

70-768 Dumps

Developing SQL Data Models (beta)

<https://www.certleader.com/70-768-dumps.html>



NEW QUESTION 1

DRAG DROP - (Topic 2)

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.

Actions	Answer Area
Set the storage mode for the latest partition to ROLAP, and set the storage mode for all other partitions to MOLAP.	<div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;">   </div> <div style="text-align: center;">   </div> </div>
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.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Partition the CoffeeSale fact 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-and-processing>

NEW QUESTION 2

HOTSPOT - (Topic 2)

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.

Answer Area

Location	Setting
Processing option	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #cccccc; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;">Default</div> <div style="padding: 2px;">Do not process</div> <div style="padding: 2px;">Full</div> </div>
Transactional deployment	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #cccccc; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;">False</div> <div style="padding: 2px;">True</div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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/multidimensional-models/deployment-script-files-specifying-processing-options>

NEW QUESTION 3

DRAG DROP - (Topic 3)

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 reconfigure the SSAS instance that hosts DB1.

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

Actions

Set the mode for the FactInternetSales table's partition to **InMemoryWithDirectQuery**.

Set the default mode for the data model to **DirectQuery**.

Set the mode for the FactInternetSales table's partition to **DirectQueryOnly**.

Run **Process Full** for the FactInternetSales partition.

Set the default mode for the data model to **Import**.

Run **Process Clear** for the FactInternetSales partition.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Set the default mode for the data model to DirectQuery.

You discover that the project has been deployed with the Direct Query Mode option set to OFF.

Step 2: Set the mode for the FactInternetSales table's partition to DirectQueryOnly. Initially, even DirectQuery models are always created in memory. The default query mode for the workspace database is also set to DirectQuery with In-Memory. This hybrid working mode lets you use the cache of imported data for improved performance during the model design process, while validating the model against DirectQuery requirements.

From Scenario: Most queries that use the SalesAnalysis data model use data from a table named FactInternetSales that is 20 gigabyte (GB) in size. Cached data must be available for the FactInternetSales table. All queries accessing the SalesAnalysis model must be executed in near real time.

Step 3: Run Process Full for the FactInternetSales partition.

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 4

- (Topic 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 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:

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

Answer: A

NEW QUESTION 5

- (Topic 4)

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 6

DRAG DROP - (Topic 4)

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.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:WITH MEMBER...

Box 2:WHERE ([Date].[Year].[CY2016])

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models/mdx/working-with-members-tuples-and-sets-mdx>

NEW QUESTION 7

- (Topic 4)

You are developing a tabular Business Intelligence Semantic Model (BISM) database based on a SQL Server database.

In the data source, the FactInternetSales table is partitioned by month. Data from the current month has been updated and new data has been inserted in the FactInternetSales table, in the DimProduct table, and in the DimCustomer table.

In the model, the FactInternetSales table is also partitioned by month.

You need to ensure that the model has the most recent data while minimizing the processing time. What should you do?

- A. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Clear processing option.
- B. Then process the database with the Process Data processing option.
- C. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Clear processing option.
- D. Then process the database with the Process Full processing option.
- E. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Defrag processing option.
- F. Then process the database with the Process Recalc processing option.
- G. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Data processing option.
- H. Then process the database with the Process Defrag processing option.
- I. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Data processing option.
- J. Then process the database with the Process Recalc processing option.

Answer: D

NEW QUESTION 8

DRAG DROP - (Topic 4)

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]
```

- A. Mastered
- B. Not Mastered

Answer: A

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 9

HOTSPOT - (Topic 4)

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"/>

- A. Mastered
- B. Not Mastered

Answer: A


Explanation:

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

NEW QUESTION 10

DRAG DROP - (Topic 4)

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])
  )
)
  
```

- A. Mastered
- B. Not Mastered

Answer: A

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 10

- (Topic 4)

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 14

- (Topic 4)

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 17

- (Topic 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 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 22

- (Topic 4)

You are managing a SQL Server Analysis Services (SSAS) tabular database.

The database must meet the following requirements:

- ? The processing must load data into partitions or tables.
- ? The processing must not rebuild hierarchies or relationships.
- ? The processing must not recalculate calculated columns.

You need to implement a processing strategy for the database to meet the requirements. Which processing mode should you use?

- A. Process Clear
- B. Process Data
- C. Process Add
- D. Process Full
- E. Process Default

Answer: C

NEW QUESTION 24

- (Topic 4)

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 29

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 70-768 Exam with Our Prep Materials Via below:

<https://www.certleader.com/70-768-dumps.html>