN 1
  WHEN ([Measures].[Total Including Tax]) / (SUM([Date].[Calendar Year].Cur-
rentMember.Lag(1), [Measures].[Total Including Tax])) <= 0.9
  AND
  [Measures].[Total Including Tax]) / (SUM([Date].[Calendar Year].Cur-
rentMember.Lag(1), [Measures].[Total Including Tax])) > 0.74
  THEN 0
  ELSE -1
END
```

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation



Explanation/Reference:

Explanation:

QUESTION 7

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 an existing multidimensional cube that provides sales analysis. The users can slice by date, product, location, customer, and employee.

The management team plans to evaluate sales employee performance relative to sales targets. You identify the following metrics for employees:

- Ninety percent or greater relative to sales target values is considered on target.
- Between 75 percent and 90 percent is considered slightly off target. ▪

Below 75 percent is considered off target.

You need to implement the KPI based on the Status expression.

Solution: You design the following solution:

```
Case
  WHEN KpiValue ("Employee Sales") / KpiGoal("Employee Sales") > .90
  THEN 1
  WHEN KpiValue ("Employee Sales") / KpiGoal("Employee Sales") <= .90
  AND
    KpiValue ("Employee Sales") / KpiGoal("Employee Sales") > .74
  THEN 0
  ELSE -1
END
```



Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 8

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 an existing multidimensional cube that provides sales analysis. The users can slice by date, product, location, customer, and employee.

The management team plans to evaluate sales employee performance relative to sales targets. You identify the following metrics for employees:

- Ninety percent or greater relative to sales target values is considered on target.
- Between 75 percent and 90 percent is considered slightly off target.
- Below 75 percent is considered off target.

You need to implement the KPI based on the Status expression.

Solution: You design the following solution:

```
Case
WHEN KpiValue ("Employee Sales") / KpiGoal("Employee Sales") >= .90
THEN 1
WHEN KpiValue ("Employee Sales") / KpiGoal("Employee Sales") < .90
AND
KpiValue ("Employee Sales") / KpiGoal("Employee Sales") > .74
THEN 0
ELSE -1
END
```

Does the solution meet the goal?

- A. Yes
- B. No



Correct Answer: A

Section: (none)

Explanation

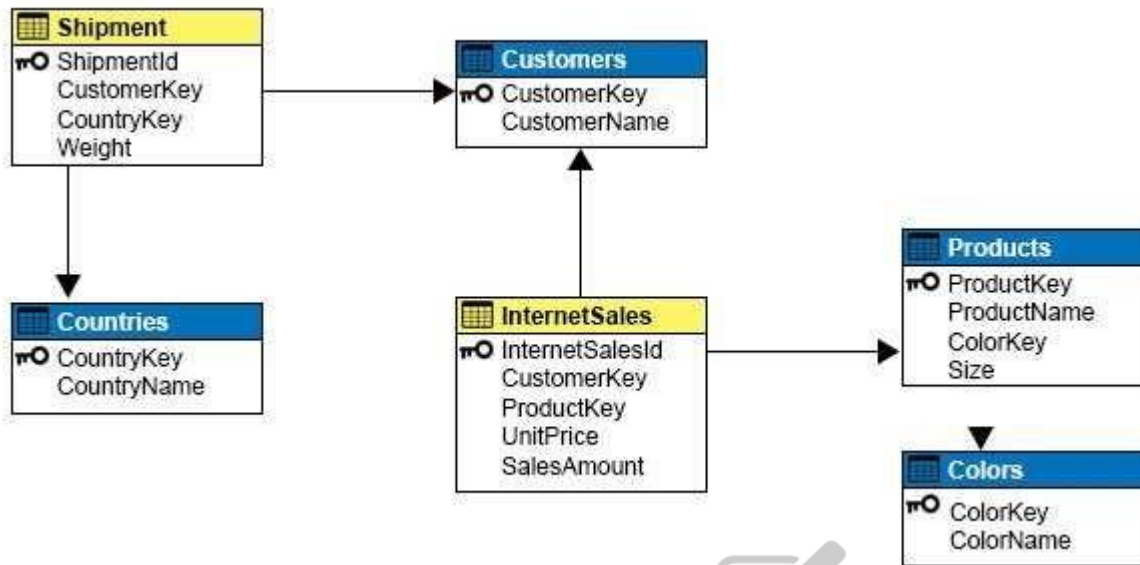
Explanation/Reference:

Explanation:

QUESTION 9

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


Section: (none)

Explanation

Explanation/Reference:

Explanation:

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.

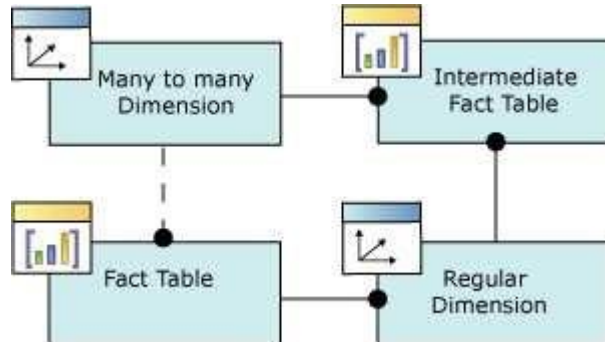
Incorrect Answers:

B: A regular dimension relationship between a cube dimension and a measure group exists when the key column for the dimension is joined directly to the fact table.

C: Fact dimensions, frequently referred to as degenerate dimensions, are standard dimensions that are constructed from attribute columns in fact tables instead of from attribute columns in dimension tables.

E: Many to Many Dimension Relationships.

In most dimensions, each fact joins to one and only one dimension member, and a single dimension member can be associated with multiple facts. In relational database terminology, this is referred to as a one-to-many relationship. However, it is frequently useful to join a single fact to multiple dimension members. For example, a bank customer might have multiple accounts (checking, saving, credit card, and investment accounts), and an account can also have joint or multiple owners. The Customer dimension constructed from such relationships would then have multiple members that relate to a single account transaction.

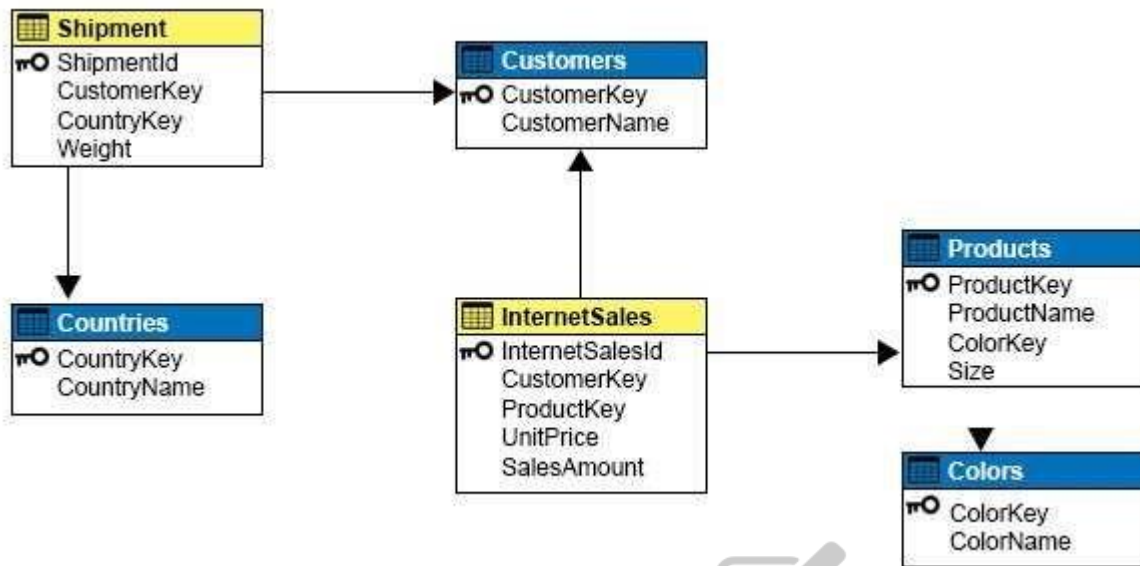


References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models-olap-logical-cube-objects/dimension-relationships>

QUESTION 10

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 color.

You need to create a dimension that contains all of the colors for products sold by the company.

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: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

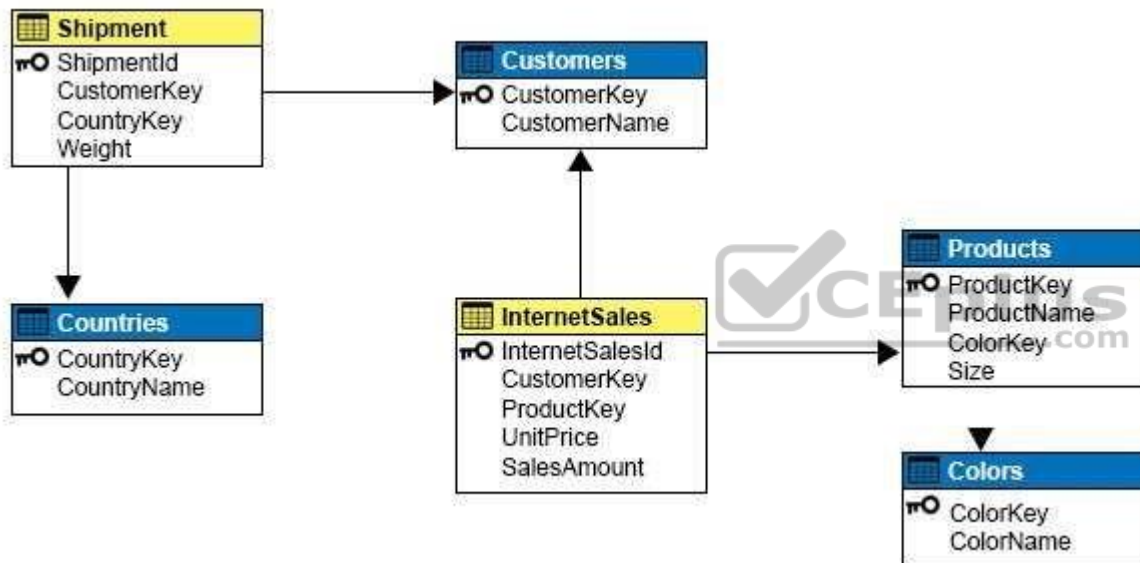
A regular dimension relationship between a cube dimension and a measure group exists when the key column for the dimension is joined directly to the fact table.

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models-olap-logical-cube-objects/dimension-relationships>

QUESTION 11

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 many 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:



You need to create a new dimension that allows users to list shipments by the country where the product is shipped.

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

- A. no relationship
- B. regular
- C. fact
- D. referenced

- E. many-to-many
- F. data mining

Correct Answer: E

Section: (none)

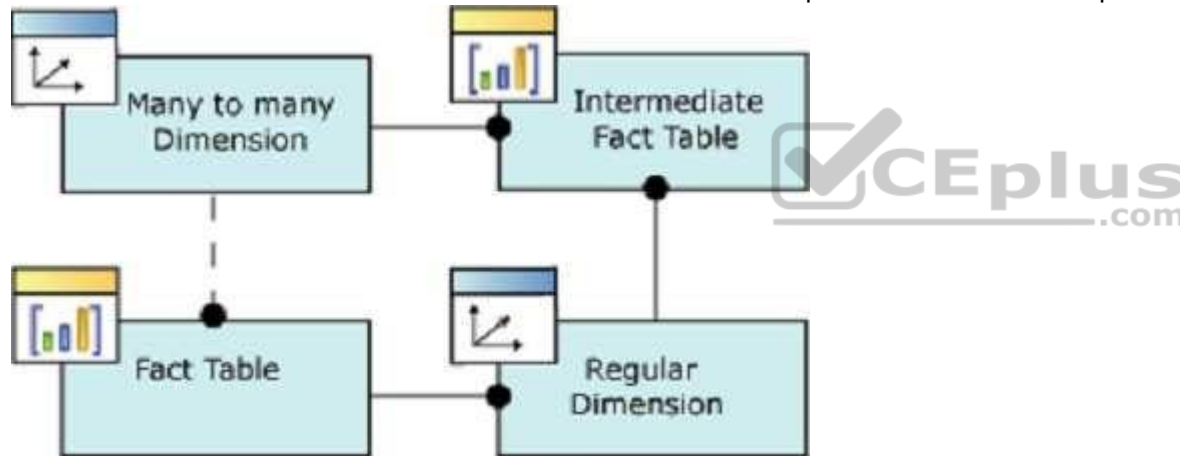
Explanation

Explanation/Reference:

Explanation:

Many to Many Dimension Relationships.

In most dimensions, each fact joins to one and only one dimension member, and a single dimension member can be associated with multiple facts. In relational database terminology, this is referred to as a one-to-many relationship. However, it is frequently useful to join a single fact to multiple dimension members. For example, a bank customer might have multiple accounts (checking, saving, credit card, and investment accounts), and an account can also have joint or multiple owners. The Customer dimension constructed from such relationships would then have multiple members that relate to a single account transaction.



References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models-olap-logical-cube-objects/dimension-relationships>

QUESTION 12

You are a business analyst for a retail company that uses a Microsoft SQL Server Analysis Services (SSAS) multidimensional database for reporting. The database contains the following objects:

Type	Name	Content
Measure	Internet Sales Amount	online sales data
Dimension	Date	the date of sales
Hierarchy	Date.Calendar.Calendar Year	the calendar year of the sale
Hierarchy	Date.Calendar.Month	the month of the sale

You must create a report that shows, for each month, the Internet sales for that month and the total Internet sales for the calendar year up to and including the current month.

You create the following MDX statement (Line numbers are included for reference only.):

```
01
02 SELECT
03   {[Measures].[Internet Sales Amount]}, [Measures].[Goal]} on 0,
04   {[Date].[Calendar].[Month].Members} on 1
05 FROM [Adventure Works];
```

You need to complete the MDX statement to return data for the report.

Which MDX segment should you use in line 01?

```
WITH MEMBER [Measures].[Goal] AS
    Aggregate(
        QTD([Date].[Calendar].CurrentMember),
        ([Measures].[Internet Sales Amount]))
```

```
WITH MEMBER [Measures].[Goal] AS
    Aggregate(PERIODSTODATE(
        [Date].[Calendar].[Calendar Year],
        [Date].[Calendar].CurrentMember
    ), ([Measures].[Internet Sales Amount]))
```

A.

B.

```
WITH MEMBER [Measures].[Goal] AS  
    Aggregate(  
        MTD([Date].[Calendar].CurrentMember),  
        ([Measures].[Internet Sales Amount]))
```

```
WITH MEMBER [Measures].[Goal] AS  
    Aggregate(OPENINGPERIOD(  
        [Date].[Calendar].[Calendar Year],  
        [Date].[Calendar].CurrentMember  
    ), ([Measures].[Internet Sales Amount]))
```

C.

D.



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The following example returns the sum of the Measures. [Order Quantity] member, aggregated over the first eight months of calendar year 2003 that are contained in the Date dimension, from the Adventure Works cube.

Copy

```
WITH MEMBER [Date].[Calendar].[First8Months2003] AS
```

```
Aggregate(  
  PeriodsToDate(  
    [Date].[Calendar].[Calendar Year],  
    [Date].[Calendar].[Month].[August 2003]  
  )  
)  
SELECT  
  [Date].[Calendar].[First8Months2003] ON COLUMNS,  
  [Product].[Category].Children ON ROWS  
FROM  
  [Adventure Works]  
WHERE  
  [Measures].[Order Quantity]
```

References: <https://docs.microsoft.com/en-us/sql/mdx/aggregate-mdx>

QUESTION 13

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.

- Project1 uses the tabular data model.
- Project2 uses SQL Server data mining to predict customer-purchasing intentions by using the Decision Trees algorithm.

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

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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 14

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 administer a Microsoft SQL Server Analysis Services (SSAS) tabular model for a travel agency that specializes in vacation packages. Vacation bookings and packages are stored in a SQL Server database. You use the model as the basis for customer emails that highlight vacation packages that are currently underbooked, or projected to be underbooked.

The company plans to incorporate cruise ship vacation packages. Cruise ship vacation packages include new features such as region availability and cruise line specialties that require changes to the tabular model.

You must ensure that the tabular model reflects the new vacation packages.

You need to configure the tabular data model.

What should you do?

- A. Ensure that DirectQuery is enabled for the model.
- B. Ensure that DirectQuery is disabled for the model.
- C. Ensure that the Transactional Deployment property is set to True.
- D. Ensure that the Transactional Deployment property is set to False.
- E. Process the model in Process Full mode.
- F. Process the model in Process Data mode.
- G. Process the model in Process Defrag mode.

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Process Full processes an Analysis Services object and all the objects that it contains. 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.

Incorrect Answers:

A, B: DirectQuery is just used to make queries run faster.

F: Process Data: 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.

G: There is no processing option named Process Defrag.

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models/processing-options-and-settings-analysis-services>

QUESTION 15

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 administer a Microsoft SQL Server Analysis Services (SSAS) tabular model for a retail company. The model is the basis for reports on inventory levels, popular products, and regional store performance.

The company recently split up into multiple companies based on product lines. Each company starts with a copy of the database and tabular model that contains data for a specific product line.

You need to optimize performance of queries that use the copied tabular models while minimizing downtime.

What should you do?

- A. Ensure that DirectQuery is enabled for the model.
- B. Ensure that DirectQuery is disabled for the model.
- C. Ensure that the Transactional Deployment property is set to True.
- D. Ensure that the Transactional Deployment property is set to False.
- E. Process the model in Process Full mode.
- F. Process the model in Process Data mode.
- G. Process the model in Process Defrag mode.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Transactional Deployment setting controls whether the deployment of metadata changes and process commands occurs in a single transaction or in separate transactions.

If this option is True (default), Analysis Services deploys all metadata changes and all process commands within a single transaction.

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

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

QUESTION 16

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 are developing Microsoft SQL Server Analysis Services (SSAS) tabular model.

The model must meet the following requirements:

- Support the RELATED and TRUNC DAX functions. ▪

Show updated data as quickly as possible.

You need to configure model.

What should you do?

- A. Ensure that DirectQuery is enabled for the model.
- B. Ensure that DirectQuery is disabled for the model.
- C. Ensure that the Transactional Deployment property is set to True.
- D. Ensure that the Transactional Deployment property is set to False.
- E. Process the model in Process Full mode.
- F. Process the model in Process Data mode.
- G. Process the model in Process Defrag mode.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:



DAX originally emerged from a Power Pivot add-in for Excel, as a formula language extension for creating calculated columns and measures for data analysis (which are also integral elements of one SSAS Tabular model database, too), but when Microsoft added support for DAX queries in SQL Server 2012, BI experts started “daxing” data from Tabular model databases.

That trend continues, because of simplicity and fast query execution (related to DirectQuery mode in SSAS Tabular).

References:<https://www.sqlshack.com/query-ssas-tabular-model-database-using-dax-functions/>

QUESTION 17

You are optimizing a Microsoft SQL Server Analysis Services (SSAS) multidimensional model over a SQL Server database. You have a table named City which has several dimensions that do not contain a space in their names. One dimension is named SalesTerritory rather than Sales Territory.

You need to ensure that Report developers can drag the attribute name to the report rather than having to re-label the attributes by implementing spaces. You must minimize administrative effort and not break any upstream processes.

What should you do?

- A. In the SQL Server database, run the system procedure sp_rename to rename the columns in the base tables with the target name.
- B. In SQL Server Management Studio, navigate to the City table, expand the columns, press F2, and rename the columns in the base tables.
- C. In the SQL Server database, implement a SYNONYM.
- D. In the SQL Server database, implement a view over the City table that aliases the columns in the tables.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 18

Case Study #1

This is a case study. Case studies are not limited separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Background

Wide World Importers imports and sells clothing. The company has a multidimensional Microsoft SQL Server Analysis Services instance. The server has 80 gigabytes (GB) of available physical memory. The following installed services are running on the server:

SQL Server Database Engine

SQL Server Analysis Services (multidimensional)

The database engine instance has been configured for a hard cap of 50 GB, and it cannot be lowered. The instance contains the following cubes: SalesAnalysis, OrderAnalysis.

Reports that are generated based on data from the OrderAnalysis cube take more time to complete when they are generated in the afternoon each day. You examine the server and observe that it is under significant memory pressure.

Processing for all cubes must occur automatically in increments. You create one job to process the cubes and another job to process the dimensions. You must configure a processing task for each job that optimizes performance. As the cubes grown in size, the overnight processing of the cubes often do not complete during the allowed maintenance time window.

SalesAnalysis

The SalesAnalysis cube is currently being tested before being used in production. Users report that day name attribute values are sorted alphabetically. Day name attribute values must be sorted chronologically. Users report that they are unable to query the cube while any cube processing operations are in progress. You need to maximize data availability during cube processing and ensure that you process both dimensions and measures.

OrderAnalysis

The OrderAnalysis cube is used for reporting and ad-hoc queries from Microsoft Excel. The data warehouse team adds a new table named Fact.Transaction to the cube. The Fact.Transaction table includes a column named Total Including Tax. You must add a new measure named Transactions – Total Including Tax to the cube. The measure must be calculated as the sum of the Total Including Tax column across any selected relevant dimensions.

Finance

The Finance cube is used to analyze General Ledger entries for the company.

Requirements

You must minimize the time that it takes to process cubes while meeting the following requirements:

- The Sales cube requires overnight processing of dimensions, cubes, measure groups, and partitions.
- The OrderAnalysis cube requires overnight processing of dimensions only. ▪

The Finance cube requires overnight processing of dimensions only.

You need to configure the server to optimize the afternoon report generation based on the OrderAnalysis cube.

Which property should you configure?

- A. LowMemoryLimit
- B. VertiPagingPolicy
- C. TotalMemoryLimit
- D. VirtualMemoryLimit

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

LowMemoryLimit: For multidimensional instances, a lower threshold at which the server first begins releasing memory allocated to infrequently used objects.

From scenario: Reports that are generated based on data from the OrderAnalysis cube take more time to complete when they are generated in the afternoon each day. You examine the server and observe that it is under significant memory pressure.

Incorrect Answers:

B: VertiPagingPolicy: For tabular instances, a lower threshold at which the server first begins releasing memory allocated to infrequently used objects.

Scenario: SQL Server Analysis Services (multidimensional) is used.

C: TotalMemoryLimit: An upper threshold at which Analysis Services begins releasing memory more aggressively to make room for requests that are in execution as well as new high priority requests.

Scenario: The database engine instance has been configured for a hard cap of 50 GB, and it cannot be lowered.

References: <https://docs.microsoft.com/en-us/sql/analysis-services/server-properties/memory-properties>

QUESTION 19

DRAG DROP

Case Study #2

This is a case study. Case studies are not limited separately. You can use as much exam time as you would like to complete each case. However, there may be additional case studies and sections on this exam. You must manage your time to ensure that you are able to complete all questions included on this exam in the time provided.

To answer the questions included in a case study, you will need to reference information that is provided in the case study. Case studies might contain exhibits and other resources that provide more information about the scenario that is described in the case study. Each question is independent of the other question on this case study.

At the end of this case study, a review screen will appear. This screen allows you to review your answers and to make changes before you move to the next sections of the exam. After you begin a new section, you cannot return to this section.

To start the case study

To display the first question on this case study, click the Next button. Use the buttons in the left pane to explore the content of the case study before you answer the questions. Clicking these buttons displays information such as business requirements, existing environment, and problem statements. If the case study has an All Information tab, note that the information displayed is identical to the information displayed on the subsequent tabs. When you are ready to answer a question, click the Question button to return to the question.

Background

Wide World Importers has multidimensional cubes named SalesAnalysis and ProductSales. The SalesAnalysis cube is refreshed from a relational data warehouse. You have a Microsoft SQL Server Analysis Services instance that is configured to use tabular mode. You have a tabular data model named CustomerAnalysis.

Sales Analysis

The SalesAnalysis cube contains a fact table named CoffeeSale loaded from a table named FactSale in the data warehouse. The time granularity within the cube is 15 minutes. The cube is processed every night at 23:00. You determine that the fact table cannot be fully processed in the expected time. Users have reported slow query response times.

The SalesAnalysis model contains tables from a SQL Server database named SalesDB. You set the DirectQueryMode option to DirectQuery. Data analyst access data from a cache that is up to 24 hours old. Data analyst report performance issues when they access the SalesAnalysis model.

When analyzing sales by customer, the total of all sales is shown for every customer, instead of the customer's sales value. When analyzing sales by product, the correct totals for each product are shown.

Customer Analysis

You are redesigning the CustomerAnalysis tabular data model that will be used to analyze customer sales. You plan to add a table named CustomerPermission to the model. This table maps the Active Directory login of an employee with the CustomerId keys for all customers that the employee manages.

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.

Data analysts must be able to analyze sales for financial years, financial quarters, months, and days. Many reports are based on analyzing sales by month.

Product Sales

The ProductSales cube allows data analysts to view sales information by product, city, and time. Data analysts must be able to view ProductSales data by Year to Date (YTD) as a measure. The measure must be formatted as currency, associated with the Sales measure group, and contained in a folder named Calculations.

Requirements

You identify the following requirements:

- Data available during normal business hours must always be up-to-date.
- Processing overhead must be minimized.
- Query response times must improve.
- All queries that access the SalesAnalysis model must use cached data by default. ▪

Data analysts must be able to access data in near real time.

You need to configure the CoffeeSale fact table environment.

Which four 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.

NOTE: More than one order of answer choices is correct. You will receive credit for any of the correct orders you select.

Select and Place:

Actions

Set the storage mode for the latest partition to ROLAP, and set the storage mode for all other partitions to MOLAP.

Alter the processing job to run every half during the day.

Alter the client application that queries the cube to query the dimensional data warehouse directly for current day data.

Set the storage mode for all partitions to ROLAP.

Test that the cube meets the functional requirement for data currency and query performance.

Partition the CoffeeSale fact table.

Set the storage mode for all partitions to HOLAP.

Alter the processing job to ensure that it rearranges the partition structure each evening.

Answer Area



Correct Answer:

Actions

Set the storage mode for the latest partition to ROLAP, and set the storage mode for all other partitions to MOLAP.

Alter the processing job to run every half during the day.

Alter the client application that queries the cube to query the dimensional data warehouse directly for current day data.

Set the storage mode for all partitions to ROLAP.

Answer Area

Partition the CoffeSale fact table.

Set the storage mode for all partitions to HOLAP.

Alter the processing job to ensure that it rearranges the partition structure each evening.

Test that the cube meets the functional requirement for data currency and query performance.

Section: (none)

Explanation

Explanation/Reference:

Step 1: Partition the CoffeSale facto table.

Step 2: Set the storage mode for all partitions to HOLAP.

Partitions stored as HOLAP are smaller than the equivalent MOLAP partitions because they do not contain source data and respond faster than ROLAP partitions for queries involving summary data.

Step 3: Alter the processing job to ensure that it rearranges the partition structure each evening.

Step 4: Test that the cube meets the functional requirement for data currency and query performance.

From scenario:

Data analysts must be able to analyze sales for financial years, financial quarters, months, and days. Many reports are based on analyzing sales by month.

The SalesAnalysis cube contains a fact table named CoffeeSale loaded from a table named FactSale in the data warehouse. The time granularity within the cube is 15 minutes. The cube is processed every night at 23:00. You determine that the fact table cannot be fully processed in the expected time. Users have reported slow query response times.

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models-olap-logical-cube-objects/partitions-partition-storage-modes-andprocessing>



<https://vceplus.com/>