



Vendor: Microsoft

Exam Code: 70-768

Exam Name: Developing SQL Data Models

Version: Demo

Topic 1, Wide World Importers

Background

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

- SQL Server Database Engine
- SQL Server Analysis Services (multidimensional)

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

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

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

SalesAnalysis

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

OrderAnalysis

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

Finance

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

Requirements

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

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

QUESTION 1

You need to configure the server to optimize the afternoon report generation based on the OrderAnalysis cube. Which property should you configure?

- A. LowMemoryLimit
- B. VertiPaqPagingPolicy

- C. TotalMemoryLimit
- D. VirtualMemoryLimit

Correct Answer: A

Explanation:

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

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

QUESTION 2

DRAG DROP

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

Correct Answer:

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

QUESTION 3

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

- Process Clear
- Process Update
- Process Index
- Process Add
- Process Data
- Process Structure

Answer Area

Job	Processing task
Incremental cube processing	Processing task
Incremental dimension processing	Processing task

Correct Answer:

Processing tasks

- Process Clear
- Process Update
- Process Index
- Process Add
- Process Data
- Process Structure

Answer Area

Job	Processing task
Incremental cube processing	Process Data
Incremental dimension processing	Process Update

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.

QUESTION 1

DRAG DROP

You need to configure the SalesAnalysis cube to correct the sales analysis by customer calculation.

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.

Actions

- Configure a relationship between the Customer dimension and the Sales measure group. Use Month as the granularity.
- Open the dimension editor, and open the Dimension Usage tab.
- Configure a relationship between the Customer dimension and the Sales measure group. Use Day as the granularity.
- Open the dimension editor for the Customer dimension.
- Open the cube editor, and open the Dimension Usage tab.
- Reprocess the Product dimension.
- Reprocess the cube.
- Deploy the project changes.

Answer Area



Correct Answer:

Actions

- Configure a relationship between the Customer dimension and the Sales measure group. Use Month as the granularity.
- Open the dimension editor, and open the Dimension Usage tab.
- Configure a relationship between the Customer dimension and the Sales measure group. Use Day as the granularity.
- Open the dimension editor for the Customer dimension.
- Open the cube editor, and open the Dimension Usage tab.
- Reprocess the Product dimension.
- Reprocess the cube.
- Deploy the project changes.



Answer Area



Open the cube editor, and open the Dimension Usage tab.

Configure a relationship between the Customer dimension and the Sales measure group. Use Day as the granularity.

Reprocess the cube.

Deploy the project changes.

QUESTION 2

HOTSPOT

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: 5px;"> <div style="background-color: #ccc; padding: 2px 5px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> <li style="padding: 2px 5px;">Default <li style="padding: 2px 5px;">Do not process <li style="padding: 2px 5px;">Full </div>
Transactional deployment	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #ccc; padding: 2px 5px; display: flex; justify-content: space-between; align-items: center;"> ▼ </div> <ul style="list-style-type: none"> <li style="padding: 2px 5px;">False <li style="padding: 2px 5px;">True </div>

Correct Answer:

Answer Area

Location	Setting
Processing option	<input type="text" value="Default"/> ▼ Default Do not process Full
Transactional deployment	<input type="text" value="False"/> ▼ False True

QUESTION 3

DRAG DROP

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



Correct Answer:

Actions

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

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

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

Set the storage mode for all partitions to ROLAP.

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

Partition the CoffeSale fact table.

Set the storage mode for all partitions to ROLAP.

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

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



Topic 3, Case Study #3

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 data the CustomerSales data model to Simplified 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.

QUESTION 1

HOTSPOT

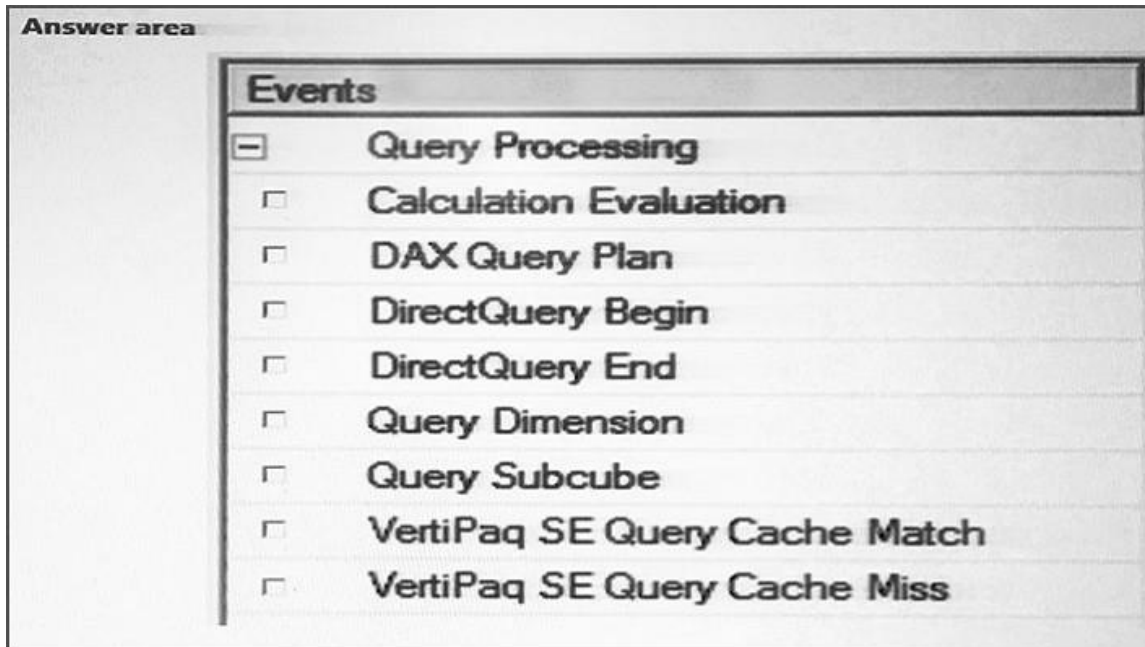
A database named DB2 uses the InMemory query mode. Users frequently run the following query:

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

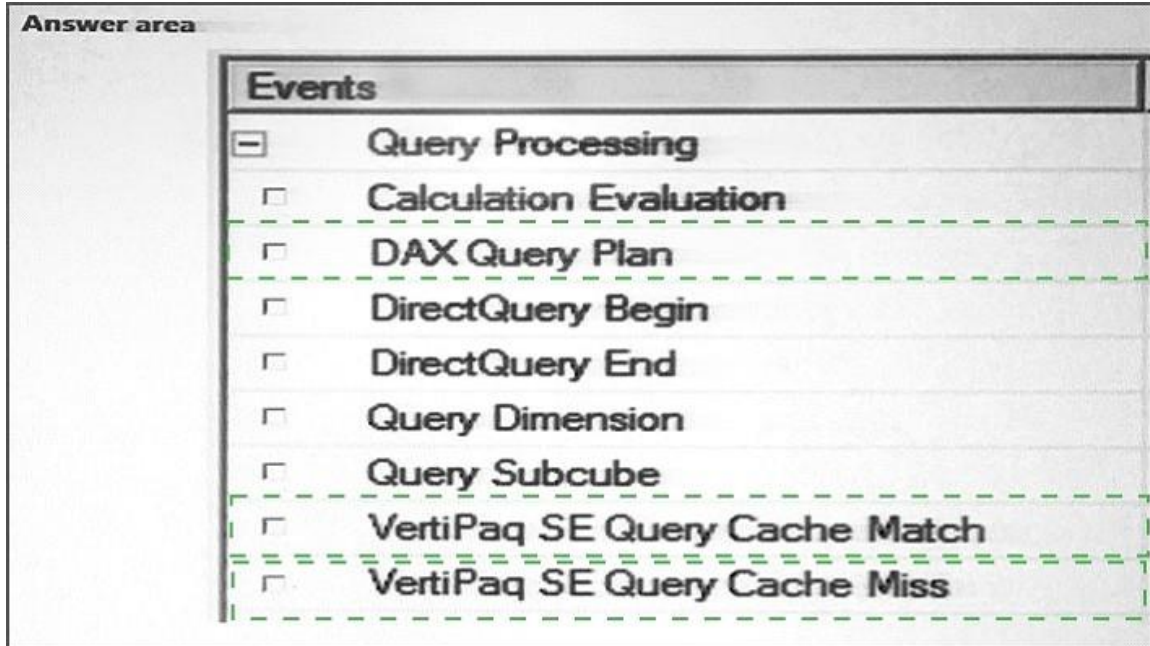
You need to configure SQL Server Profiler to determine why the query is performing poorly.

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.



Correct Answer:



QUESTION 2

DRAG DROP

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



Correct Answer:

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

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.



QUESTION 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 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.

Correct 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