DUMPSBOSS.

Developing SQL Data Models

Microsoft 70-768

Version Demo

Total Demo Questions: 5

Total Premium Questions: 38

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1 - (HOTSPOT)

HOTSPOT

Case Study #3

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

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

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

To start the case study

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

Background

You are a developer for a Seattle-based company. The company is expanding globally. Many company employees speak fluent Mandarin and read Simplified Chinese.

You have six tabular data models that are deployed to two instances of Microsoft SQL Server Analysis Services (SSAS).

Users report that the query takes a long time to complete.

You are planning the disk space allocations for a new Microsoft SQL Server Analysis Services deployment. You plan to move several relational data file databases to the new SSAS instance. The databases require a total of 10 GB of disk space.

You also plan to deploy Cubes and Aggregations and use Object Processing. Cubes will have small fact tables and few dimension members. No unnecessary aggregations will be created. You plan to process an entire cube in a single transaction.

Data Models

One of the data models is named CustomerSales. This data model contains eight tables. The model includes a table named Sales that defines several measures, including a measure named PriorYearSales. The PriorYearSales measure is referenced by other measures, and is not intended to be analyzed directly by users. You must translate the metadata for all fata the CustomerSales data model to Simplifies Chinese. Team members from the Shanghai office assist with identifying appropriate translations.

A data model named OrderAnalysis is deployed to one of the SSAS instances. Order data is loaded into the OrderAnalysis data as part of an overnight process. You observe that the model is not up-to-date.

The business analysis team uses a variety of client applications to issue MDX queries against OrderAnalysis. Order data must be completely up-to-date.

The OrderAnalysis model has two user-defined hierarchies that are defined in a table named Order. New customers are only added once per day. The overnight process is sufficiently up-to-date for the Customer data to provide optimal performance while achieving the data currency goals whenever possible.

Databases

You deploy a database named DB1 to an SSAS instance as a project by using SQL Server Data Tools. Data analysts report that they cannot access near real time data from the SSAS SalesAnalysis model from DB1. You discover that the project has been deployed with the Direct Query Mode option set to OFF.

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.

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

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

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

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

Hot Area:

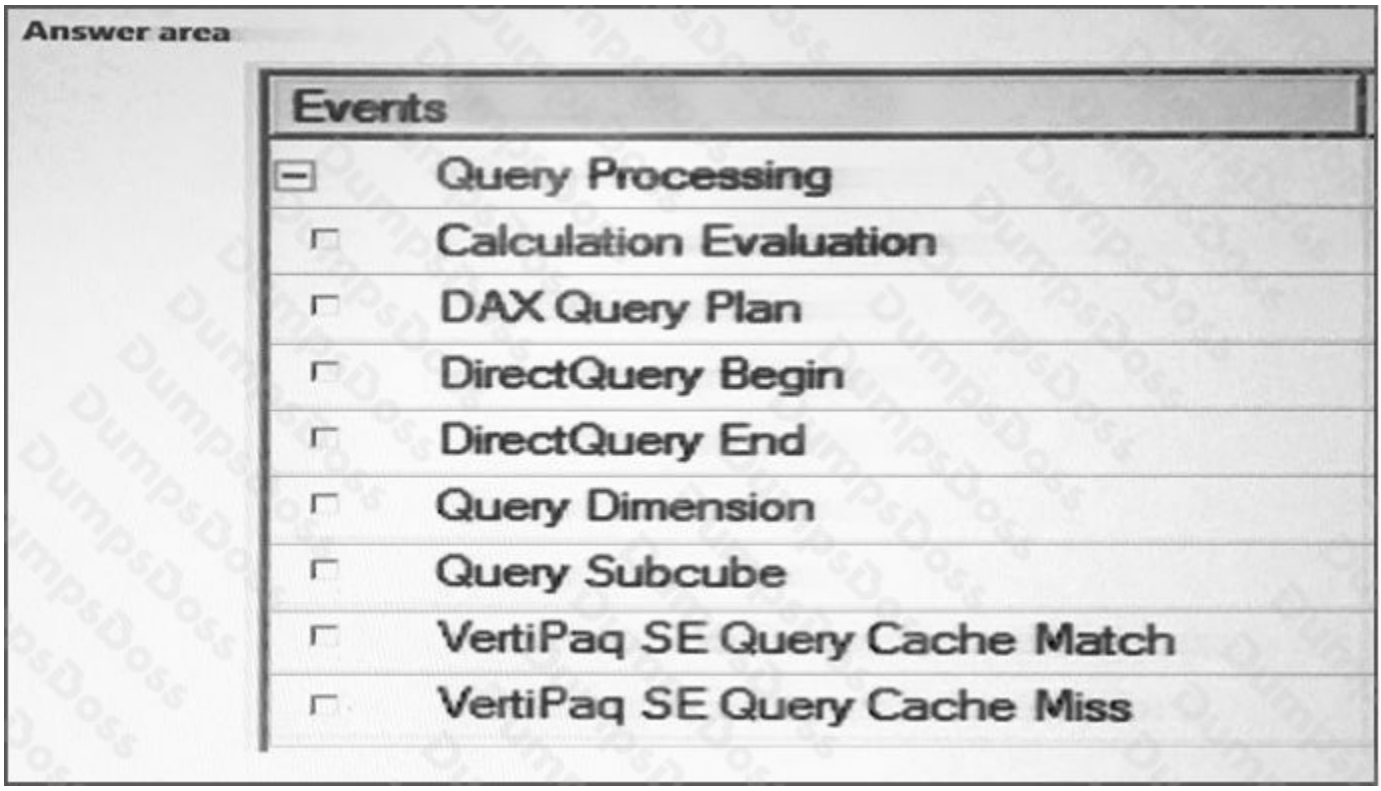
Answer area
Events
Query Processing
Calculation Evaluation
DAX Query Plan
DirectQuery Begin
DirectQuery End
Query Dimension
Query Subcube
VertiPaq SE Query Cache Match
VertiPaq SE Query Cache Miss

ANSWER:

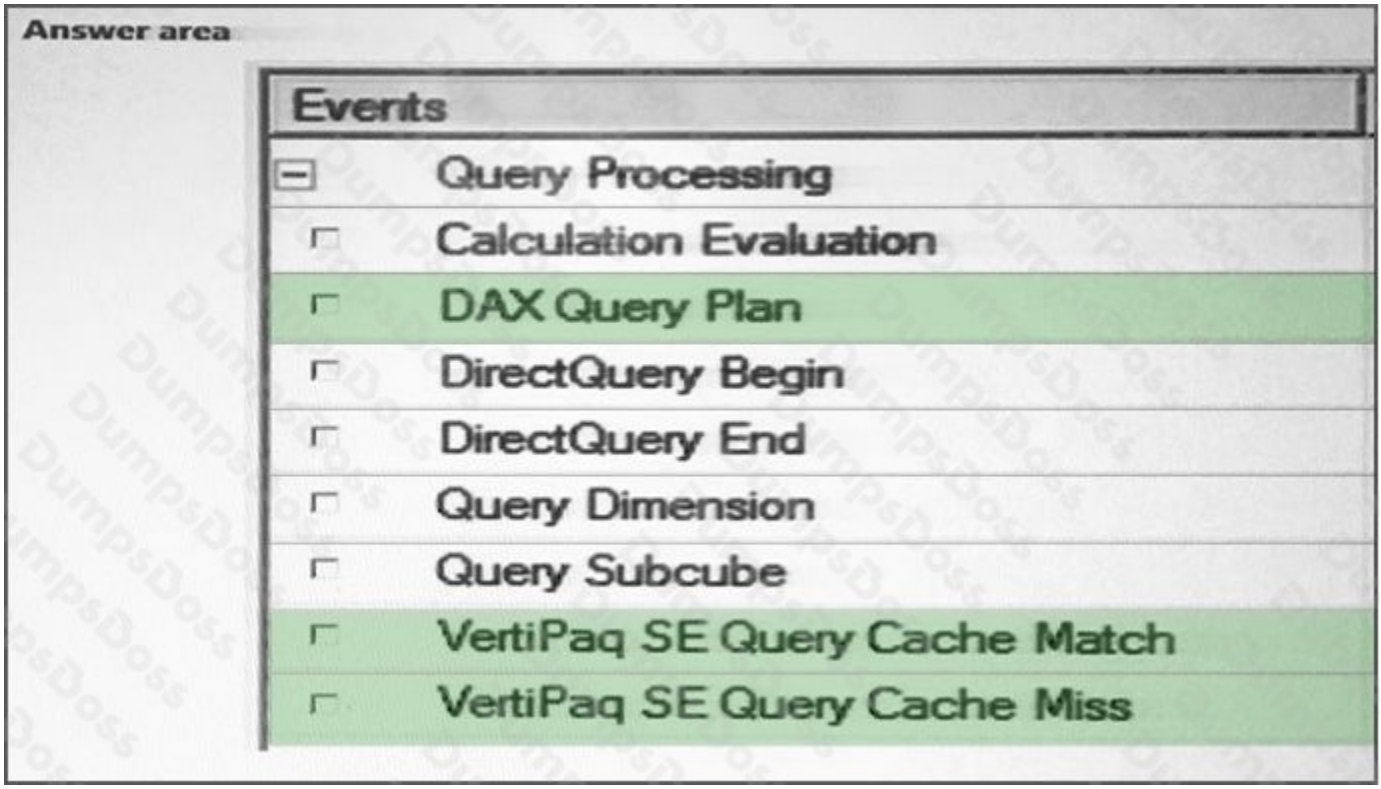
Answer area
Events
Query Processing
Calculation Evaluation
DAX Query Plan
DirectQuery Begin
DirectQuery End
Query Dimension
Query Subcube
VertiPaq SE Query Cache Match
VertiPaq SE Query Cache Miss

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.

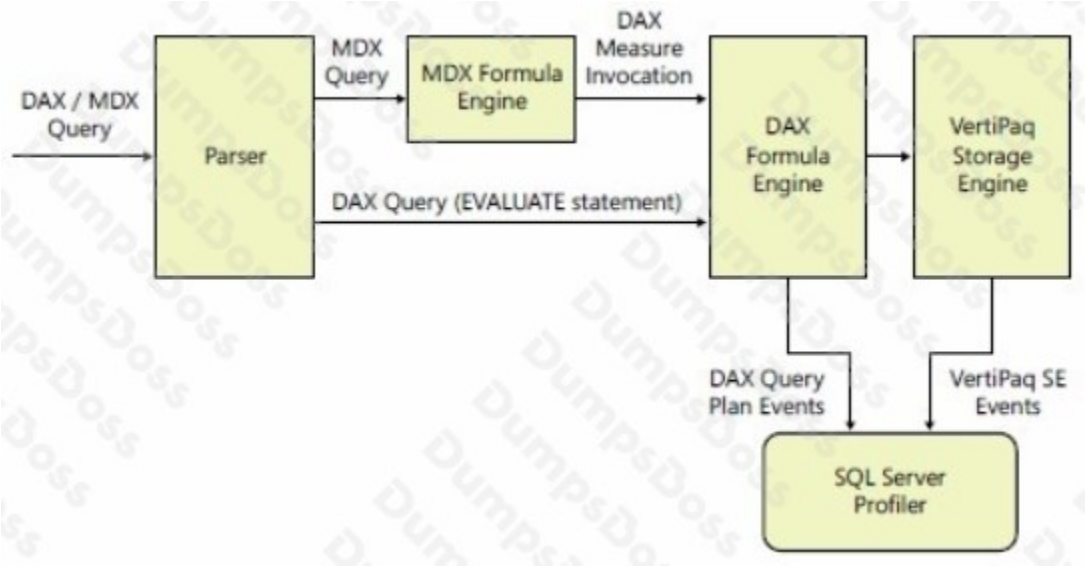

References: Microsoft SQL Server 2012 Analysis Services, The BISM Tabular Model, Microsoft Press (July 2012), page 331

From Scenario: Users report that the query takes a long time to complete.

QUESTION NO: 2

You are responsible for installing new database server instances.

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

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

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

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

A. Install one tabular instance of SSAS and enable the Data Mining Extensions.

B. Install one multidimensional instance of SSAS.

C. Install one tabular instance of SSAS.

D. Install a multidimensional instance and a Power Pivot instance of SSAS on the same server.

E. Install two separate tabular instances of SSAS.

ANSWER: B C

Explanation:

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

References:https://docs.microsoft.com/en-us/sql/analysis-services/comparing-tabular-and-multidimensional-solutions-ssas

QUESTION NO: 3 - (DRAG DROP)

DRAG DROP

Case Study #2

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

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

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

To start the case study

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

Background

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

Sales Analysis

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

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

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

Customer Analysis

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

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

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

Product Sales

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

Requirements

You identify the following requirements:

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

You need to configure the CoffeeSale fact table environment.

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

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

Select and Place:

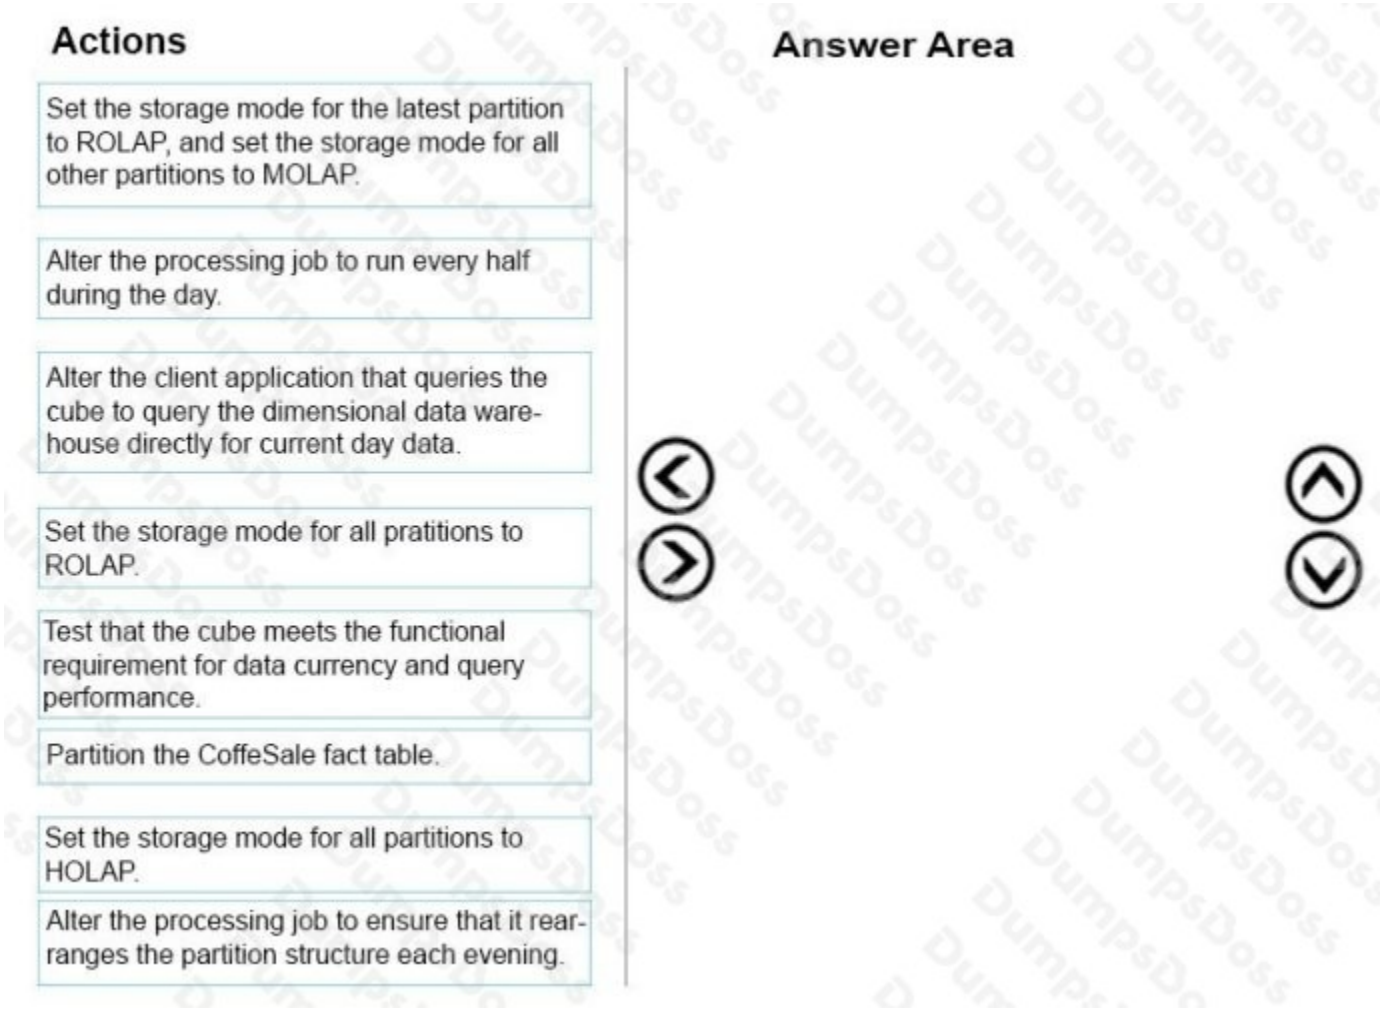
Actions
Set the storage mode for the latest partition to ROLAP, and set the storage mode for all other partitions to MOLAP.
Alter the processing job to run every half during the day.
Alter the client application that queries the cube to query the dimensional data warehouse directly for current day data.
Set the storage mode for all pratitions 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

ANSWER:

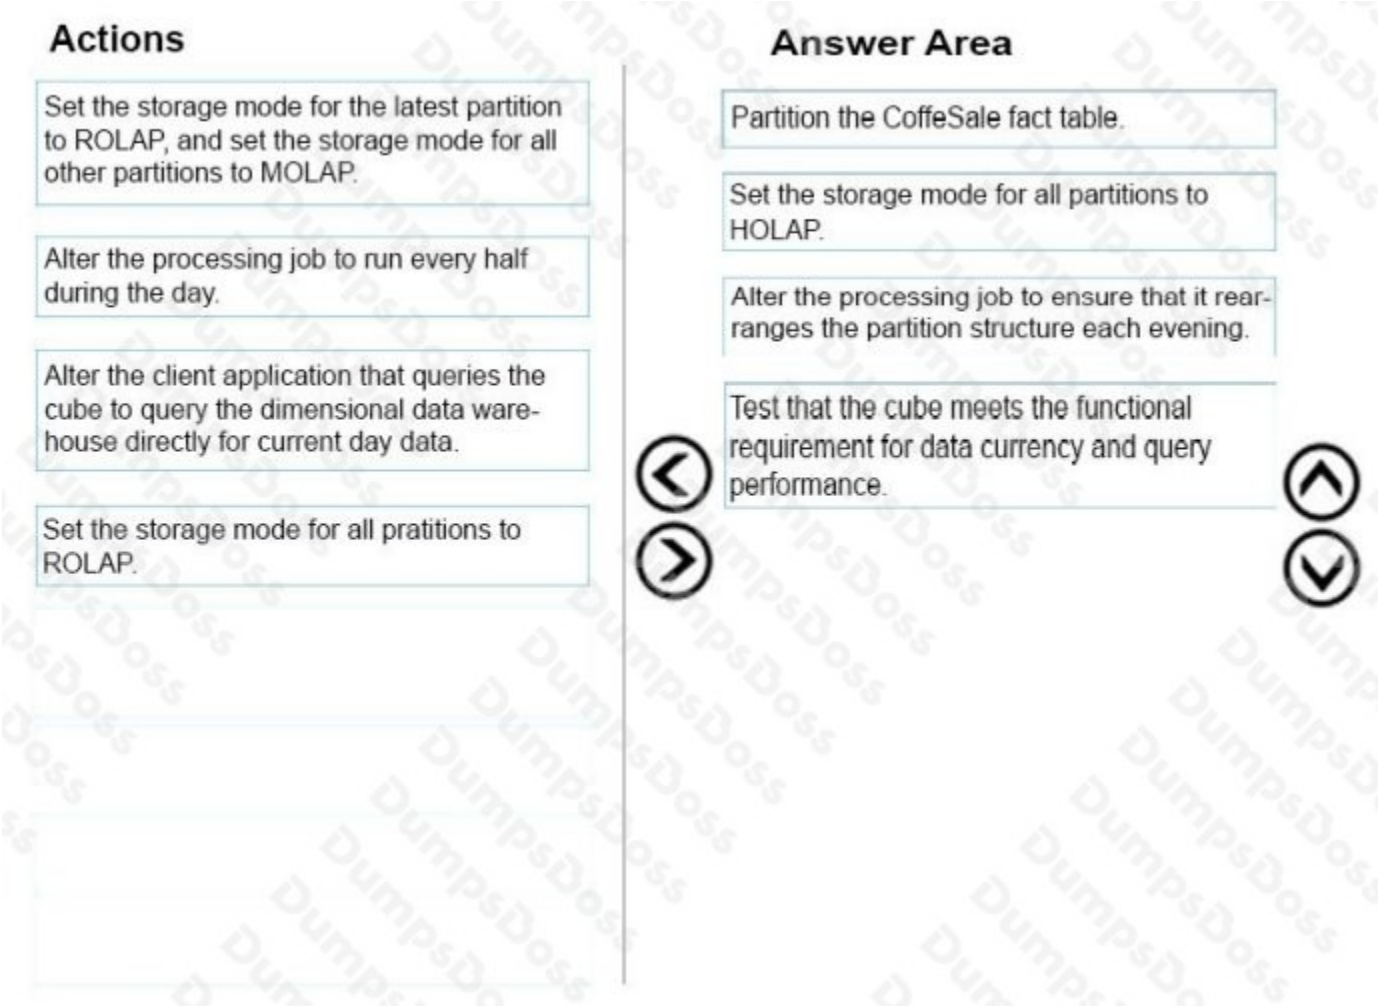

Explanation:

Step 1: Partition the CoffeSale facto table.

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

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

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

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

From scenario:

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

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

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

QUESTION NO: 4

Case Study #3

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

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

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

To start the case study

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

Background

You are a developer for a Seattle-based company. The company is expanding globally. Many company employees speak fluent Mandarin and read Simplified Chinese.

You have six tabular data models that are deployed to two instances of Microsoft SQL Server Analysis Services (SSAS).

Users report that the query takes a long time to complete.

You are planning the disk space allocations for a new Microsoft SQL Server Analysis Services deployment. You plan to move several relational data file databases to the new SSAS instance. The databases require a total of 10 GB of disk space.

You also plan to deploy Cubes and Aggregations and use Object Processing. Cubes will have small fact tables and few dimension members. No unnecessary aggregations will be created. You plan to process an entire cube in a single transaction.

Data Models

One of the data models is named CustomerSales. This data model contains eight tables. The model includes a table named Sales that defines several measures, including a measure named PriorYearSales. The PriorYearSales measure is referenced by other measures, and is not intended to be analyzed directly by users. You must translate the metadata for all fata the CustomerSales data model to Simplifies Chinese. Team members from the Shanghai office assist with identifying appropriate translations.

A data model named OrderAnalysis is deployed to one of the SSAS instances. Order data is loaded into the OrderAnalysis data as part of an overnight process. You observe that the model is not up-to-date.

The business analysis team uses a variety of client applications to issue MDX queries against OrderAnalysis. Order data must be completely up-to-date.

The OrderAnalysis model has two user-defined hierarchies that are defined in a table named Order. New customers are only added once per day. The overnight process is sufficiently up-to-date for the Customer data to provide optimal performance while achieving the data currency goals whenever possible.

Databases

You deploy a database named DB1 to an SSAS instance as a project by using SQL Server Data Tools. Data analysts report that they cannot access near real time data from the SSAS SalesAnalysis model from DB1. You discover that the project has been deployed with the Direct Query Mode option set to OFF.

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.

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 ensure no users see the PriorYearSales measure in the field list for the Sales table.

What should you do?

A. Create a perspective, and ensure that the PriorYearSales measure is not added to the perspective. Ensure that users connect to the model by using the perspective.

B. Set the Display Folder property for PriorYearSales toHidden.

C. Remove the PriorYearSales measure from the default field set of the Sales table.

D. Create a role using Read permissions, and define a DAX expression to filter out the PriorYearSales measure. Add all users to the role.

ANSWER: A

Explanation:

Using perspectives in the data model might help you expose a subset of tables, columns, and measures that are useful for a particular type of analysis. Usually, every user needs only a subset of data you create, and showing him or her the model through perspectives can offer a better user experience.

From scenario; The PriorYearSales measure is referenced by other measures, and is not intended to be analyzed directly by users.

References: Microsoft SQL Server 2012 Analysis Services, The BISM Tabular Model, Microsoft Press (July 2012), page 305

QUESTION NO: 5

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