DUMPSARENA

Developing SQL Data Models

Microsoft 70-768

Version Demo

Total Demo Questions: 5

Total Premium Questions: 38

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

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: 2

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: 3

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

QUESTION NO: 4 - (HOTSPOT)

HOTSPOT

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.

Hot Area:

ANSWER:

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

QUESTION NO: 5 - (DRAG DROP)

DRAG DROP

Case Study #1

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 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.

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.

Select and Place:

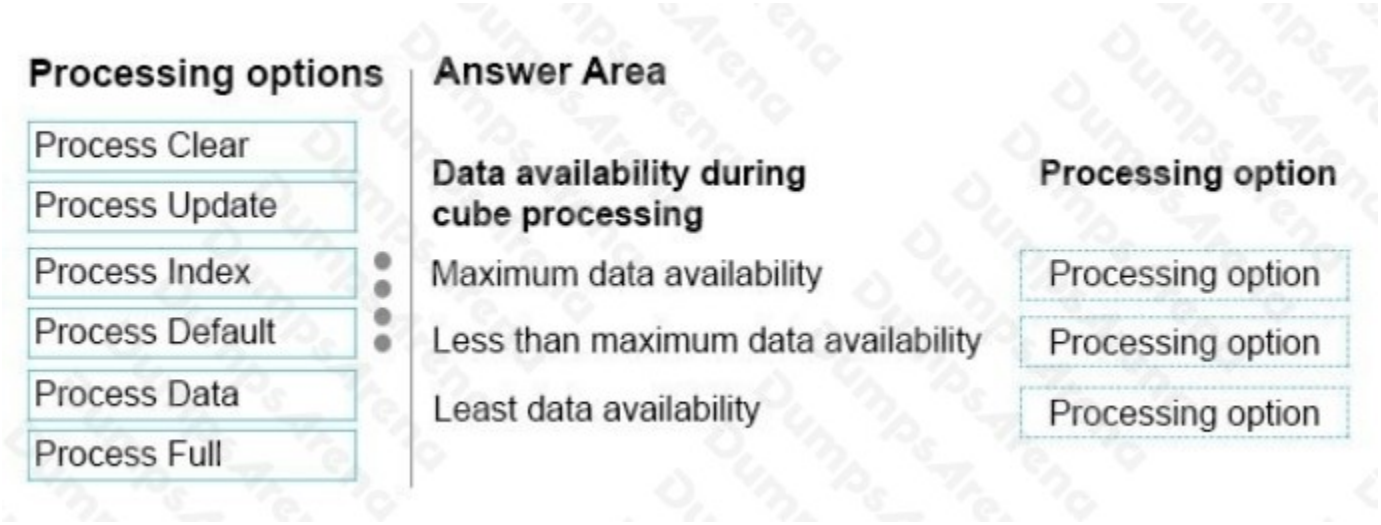

ANSWER:

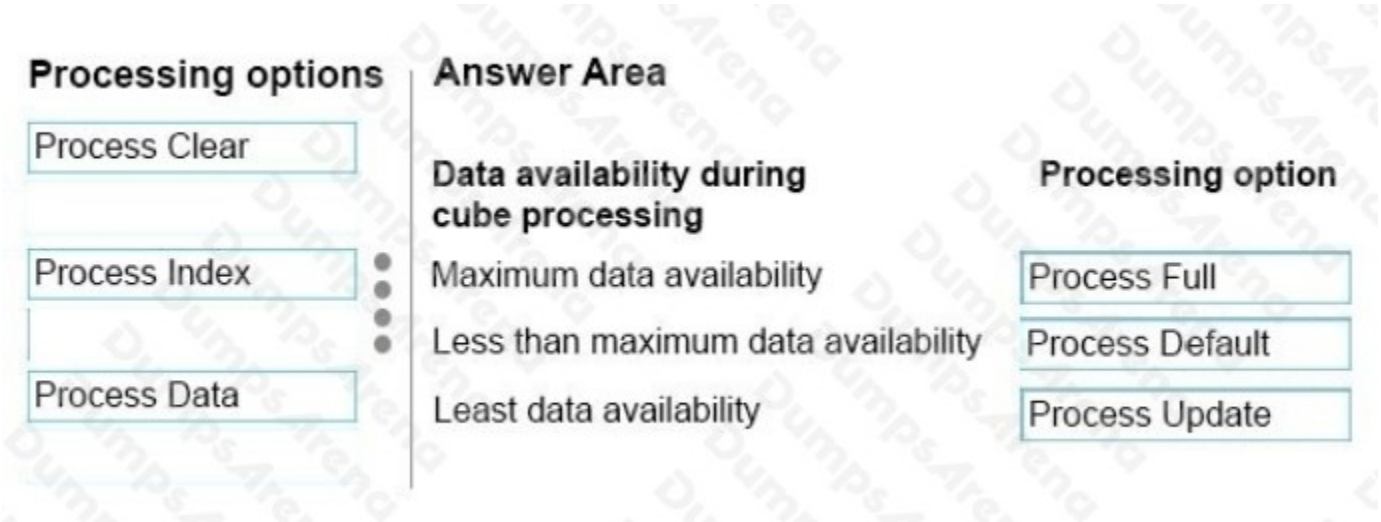

Explanation:

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.

Incorrect Answers:

Not Process Clear: Drops the data in the object specified and any lower-level constituent objects. After the data is dropped, it is not reloaded.

Not Process Data: Processes data only without building aggregations or indexes. If there is data is in the partitions, it will be dropped before re-populating the partition with source data.

Not Process Index: Creates or rebuilds indexes and aggregations for all processed partitions. For unprocessed objects, this option generates an error.

Processing with this option is needed if you turn off Lazy Processing.

References:https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models/processing-options-and-settings-analysis-services