



-The original certification question!

<https://www.it-exams.com>

Exam Number/Code:70-768

Exam Name: Developing SQL Data
Models

Version: Demo

Case Study: 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.

Order Analysis

The OrderAnalysis cube is used for reporting and ad-hoc queries from Microsoft Excel. The data warehouse team adds a new table named FactTransaction to the cube. The FactTransaction 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.

Q1

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

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.

Q2

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
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

Answer:

Answer Area

Data availability during cube processing

Maximum data availability

Less than maximum data availability

Least data availability

Processing option

Process Full

Process Default

Process Update

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.

Q3

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

Answer:

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

Explanation:

Box 1: ProcessData

Processes data only without building aggregations or indexes. If there is data in the partitions, it will be dropped before re-populating the partition with source data.

Box 2: Process Update

Forces a re-read of data and an update of dimension attributes. Flexible aggregations and indexes on related partitions will be dropped.

References:

<https://docs.microsoft.com/en-us/sql/analysis-services/multidimensionalmodels/processing-options-and-settings-analysis-services>

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.

Q4

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	<table border="1"><tr><td></td><td>▼</td></tr><tr><td colspan="2">Default</td></tr><tr><td colspan="2">Do not process</td></tr><tr><td colspan="2">Full</td></tr></table>		▼	Default		Do not process		Full	
	▼								
Default									
Do not process									
Full									
Transactional deployment	<table border="1"><tr><td></td><td>▼</td></tr><tr><td colspan="2">False</td></tr><tr><td colspan="2">True</td></tr></table>		▼	False		True			
	▼								
False									
True									

Answer:

Answer Area

Location	Setting								
Processing option	<table border="1"><tr><td></td><td>▼</td></tr><tr><td colspan="2">Default</td></tr><tr><td colspan="2">Do not process</td></tr><tr><td colspan="2">Full</td></tr></table>		▼	Default		Do not process		Full	
	▼								
Default									
Do not process									
Full									
Transactional deployment	<table border="1"><tr><td></td><td>▼</td></tr><tr><td colspan="2">False</td></tr><tr><td colspan="2">True</td></tr></table>		▼	False		True			
	▼								
False									
True									

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/multidimensionalmodels/deploying-nt-script-files-specifying-processing-options>

Q5

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



Answer:

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.

Explanation:

Step 1: Open the cube editor, and open the Dimension Usage tab. Step 2: Configure a relationship between the Customer dimension and the Sales measure group. Use Day as the granularity.

From scenario: 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.

Step 3: Reprocess the cube.

Step 4: Deploy the project changes.

Q6

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



Answer:

Answer Area

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.

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

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-modelsolap-logical-cube-objects/partitions-partition-storage-modes-and-processing>