



## Microsoft

### Exam Questions 70-768

Developing SQL Data Models (beta)

**NEW QUESTION 1**

DRAG DROP

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.

Processing tasks	Answer Area	
	Job	Processing task
Process Clear	Incremental cube processing	Processing task
Process Update	Incremental dimension processing	Processing task
Process Index		
Process Add		
Process Data		
Process Structure		

**Answer:**

**Explanation:** 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/multidimensional-models/processing-options-and-settings-analysis-services>

Topic 2, Case Study #2

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.

**NEW QUESTION 2**

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.

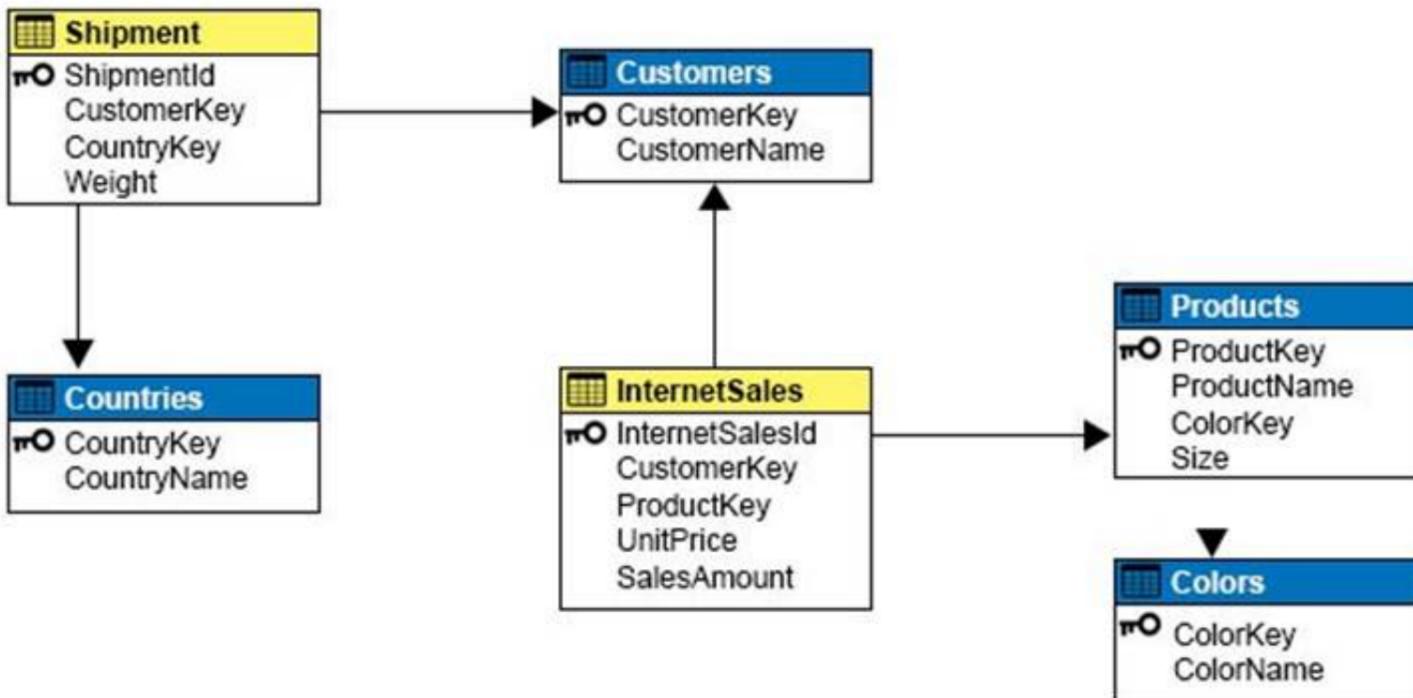
- C. Ensure that the Transactional Deployment property is set to True.
- D. Ensure that the Transactional Deployment property is set to False.
- E. Process the model in Process Full mode.
- F. Process the model in Process Data mode.
- G. Process the model in Process Defrag mode.

**Answer:** E

**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.

**NEW QUESTION 3**

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

**Answer:** B

**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>

**NEW 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.

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

Answer: B

**NEW QUESTION 5**

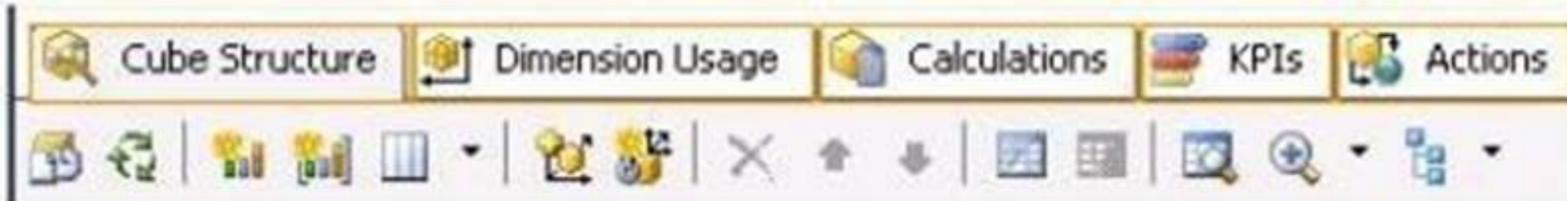
**HOTSPOT**

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

Revenue must be compared to a goal and described by a status and a trend. Revenue, goal, status, and trend will be defined by Multidimensional Expressions (MDX) expressions.

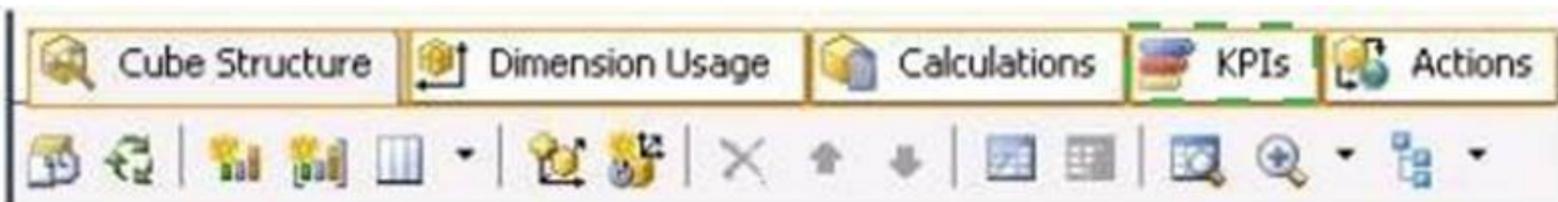
You need to add the Revenue indicator.

Which tab should you select? (To answer, select the appropriate tab in the work area.)



Answer:

Explanation:



**NEW QUESTION 6**

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?

- A. [MISSING]
- B. [MISSING]
- C. [MISSING]
- D. [MISSING]

**Answer:** B

**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
```

**NEW QUESTION 7**

HOTSPOT

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.

**Answer Area**

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

**Answer:**

**Explanation:** 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/multidimensional-models/define-named-queries-in-a-data-source-view-analysis-services>

**NEW 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.

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

**Answer:** A

**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)

**NEW QUESTION 9**

You are developing a SQL Server Analysis Services (SSAS) tabular project that will be used by the finance, sales, and marketing teams. The sales team reports that the model is too complex and difficult to use. The sales team does not need any information other than sales-related resources in the

tabular model. The finance and marketing teams need to see all the resources in the tabular model. You need to implement a solution that meets the needs of the sales team while minimizing development and administrative effort. What should you do?

- A. Create a separate partition for each team.
- B. Create a separate data source for each team.
- C. Create a perspective for the sales team.
- D. Enable client side security to filter non-sales data.

**Answer: C**

**NEW QUESTION 10**

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

**Answer: F**

**NEW QUESTION 10**

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 with DirectQuery.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:** With InMemory with DirectQuery: Queries use the cache 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)

**NEW QUESTION 11**

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("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

**Answer:** A

#### NEW 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 are developing Microsoft SQL Server Analysis Services (SSAS) tabular model. The model must meet the following requirements: 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.

**Answer:** A

**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/>

#### NEW QUESTION 17

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 deploy a tabular data model to an instance of Microsoft SQL Server Analysis Services (SSAS). The model uses an in-memory cache to store and query data. The data set is already the same size as the available RAM on the server. Data volumes are likely to continue to increase rapidly.

Your data model contains multiple calculated tables.

The data model must begin processing each day at 2:00 and processing should be complete by 4:00 the same day. You observe that the data processing operation often does not complete before 7:00. This is adversely affecting team members.

You need to improve the performance.

Solution: Install solid-state disk drives to store the tabular data model. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** B

**Explanation:** By default, tabular models use an in-memory cache to store and query data. When tabular models query data residing in-memory, even complex queries can be incredibly fast. However, there are some limitations to using cached data. Namely, large data sets can exceed available memory, and data freshness requirements can be difficult if not impossible to achieve on a regular processing schedule.

DirectQuery overcomes these limitations while also leveraging RDBMS features making query execution more efficient.

With DirectQuery: +

References:<https://docs.microsoft.com/en-us/sql/analysis-services/tabular-models/directquery-mode-ssas-tabular>

#### NEW QUESTION 20

DRAG DROP

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

Type	Name	Content
Measure	Reseller Sales Amount	the total sales made by a reseller
Dimension	Geography	the location of the reseller
Hierarchy	Geography.City	the city where the reseller is located
Member	Geography.City.&[London]&[UK], Geography.City.&[Tokyo]&[JP]	a specific city and region

Your company is developing a promotional plaque to recognize the top resellers in the top 10 cities where the company does business. Each plaque must display the sales total for all resellers in the city. In addition, the plaque must display a total for all cities not in the top 10.

You have the following requirements:

You need to provide the information needed for the promotional plaques.

How should you complete the MDX statement? To answer, drag the appropriate MDX segments 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.

NOTE: Each correct selection is worth one point.

**MDX segments**

- MEMBER
- DYNAMIC SET
- [Geography].[City].CURRENTMEMBER
- [Geography].[City].[City].members
- [Measures].[Reseller Sales Amount]

**Answer Area**

```

WITH [MDX segment] [Top 10] AS
    TOPCOUNT([Geography].[City].[City].members, 10,
    [Measures].[Reseller Sales Amount])
[MDX segment] [Geography].[City].[Others] AS
    Aggregate(Except([Geography].[City].[City].members, [Top 10]))
[MDX segment] [ALL] AS
    {[Top 10], [Geography].[City].[Others] }
[MDX segment] [Measures].[Rank] AS
    RANK([MDX segment], [All])
SELECT {[Measures].[Reseller Sales Amount], [Measure].[Rank]} ON 0, [All] on 1
FROM [AdventureWorks]
    
```

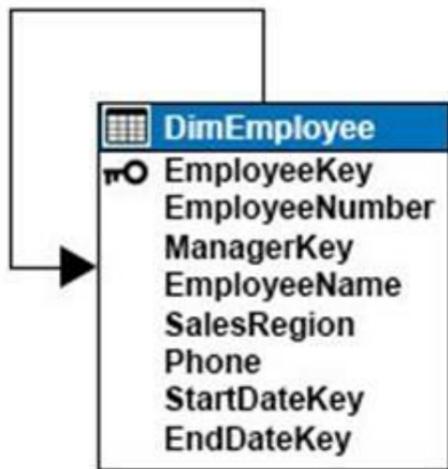
**Answer:**

**Explanation:** Box 1:DYNAMIC SET Box 2:MEMBER  
 Box 3:DYNAMIC SET  
 Box 4:[Geography].[City].[City].members Box 5:[Measures].[Reseller Sales Amount]  
 References: <https://docs.microsoft.com/en-us/sql/mdx/aggregate-mdx>

**NEW QUESTION 21**

HOTSPOT

You have a Microsoft SQL Server Analysis Services (SSAS) multidimensional project. You are developing a dimension that uses data from the following table:



The ManagerKey column defines a foreign key constraint that references the EmployeeKey column. The table stores employee history information by using slowly changing dimensions (SCD). Changes to EmployeeName, Phone, or ManagerKey are managed as SCD Type 1 changes. Changes to SalesRegion are managed as SCD Type 2 changes.

You create the following attributes, and set the KeyColumns and NameColumn properties to the columns listed in the table below:

Attribute	KeyColumns	NameColumn
Employee	EmployeeKey	EmployeeName
Employee Number	EmployeeNumber	
Phone	Phone	
Manager	ManagerKey	
Sales Region	SalesRegion	

You need to add a parent-child hierarchy to the dimension to enable navigating the organization hierarchy. In the table below, identify the attribute that you must use for each attribute usage type. NOTE: Make only one selection in each column.

## Answer Area

Attribute	Key	Parent
Employee	<input type="radio"/>	<input type="radio"/>
Employee Number	<input type="radio"/>	<input type="radio"/>
Manager	<input type="radio"/>	<input type="radio"/>
Phone	<input type="radio"/>	<input type="radio"/>
Sales Region	<input type="radio"/>	<input type="radio"/>

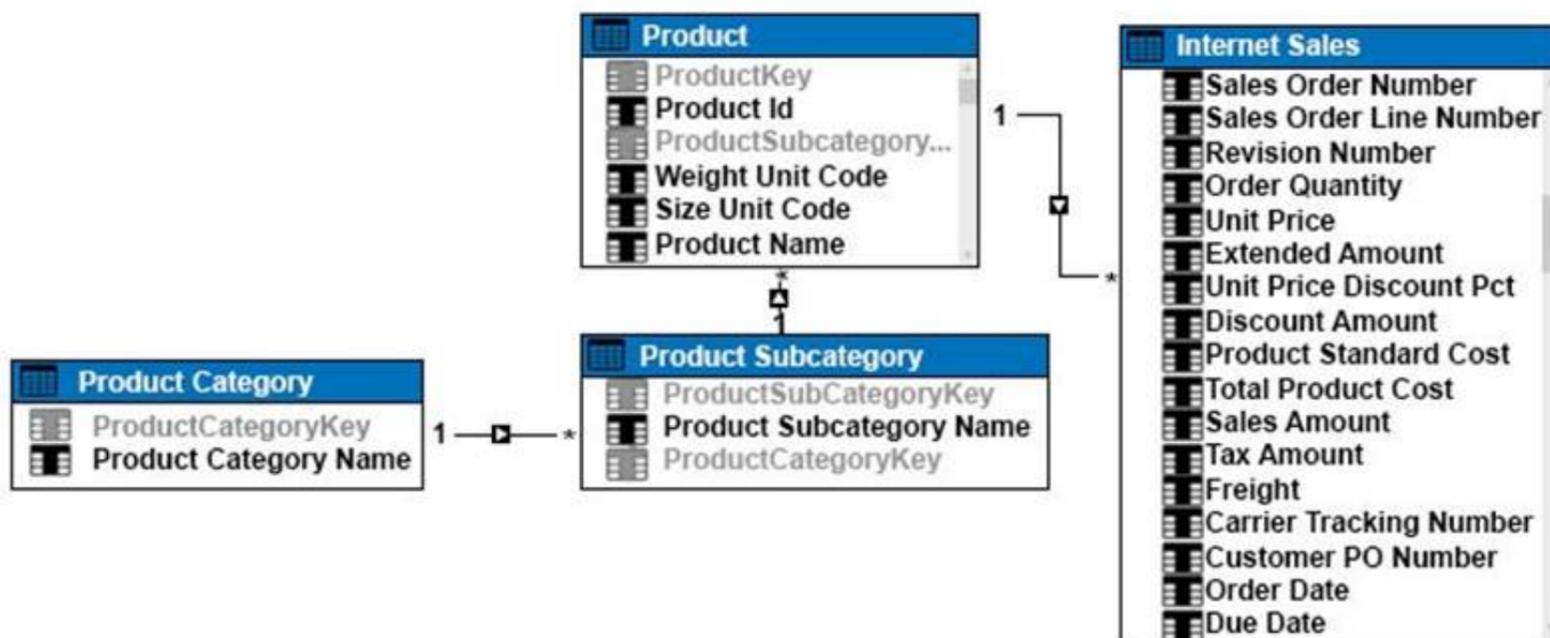
Answer:

**Explanation:** The ManagerKey column, the Manager attribute, defines a foreign key constraint that references the EmployeeKey column, the Employee attribute.

### NEW QUESTION 23

DRAG DROP

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



You have been asked to write a query for a report that returns the total sales for each product subcategory, as well as for each product category. You need to write the query to return the data for the report.

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

### MDX segments

order by

evaluate

summarize

ROLLUP

SUM

'Product Subcategory' [Product Subcategory Name]

'Product Category' [Product Category Name]

### Answer Area

DAX segment

(

DAX segment

(

'Internet Sales',

DAX segment

( DAX segment ),

'Product Category' [Product Category Name],

"Total Sales Amount", SUM('Internet Sales' [Sales Amount])

)

)

**Answer:**

**Explanation:** Box 1:EVALUATE Box 2:SUMMERIZE Box 3:ROLLUP

Box 4:'Product Subcategory' [Product Subcategory Name]

Note: The behavior of SUMMARIZE is similar to the GROUP BY syntax of a SELECT statement in SQL. For example, consider the following query.

```
EVALUATE SUMMARIZE(
'Internet Sales',
'Internet Sales'[Order Date],
"Sales Amount", SUM( 'Internet Sales'[Sales Amount] )
)
```

This query calculates the total of Sales Amount for each date in which there is at least one order, producing this result. References:

**NEW QUESTION 26**

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 tabl
- B. Create four role
- C. Use the same row filter queries for each role and partition.
- D. Convert the Sales table into four smaller tables by using row filter querie
- E. Use one perspective for all four tables.
- F. Create four partitions for the Sales tabl
- G. Use row filter queries for each partition.
- H. Convert the Sales table into four smaller tables by using row filter querie
- I. Use one perspective for each of the four tables.

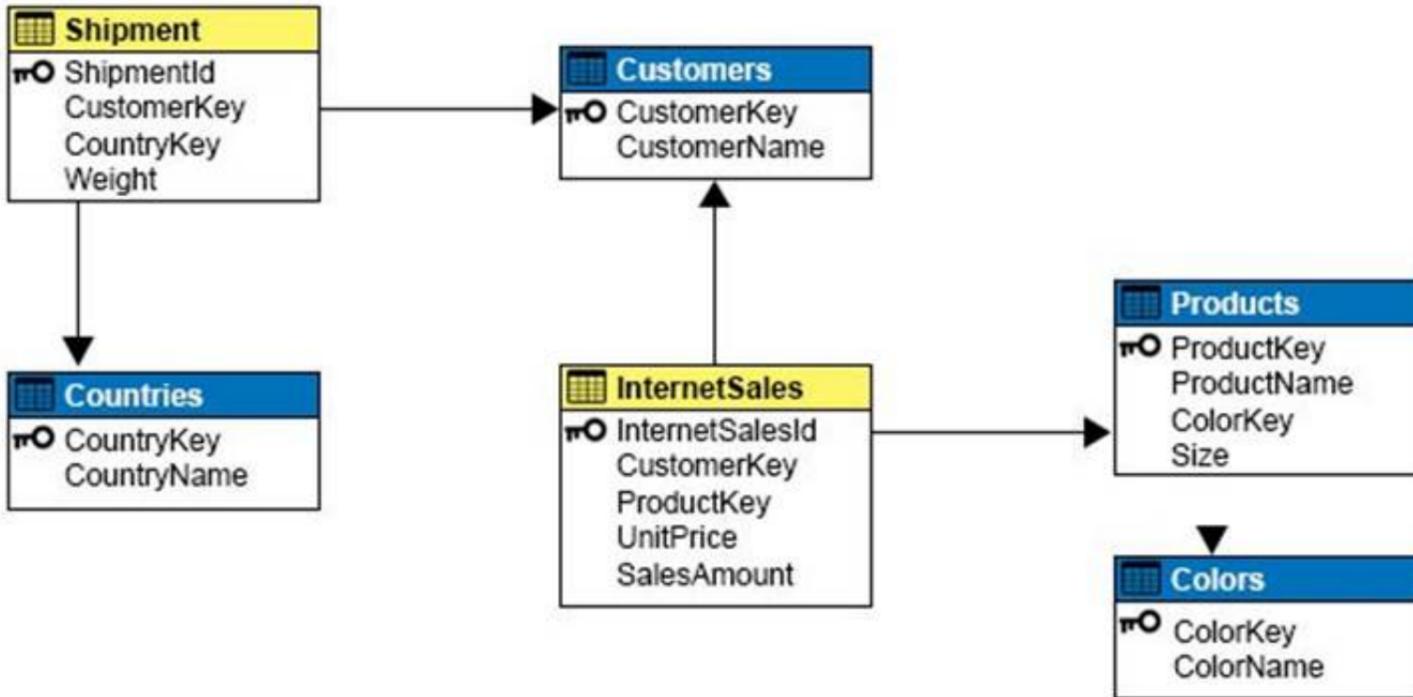
**Answer: C**

**NEW QUESTION 29**

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:

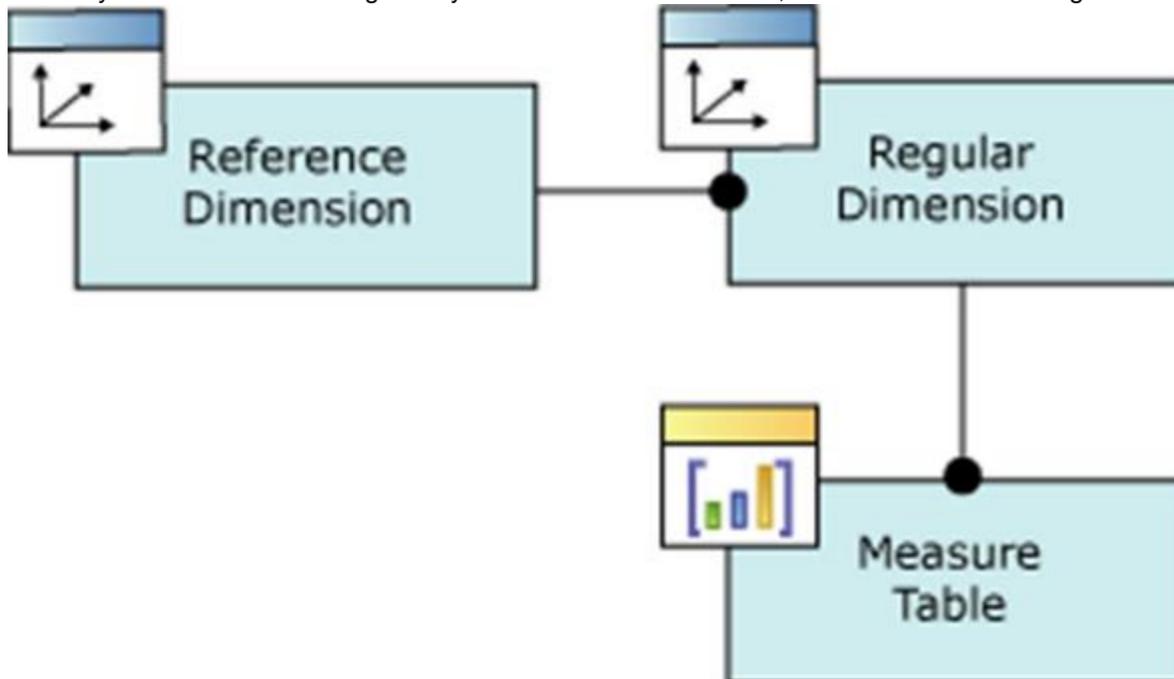


Users need to be able to analyze sales by product and color. You need to create the dimension. Which relationship type should you use between the InternetSales table and the new dimension?

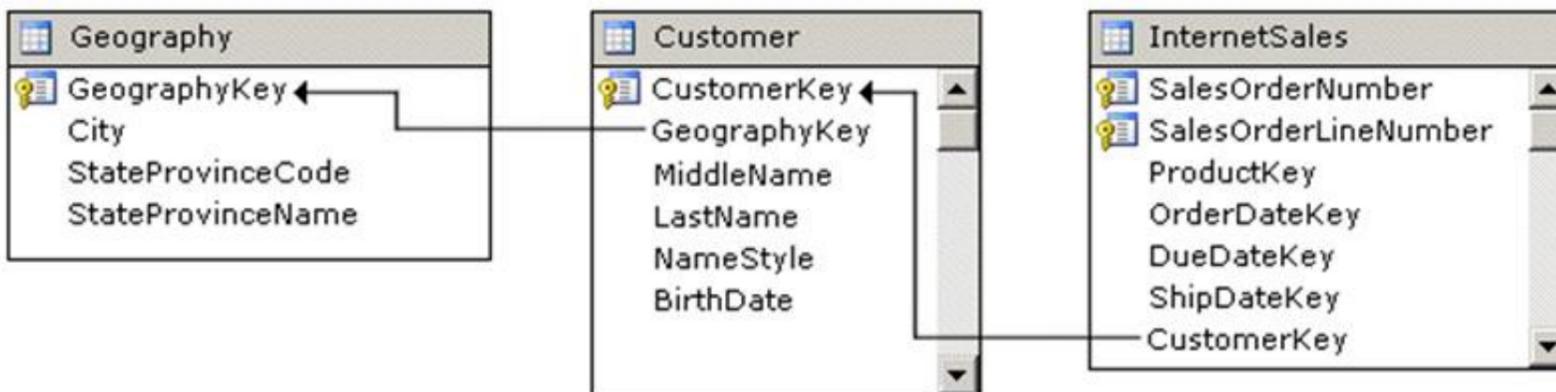
- A. no relationship
- B. regular
- C. fact
- D. referenced
- E. many-to-many
- F. data mining

**Answer:** D

**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.

**NEW QUESTION 33**

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.

**Answer:** C

**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>

### NEW QUESTION 38

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])
- D. =VALUES(Geography[City])
- E. City:=VALUES(Geography[City])

**Answer:** C

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

### NEW QUESTION 39

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 deploy a tabular data model to an instance of Microsoft SQL Server Analysis Services (SSAS). The model uses an in-memory cache to store and query data. The data set is already the same size as the available RAM on the server. Data volumes are likely to continue to increase rapidly. Your data model contains multiple calculated tables. The data model must begin processing each day at 2:00 and processing should be complete by 4:00 the same day. You observe that the data processing operation often does not complete before 7:00. This is adversely affecting team members. You need to improve the performance. Solution: Change the storage mode for the data model to DirectQuery. Does the solution meet the goal?

- A. Yes
- B. No

**Answer:** A

**Explanation:** By default, tabular models use an in-memory cache to store and query data. When tabular models query data residing in-memory, even complex queries can be incredibly fast. However, there are some limitations to using cached data. Namely, large data sets can exceed available memory, and data freshness requirements can be difficult if not impossible to achieve on a regular processing schedule. DirectQuery overcomes these limitations while also leveraging RDBMS features making query execution more efficient. With DirectQuery: + Data is up-to-date, and there is no extra management overhead of having to maintain a separate copy of the data (in the in-memory cache). Changes to the underlying source data can be immediately reflected in queries against the data model. Datasets can be larger than the memory capacity of an Analysis Services server. Etc. References: <https://docs.microsoft.com/en-us/sql/analysis-services/tabular-models/directquery-mode-ssas-tabular>

### NEW QUESTION 42

You are building a Microsoft SQL Server Analysis Services multidimensional model over a SQL Server database. In a cube named OrderAnalysis, there is a standard cube dimension named Stock Item. This dimension has the following attributes: Users report that the attributes Stock Item Key and Photo are distracting and are not providing any value. They have asked for the attributes to be removed. However, these attributes are needed by other cubes. You need to hide the specified attributes from the end users of the OrderAnalysis cube. You do not want to change the structure of the dimension. Which change should you make to the properties for the Stock Item Key and Photo attributes?

- A. Set the AttributeHierarchyVisible property to False.
- B. Set the AttributeHierarchyEnabledproperty to False.
- C. Set the AttributeVisibility property to Hidden.

- D. Set the Usage property to Regular.
- E. Set the AttributeHierarchyDisplayFolder property to Hidden.

**Answer:** A

**Explanation:** The value of the AttributeHierarchyEnabled property determines whether an attribute hierarchy is created. If this property is set to False, the attribute hierarchy is not created and the attribute cannot be used as a level in a user hierarchy; the attribute hierarchy exists as a member property only. However, a disabled attribute hierarchy can still be used to order the members of another attribute. If the value of the AttributeHierarchyEnabled property is set to True, the value of the AttributeHierarchyVisible property determines whether the attribute hierarchy is visible independent of its use in a user-defined hierarchy.  
References:[https://technet.microsoft.com/en-us/library/ms166717\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/ms166717(v=sql.110).aspx)

#### NEW QUESTION 43

.....

## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### 70-768 Practice Exam Features:

- \* 70-768 Questions and Answers Updated Frequently
- \* 70-768 Practice Questions Verified by Expert Senior Certified Staff
- \* 70-768 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* 70-768 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
[Order The 70-768 Practice Test Here](#)