

Microsoft

Exam Questions 70-768

Developing SQL Data Models (beta)



NEW QUESTION 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	Answer Area
Process Clear	<p>Data availability during cube processing</p> <p>Maximum data availability</p> <p>Less than maximum data availability</p> <p>Least data availability</p>
Process Update	
Process Index	
Process Default	
Process Data	
Process Full	
	<p>Processing option</p> <p>Processing option</p> <p>Processing option</p>

- A. Mastered
- B. Not Mastered

Answer: A

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.

NEW QUESTION 2

DRAG DROP - (Topic 1)

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	Answer Area						
Process Clear	<table border="1"> <thead> <tr> <th>Job</th> <th>Processing task</th> </tr> </thead> <tbody> <tr> <td>Incremental cube processing</td> <td>Processing task</td> </tr> <tr> <td>Incremental dimension processing</td> <td>Processing task</td> </tr> </tbody> </table>	Job	Processing task	Incremental cube processing	Processing task	Incremental dimension processing	Processing task
Job		Processing task					
Incremental cube processing		Processing task					
Incremental dimension processing		Processing task					
Process Update							
Process Index							
Process Add							
Process Data							
Process Structure							

- A. Mastered
- B. Not Mastered

Answer: A

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/multidimensional-models/processing-options-and-settings-analysis-services>

NEW QUESTION 3

DRAG DROP - (Topic 2)

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		Answer Area
Set the storage mode for the latest partition to ROLAP, and set the storage mode for all other partitions to MOLAP.	<div style="display: flex; flex-direction: column; align-items: center; gap: 10px;"> ⬅ ➡ </div>	<div style="display: flex; flex-direction: column; align-items: center; gap: 10px;"> ⬆ ⬇ </div>
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 CoffeeSale 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.		

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Partition the CoffeeSale 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-models-olap-logical-cube-objects/partitions-partition-storage-modes-and-processing>

NEW QUESTION 4

HOTSPOT - (Topic 2)

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: 2px;"> <div style="background-color: #cccccc; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;">Default</div> <div style="padding: 2px;">Do not process</div> <div style="padding: 2px;">Full</div> </div>
Transactional deployment	<div style="border: 1px solid gray; padding: 2px;"> <div style="background-color: #cccccc; padding: 2px; text-align: right;">▼</div> <div style="padding: 2px;">False</div> <div style="padding: 2px;">True</div> </div>

- A. Mastered
- B. Not Mastered

Answer: A

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/multidimensional-models/deployment-script-files-specifying-processing-options>

NEW QUESTION 5

DRAG DROP - (Topic 2)

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



- A. Mastered
- B. Not Mastered

Answer: A

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.

NEW QUESTION 6

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



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 7

- (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 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.
- B. Ensure that users connect to the model by using the perspective.
- C. Set the Display Folder property for PriorYearSales to Hidden.
- D. Remove the PriorYearSales measure from the default field set of the Sales table.
- E. Create a role using Read permissions, and define a DAX expression to filter out the PriorYearSales measure.
- F. 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

NEW QUESTION 8

- (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

NEW QUESTION 9

- (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: BC

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>

NEW QUESTION 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 administer a Microsoft SQL Server Analysis Services (SSAS) tabular model for a travel agency that specializes in vacation packages. Vacation bookings and packages are stored in a SQL Server database. You use the model as the basis for customer emails that highlight vacation packages that are currently

underbooked, or projected to be underbooked.

The company plans to incorporate cruise ship vacation packages. Cruise ship vacation packages include new features such as region availability and cruise line specialties that require changes to the tabular model.

You must ensure that the tabular model reflects the new vacation packages. You need to configure the tabular data 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: E

Explanation:

Process Full processes an Analysis Services object and all the objects that it contains. 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.

NEW QUESTION 10

HOTSPOT - (Topic 4)

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.

Answer Area

Option	Dimension Picker	Dimension Salesperson
Create a second data source view.	<input type="radio"/>	<input type="radio"/>
Rename the Employee dimension.	<input type="radio"/>	<input type="radio"/>
Create a new named query for both dimensions.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

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>

NEW QUESTION 13

- (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: B

NEW QUESTION 18

- (Topic 4)

You are developing a SQL Server Analysis Services (SSAS) tabular project. You need to grant the minimum permissions necessary to enable users to query data in a data model. Which role permission should you use?

- A. Explorer
- B. Process
- C. Browser
- D. Administrator
- E. Select
- F. Read

Answer: F

NEW QUESTION 22

- (Topic 4)

You are troubleshooting query performance for a SQL Server Analysis Services (SSAS) cube. A user reports that a Multidimensional Expressions (MDX) query is very slow. You need to identify the MDX query statement in a trace by using SQL Server Profiler. Which event class should you use?

- A. Get Data From Aggregation
- B. Query Subcube
- C. Query Begin
- D. Progress Report Begin
- E. Calculate Non Empty Begin
- F. Execute MDX Script Begin

Answer: C

NEW QUESTION 25

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 VertiPaqPagingPolicy** configuration option to **1**.

Change the value of the **OLAP Process AllowDiskPaging** configuration option to **1**.

Change the value of the **Memory VertiPaqPagingPolicy** 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 VertiPaqMemoryLimit** configuration option to **0**.

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

- A. Mastered
- B. Not Mastered

Answer: A

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 VertiPagingPolicy** 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.

* VertiPagingPolicy

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 VertiPagingPolicy 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 VertiPagingPolicy 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 VertiPagingMemoryLimit.

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: VertiPagingMemoryLimit

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 VertiPagingPolicyProperty, which must be set to 1 in order for paging to occur.

Applies to tabular server mode only.

NEW QUESTION 30

- (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/>

NEW QUESTION 31

- (Topic 4)

You are developing a SQL Server Analysis Services (SSAS) tabular project. The model includes a table named Sales. The Sales table includes a single date column.

The Sales table must meet the following requirements:

? Queries must be able to return all rows.

? Must be able to support four different processing schedules for different date ranges.

? Date ranges must not include any overlapping data.

You need to implement a solution that meets the requirements. What should you do?

- A. Create four partitions for the Sales tabl
- B. Create four role
- C. Use the same row filter queries for each role and partition.
- D. Convert the Sales table into four smaller tables by using row filter querie
- E. Use one perspective for all four tables.
- F. Create four partitions for the Sales tabl
- G. Use row filter queries for each partition.

- H. Convert the Sales table into four smaller tables by using row filter queries
- I. Use one perspective for each of the four tables.

Answer: C

NEW QUESTION 33

- (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.

A company has an e-commerce website. When a customer places an order, information about the transaction is inserted into tables in a Microsoft SQL Server relational database named OLTP1. The company has a SQL Server Analysis Services (SSAS) instance that is configured to use Tabular mode. SSAS uses data from OLTP1 to populate a data model.

Sales analysts build reports based on the SSAS model. Reports must be able to access data as soon as it is available in the relational database.

You need to configure and deploy an Analysis Services project to the Analysis Services instance that allows near real-time data source access.

Solution: In the Deployment Option property for the report, you set the Query Mode to InMemory.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

With InMemory the queries can use the cache only. References:[https://msdn.microsoft.com/en-us/library/hh230898\(v=sql.120\).aspx](https://msdn.microsoft.com/en-us/library/hh230898(v=sql.120).aspx)

NEW QUESTION 38

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

70-768 Practice Exam Features:

- * 70-768 Questions and Answers Updated Frequently
- * 70-768 Practice Questions Verified by Expert Senior Certified Staff
- * 70-768 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 70-768 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 70-768 Practice Test Here](#)