

70-768 Dumps

Developing SQL Data Models (beta)

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



NEW QUESTION 1

DRAG DROP - (Topic 1)

You need to resolve the issues that the users report.

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

Processing options		Answer Area
Process Clear		
Process Update		
Process Index	●	Data availability during cube processing
Process Default	●	Maximum data availability
Process Data	●	Less than maximum data availability
Process Full		Least data availability
		Processing option
		Processing option
		Processing option

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box1: Process Full:

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

Box 2: Process Default

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

Box 3:

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

NEW QUESTION 2

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

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 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.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Partition the CoffeSale 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 3

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

Processing option

Setting

	▼
Default	
Do not process	
Full	

Transactional deployment

	▼
False	
True	

- 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 4

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 5

HOTSPOT - (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 configure SQL Server Profiler to determine why the query is performing poorly. Which three event should you monitor on the SQL Server Profiler trace events configuration page? To answer, select the appropriate options in the answer area.

Answer area

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

- A. Mastered
- B. Not Mastered

Answer: A

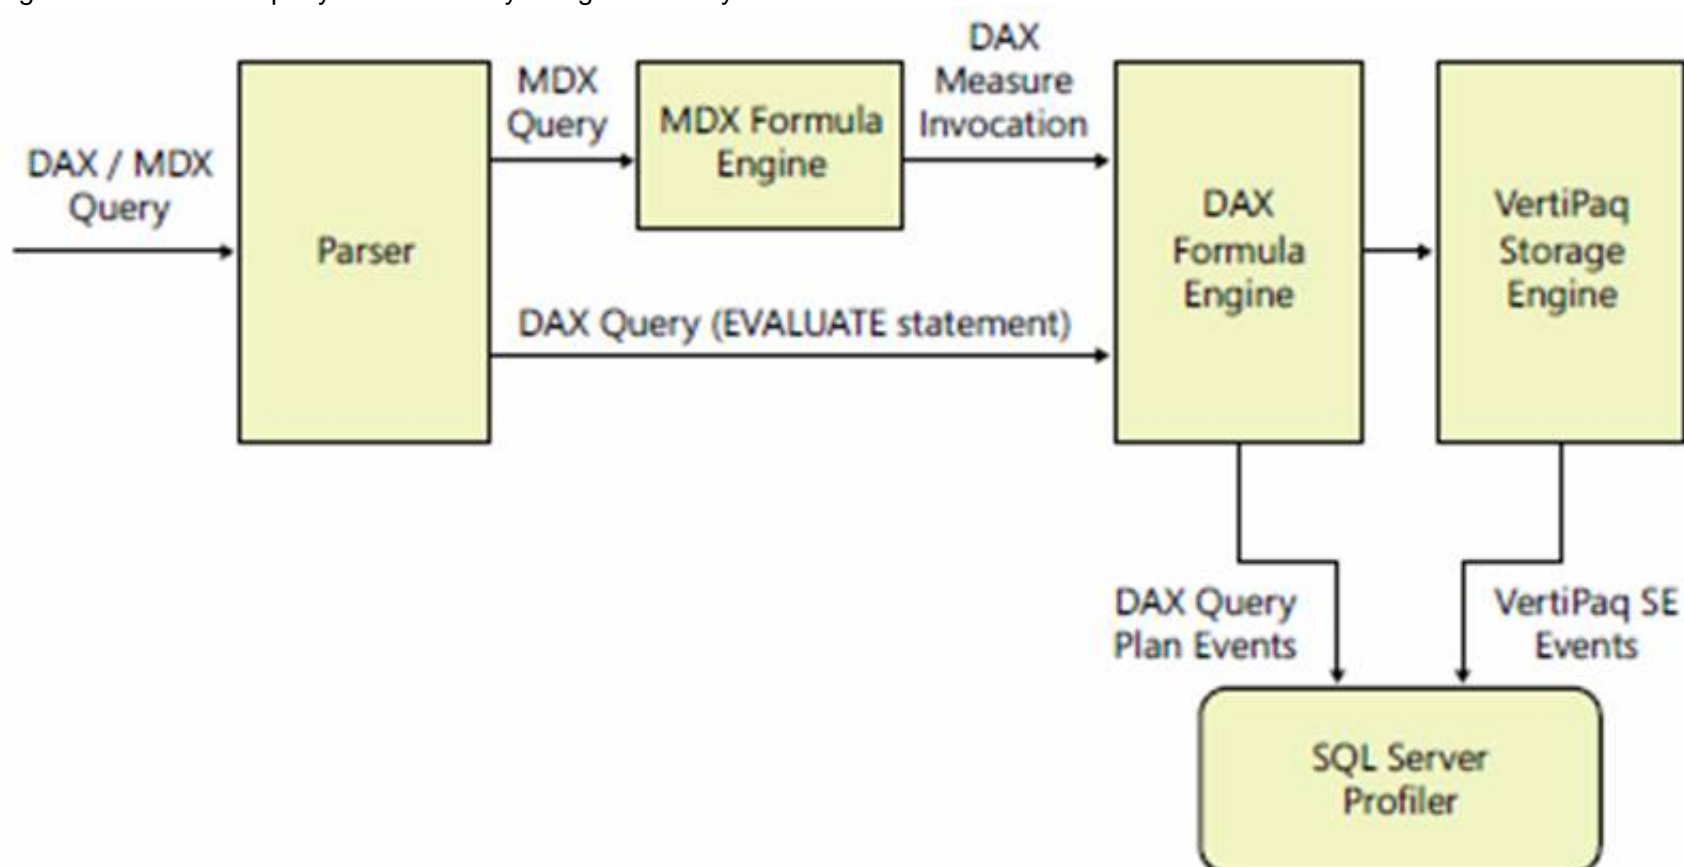
Explanation:

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

DAX Query Plan events are generated by the DAX formula.

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

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



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

NEW QUESTION 6


HOTSPOT - (Topic 4)

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.)



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:



NEW QUESTION 7

- (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 8

HOTSPOT - (Topic 4)

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"/>

- A. Mastered
- B. Not Mastered

Answer: A

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 9

- (Topic 4)

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

- (Topic 4)

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

- (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 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 15

- (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: 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 17

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 18

HOTSPOT - (Topic 4)

You are a database administrator in a company that uses Microsoft SharePoint Server for all intranet sites. You are responsible for the installation of new database server instances.

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

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

What should you do? In the table below, select the appropriate services for each project. NOTE: Make only one selection in each column. Each correct selection is worth one point.

Answer Area

Action

Project1

Project2

Install one tabular instance of SSAS.

☐
☐

Install one multidimensional instance of SSAS.

☐
☐

Install a Power Pivot instance of SSAS.

☐
☐

Install two separate tabular instances of SSAS.

☐
☐

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Project1:

Project2: Multidimensional

Note: Analysis Services can be installed in one of three server modes: Multidimensional and Data Mining (default), Power Pivot for SharePoint, and Tabular.

NEW QUESTION 22

- (Topic 4)

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

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

You need to add the City column to the Customer table. How should you write the calculation?

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

Answer: C

Explanation:

* RELATED Function Returns a related value from another table.

NEW QUESTION 26

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

Does the solution meet the goal?

- A. Yes
B. No

Answer: B

NEW QUESTION 29

- (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 33

- (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: Enable Buffer Cache Extensions. Does the solution meet the goal?

A. Yes

B. No

Answer: B

Explanation:

In this scenario we would need both Buffer Cache Extensions and SSD.

The buffer pool extension provides the seamless integration of a nonvolatile random access memory (that is, solid-state drive) extension to the Database Engine buffer pool to significantly improve I/O throughput.

References:<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/buffer-pool-extension>

NEW QUESTION 37

- (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 38

.....

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>