



Microsoft

70-466 Exam

**Microsoft Implementing Data Models and Reports with
Microsoft SQL Server 2012 Exam**

Thank you for Downloading 70-466 exam PDF Demo

You can also try our 70-466 practice exam software

Download Free Demo:

<https://www.dumpsschool.com/70-466-exam-dumps.html>

**Questions & Answers
(Demo Version – Limited Content)**

Version: 17.0

Question: 1

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 table. Create four roles. Use the same row filter queries for each role and partition.
- B. Convert the Sales table into four smaller tables by using row filter queries. Use one perspective for all four tables.
- C. Create four partitions for the Sales table. Use row filter queries for each partition.
- D. Convert the Sales table into four smaller tables by using row filter queries. Use one perspective for each of the four tables.

Answer: C

Question: 2

You are administrating a SQL Server Analysis Services (SSAS) tabular database.

You need to create a new role that allows its members to query data and to refresh data in the model.

Which permission should you use? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. Browse and Manage
- B. Administrator
- C. Read and Process
- D. Explore and Manage

Answer: C

Explanation:

* Giving a database role permission to process an Analysis Services database means that the role has permission to perform all processing options on the database. This includes the processing of all cubes, dimensions, mining structures, and mining models in the database. However, the role does not have permission to read database metadata or access any data in the database itself.

Question: 3

You are developing a SQL Server Analysis Services (SSAS) tabular project.

A column named City must be added to the table named Customer. The column will be used in the definition of a hierarchy. The City column exists in the Geography table that is related to the Customer table.

You need to add the City column to the Customer table.

How should you write the calculation?

- A. City:= LOOKUP(Geography[City],Geography[GeographyKey],[GeographyKey])
- B. City:= LOOKUPVALUE(Geography[City],Geography[GeographyKey],[GeographyKey])
- C. =RELATED(Geography[City])
- D. =RELATED(Geography.City)
- E. =VALUES(Geography[City])
- F. City:=VALUES(Geography[City])

Answer: C

Explanation:

* RELATED Function

Returns a related value from another table.

Question: 4

You are managing a SQL Server Analysis Services (SSAS) tabular database.

The database must meet the following requirements:

The processing must load data into partitions or tables.

The processing must not rebuild hierarchies or relationships.

The processing must not recalculate calculated columns.

You need to implement a processing strategy for the database to meet the requirements.

Which processing mode should you use?

- A. Process Clear
- B. Process Data
- C. Process Add
- D. Process Full
- E. Process Default

Answer: C

Question: 5

You are developing a SQL Server Analysis Services (SSAS) tabular project that will be used by the finance, sales, and marketing teams.

The sales team reports that the model is too complex and difficult to use. The sales team does not need any information other than sales-related resources in the tabular model. The finance and marketing teams need to see all the resources in the tabular model.

You need to implement a solution that meets the needs of the sales team while minimizing development and administrative effort.

What should you do?

- A. Create a separate partition for each team.
- B. Create a separate data source for each team.
- C. Create a perspective for the sales team.
- D. Enable client side security to filter non-sales data.

Answer: C

Question: 6

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

Question: 7

You are developing a SQL Server Analysis Services (SSAS) tabular project.

In the data warehouse, a table named Sales Persons and Territories defines a relationship between a salesperson's name, logon ID, and assigned sales territory.

You need to ensure that each salesperson has access to data from only the sales territory assigned to that salesperson. You need to use the least amount of development effort to achieve this goal.

What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. Create a new role named Sales Persons with Read permission. Add each salesperson's logon as a member to the role.
- B. Add the Sales Persons and Territories table to the model, define the relationships, and then implement dynamic security by using row filters. Grant each salesperson access to the model.
- C. Create a new Active Directory Domain Services (AD DS) security group and add each salesperson as a member. Then create a new role named Sales Persons with Read permission. Add the group as a member to the new role.
- D. Create a separate tabular model for each sales territory and assign each tabular model a corresponding sales territory name. Grant each salesperson access to the corresponding tabular model of the assigned sales territory.

Answer: B

Question: 8

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

Question: 9

DRAG DROP

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:

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.

Explanation:

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.

Reference: Memory Properties; Configure Server Properties in Analysis Services

Question: 10

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: 11

You are developing a tabular Business Intelligence Semantic Model (BISM) database based on a SQL

Server database.

In the data source, the FactInternetSales table is partitioned by month. Data from the current month has been updated and new data has been inserted in the FactInternetSales table, in the DimProduct table, and in the DimCustomer table.

In the model, the FactInternetSales table is also partitioned by month.

You need to ensure that the model has the most recent data while minimizing the processing time. What should you do?

- A. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Clear processing option. Then process the database with the Process Data processing option.
- B. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Clear processing option. Then process the database with the Process Full processing option.
- C. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Defrag processing option. Then process the database with the Process Recalc processing option.
- D. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Data processing option. Then process the database with the Process Defrag processing option.
- E. Process the latest FactInternetSales model table partition, the DimProduct table, and the DimCustomer table with the Process Data processing option. Then process the database with the Process Recalc processing option.

Answer: D

Question: 12

HOTSPOT


You maintain a multidimensional Business Intelligence Semantic Model (BISM) that was developed with default settings.

The model has one cube and the cube has one measure group. The measure group is based on a very large fact table and is partitioned by month. The fact table is incrementally loaded each day with approximately 800,000 new rows.

You need to ensure that all rows are available in the cube while minimizing the processing time.

Which processing option should you use? (To answer, select the appropriate option in the answer area.)

Work Area



Answer:

Explanation:
Process Add

Question: 13

You are modifying a SQL Server Analysis Services (SSAS) multidimensional database. You have identified a dimension that is no longer used by any cubes. You need to delete the dimension. What should you do?

- A. Write a Multidimensional Expressions (MDX) command to drop the dimension from the database.
- B. Write a Data Mining Extensions (DMX) command to drop the dimension from the database.
- C. Script the deletion of the dimension as an XMLA command for execution against the production model.
- D. Write a T-SQL command to drop the dimension from the database.

Answer: C

Question: 14

DRAG DROP

You are developing a SQL Server Analysis Services (SSAS) multidimensional project that is configured to source data from a Microsoft Azure SQL Database database. The cube is processed each night at midnight.

The largest partition in the cube takes 12 hours to process, and users are unable to access the cube until noon. The partition must be available for querying as soon as possible after processing commences.

You need to ensure that the partition is available for querying as soon as possible, without using source data to satisfy the query.

Which three 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.)

Open the cube for editing and then select the **Partitions** tab.

On the Partitions tab, select the partition to edit.

Enable proactive caching and then select the **Bring Online Immediately** option.

Click **Storage Settings** and then click **Options** to open the Storage Options dialog box.

On the Properties window, change the **ProcessingMode** property to **LazyAggregations**.

On the Properties window, change the **ProcessingMode** property to **(default)**.

Answer:

Box 1:

Open the cube for editing and then select the **Partitions** tab.

Box 2:

On the Partitions tab, select the partition to edit.

Box 3:

On the Properties window, change the **ProcessingMode** property to **LazyAggregations**.

Explanation:

Note:

* Processing mode has two possible options.

Regular. This is the default setting. When set to regular, partitions will be available to users after data has been loaded and aggregations are created completely.

Lazy Aggregations. When set to lazy aggregations, partitions will be available to user queries immediately after data has been loaded. Aggregations will be created as a separate background process while users start to query the partition.

* Lazy processing performs the task of building indexes and aggregations for dimensions and measure group partitions at a lower priority to reduce foreground processing time and to allow users to query the cube sooner. For lazy processing to occur, you must switch the ProcessingMode = LazyAggregations of your measure group partitions; by default this value is Regular (lazy processing is turned off). When processing a dimension with flexible aggregations such as parent-child or virtual dimension by using the processing enumeration of ProcessUpdate (such as to take into account of member name or hierarchy changes), lazy processing is initiated to ensure that the aggregations are rebuilt on the associated measure group partitions.

* Configure Lazy Processing for the cube, measure group, or partition. If you configure Lazy Processing, the dropped aggregations are recalculated as a background task. While the flexible aggregations are being recalculated, users can continue to query the cube (without the benefit of the flexible aggregations). While the flexible aggregations are being recalculated, queries that would benefit from the flexible aggregations run slower because Analysis Services resolves these queries by scanning the fact data and then summarizing the data at query time. As the flexible aggregations are recalculated, they become available incrementally on a partition-by-partition basis. For a given cube,

Lazy Processing is not enabled by default. You can configure it for a cube, measure group, or partition by changing the ProcessingMode property from Regular to LazyAggregations. To manage Lazy Processing, there are a series of server properties such as the LazyProcessing \ MaxObjectsInParallel setting, which controls the number of objects that can be lazy processed at a given time. By default it is set to 2. By increasing this number, you increase the number of objects processed in parallel; however, this also impacts query performance and should therefore be handled with care.

* Incorrect: With Bring Online Immediately enabled, during cache refresh all queries are directed to the relational source database to retrieve the latest data for end users. While this provides users with refreshed data, it can also result in reduced query performance given that Analysis Services needs to redirect queries to the relational source database.

Question: 15

DRAG DROP

You are developing a SQL Server Analysis Services (SSAS) multidimensional project.

You identify that a dimension, which has a large number of attributes, is issuing a separate (SELECT DISTINCT) query for each attribute even though they all come from a single database table. The table does not have a large number of rows.

You need to configure the processing of the dimension to issue only a single SQL query to the underlying database, and processing must continue if any errors are encountered.

Which three 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	Answer area
In the Properties window, set the ProcessingGroup property to ByTable and then change the KeyDuplicate property to ReportAndStop .	
In the Properties window, set the ProcessingMode property to ByTable and then change the KeyDuplicate property to IgnoreError .	
Open the dimension designer.	
In the Properties window, set the ProcessingMode property to ByTable and then change the KeyDuplicate property to ReportAndStop .	
Select the dimension.	
In the Properties window, set the ProcessingGroup property to ByTable and then change the KeyDuplicate property to IgnoreError .	
Open the cube designer.	

Answer:

Box 1:

Open the dimension designer.

Box 2:

Select the dimension.

Box 3:

In the Properties window, set the **ProcessingGroup** property to **ByTable** and then change the **KeyDuplicate** property to **IgnoreError**.

Thank You for trying 70-466 PDF Demo

To try our 70-466 practice exam software visit link below

<https://www.dumpsschool.com/70-466-exam-dumps.html>

Start Your 70-466 Preparation

Use Coupon “20OFF” for extra 20% discount on the purchase of Practice Test Software. Test your 70-466 preparation with actual exam questions.