



SUCCEED AT THE
FIRST ATTEMPT!

<https://developer certification.com>



Questions and Answers Demo PDF

Microsoft

Exam 70-768

Developing SQL Data Models

Version: Demo

[Total Questions: 10]

Topic break down

Topic	No. of Questions
Topic 1: Wide World Importers	1
Topic 3: Case Study #3	1
Topic 4: Mix Questions	8



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 No : 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

- Process Clear
- Process Update
- Process Index
- Process Default
- Process Data
- Process Full

Answer Area

Data availability during cube processing

- Maximum data availability
- Less than maximum data availability
- Least data availability

Processing option

- Processing option
- Processing option
- Processing option

Answer:

Processing options

- Process Clear
- Process Update
- Process Index
- Process Default
- Process Data
- Process Full

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:**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

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.

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 in 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 No : 2 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

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.

Explanation:**Answer Area**

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

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

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

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

Topic 4, Mix Questions

Question No : 3 - (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?

<https://developer certification.com>

- A. Yes
- B. No

Answer: B

Question No : 4 DRAG DROP - (Topic 4)

You install a SQL Server Analysis Services (SSAS) instance in tabular mode on a server.

While processing a very large tabular model, you receive an out-of-memory error. You identify that the amount of physical memory in the server is insufficient. Additional physical memory cannot be installed in the server.

You need to configure the server to allow paging to disk by using the operating system page file (pagefile.sys).

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.)

- Change the value of the **Memory VertiPagingPolicy** configuration option to **1**.
- Change the value of the **OLAP Process AllowDiskPaging** configuration option to **1**.
- Change the value of the **Memory VertiPagingPolicy** configuration option to **2**.
- Restart the Analysis Services instance.
- In Object Explorer, right-click the Analysis Services instance and then click **Properties**.
- Change the value of the **Memory VertiPagingMemoryLimit** configuration option to **0**.
- Select the **Show Advanced (All) Properties** checkbox.

Answer:

- Change the value of the **Memory VertiPagingPolicy** configuration option to **1**.
 - Change the value of the **OLAP Process AllowDiskPaging** configuration option to **1**.
 - Change the value of the **Memory VertiPagingPolicy** configuration option to **2**.
 - Restart the Analysis Services instance.
 - In Object Explorer, right-click the Analysis Services instance and then click **Properties**.
 - Change the value of the **Memory VertiPagingMemoryLimit** configuration option to **0**.
 - Select the **Show Advanced (All) Properties** checkbox.
- <https://developer certification.com>
- In Object Explorer, right-click the Analysis Services instance and then click **Properties**.
 - Select the **Show Advanced (All) Properties** checkbox.
 - Change the value of the **Memory VertiPagingPolicy** configuration option to **1**.
 - Restart the Analysis Services instance.

Explanation:

Box 1:

In Object Explorer, right-click the Analysis Services instance and then click **Properties**.

Box 2:

Select the **Show Advanced (All) Properties** checkbox.

Box 3:


Change the value of the **Memory VertiPaqPagingPolicy** configuration option to **1**.

Box 4:

Restart the Analysis Services instance.


Note:

* View or set configuration properties in Management Studio

 In SQL Server Management Studio, connect to an Analysis Services instance.

In Object Explorer, right-click the Analysis Services instance, and then click **Properties**.

The General page appears, displaying the more commonly used properties.

 To view additional properties, click the **Show Advanced (All) Properties** checkbox at the bottom of the page.

Modifying server properties is supported only for tabular mode and multidimensional mode servers. If you installed PowerPivot for SharePoint, always use the default values unless you are directed otherwise by a Microsoft product support engineer.

* **VertiPaqPagingPolicy**

Specifies the paging behavior in the event the server runs low on memory. Valid values are as follows:

Zero (0) is the default. No paging is allowed. If memory is insufficient, processing fails with an out-of-memory error.

1 enables paging to disk using the operating system page file (pagefile.sys).

When **VertiPaqPagingPolicy** is set to 1, processing is less likely to fail due to memory constraints because the server will try to page to disk using the method that you specified.

Setting the **VertiPaqPagingPolicy** property does not guarantee that memory errors will never happen. Out of memory errors can still occur under the following conditions:

There is not enough memory for all dictionaries. During processing, Analysis Services locks the dictionaries for each column in memory, and all of these together cannot be more than the value specified for **VertiPaqMemoryLimit**.

There is insufficient virtual address space to accommodate the process.

To resolve persistent out of memory errors, you can either try to redesign the model to reduce the amount of data that needs processing, or you can add more physical memory to the computer.

Applies to tabular server mode only

* Incorrect: VertiPaqMemoryLimit

If paging to disk is allowed, this property specifies the level of memory consumption (as a percentage of total memory) at which paging starts. The default is 60. If memory consumption is less than 60 percent, the server will not page to disk.

This property depends on the VertiPaqPagingPolicyProperty, which must be set to 1 in order for paging to occur.

Applies to tabular server mode only.

Question No : 5 - (Topic 4)

You maintain SQL Server Analysis Services (SSAS) instances.

You need to configure an installation of PowerPivot for Microsoft SharePoint in a SharePoint farm.

Which tool should you use? (Each correct answer presents a complete solution. Choose all that apply.)

- A. SQL Server Configuration Manager
- B. PowerPivot Configuration Tool
- C. SharePoint Products Configuration Wizard
- D. SharePoint Central Administration
- E. PowerShell

Answer: A,B,D

Question No : 6 - (Topic 4)

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.

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 : 7 - (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.

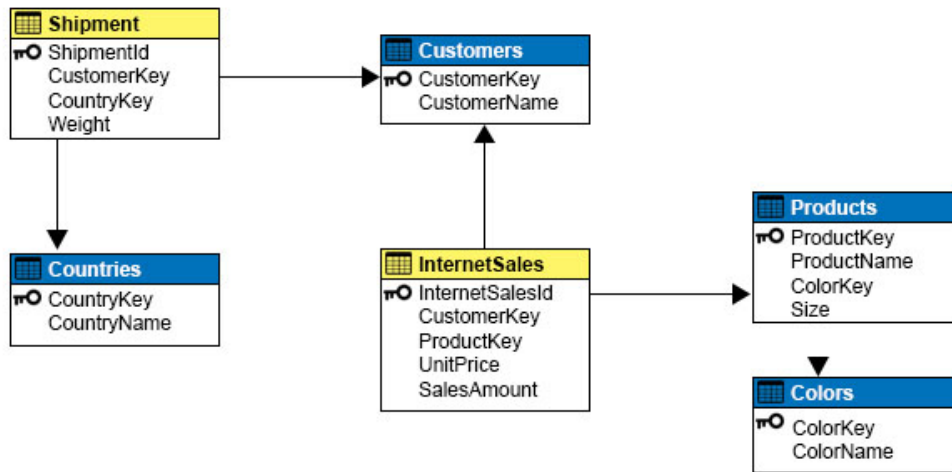
References:<https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/buffer-pool-extension>

Question No : 8 - (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.

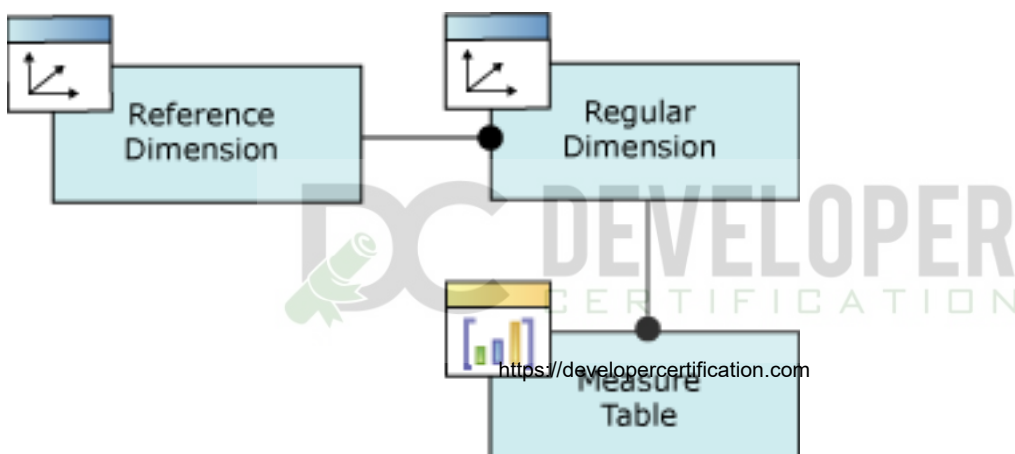
Which relationship type should you use between the InternetSales table and the new dimension?

- A. no relationship
- B. regular
- C. fact
- D. referenced
- E. many-to-many
- F. data mining

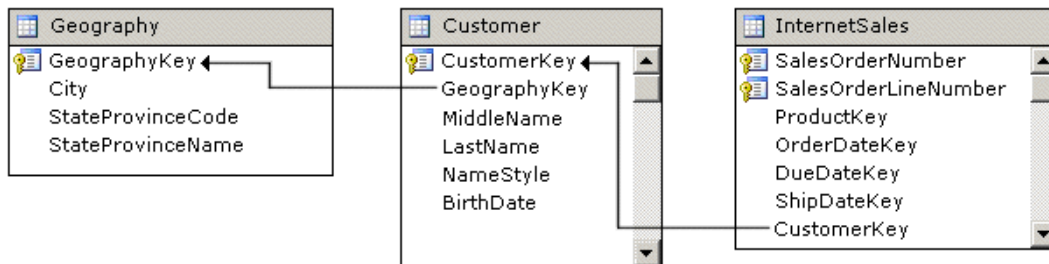
Answer: D

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.

Question No : 9 - (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

Question No : 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 are developing Microsoft SQL Server Analysis Services (SSAS) tabular model.

The model must meet the following requirements:

You need to configure model.

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: A

Explanation:

DAX originally emerged from a Power Pivot add-in for Excel, as a formula language extension for creating calculated columns and measures for data analysis (which are also integral elements of one SSAS Tabular model database, too), but when Microsoft added support for DAX queries in SQL Server 2012, BI experts started “daxing” data from Tabular model databases.

That trend continues, because of simplicity and fast query execution (related to DirectQuery mode in SSAS Tabular).

References:<https://www.sqlshack.com/query-ssas-tabular-model-database-using-dax-functions/>





<https://developer certification.com>