DumpKiller

Latest IT Exam Questions & Answers

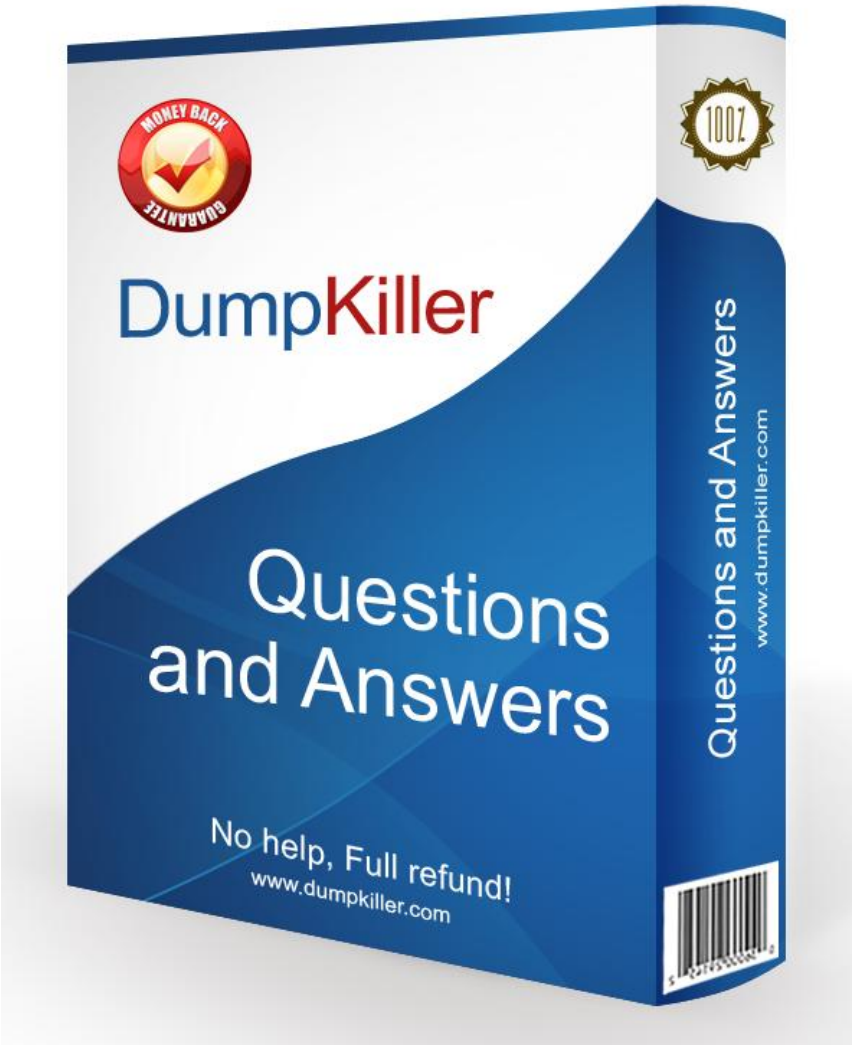

http://www.dumpkiller.com

No help, Full refund!

Exam : **70-768**

Title : Developing SQL Data Models

Vendor : Microsoft

Version : DEMO

NO.1 You are building a Microsoft SQL Server Analysis Services multidimensional model over a SQL Server database.
In a cube named OrderAnalysis, there is a standard cube dimension named Stock Item.
This dimension has the following attributes:
Users report that the attributes Stock Item Key and Photo are distracting and are not providing any value.
They have asked for the attributes to be removed. However, these attributes are needed by other cubes.
You need to hide the specified attributes from the end users of the OrderAnalysis cube. You do not want to change the structure of the dimension.
Which change should you make to the properties for the Stock Item Key and Photo attributes?

A. Set the AttributeHierarchyEnabledproperty to False.

B. Set the AttributeHierarchyDisplayFolder property to Hidden.

C. Set the Usage property to Regular.

D. Set the AttributeVisibility property to Hidden.

E. Set the AttributeHierarchyVisible property to False.

Answer: E
Explanation:
The value of the AttributeHierarchyEnabled property determines whether an attribute hierarchy is created. If this property is set to False, the attribute hierarchy is not created and the attribute cannot be used as a level in a user hierarchy; the attribute hierarchy exists as a member property only. However, a disabled attribute hierarchy can still be used to order the members of another attribute.
If the value of the AttributeHierarchyEnabled property is set to True, the value of the AttributeHierarchyVisible property determines whether the attribute hierarchy is visible independent of its use in a user-defined hierarchy.
https://technet.microsoft.com/en-us/library/ms166717(v=sql.110).aspx

NO.2 You need to configure the server to optimize the afternoon report generation based on the OrderAnalysis cube.
Which property should you configure?

A. VirtualMemoryLimit

B. LowMemoryLimit

C. TotalMemoryLimit

D. VertiPaqPagingPolicy

Answer: B
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.

NO.3 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?

A. Install two separate tabular instances of SSAS.
B. Install one tabular instance of SSAS and enable the Data Mining Extensions.
C. Install a multidimensional instance and a Power Pivot instance of SSAS on the same server.
D. Install one tabular instance of SSAS.
E. Install one multidimensional instance of SSAS.
Answer: D,E

NO.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:

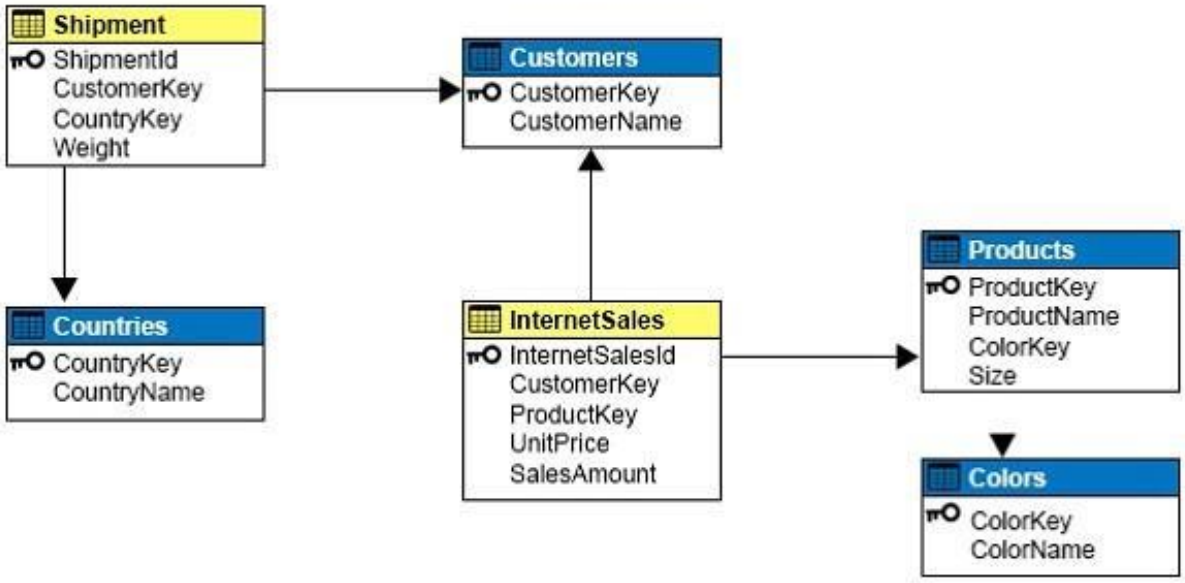

You need to create a new dimension that allows users to list shipments by the country where the product is shipped.

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

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

Many to Many Dimension Relationships.
In most dimensions, each fact joins to one and only one dimension member, and a single dimension member can be associated with multiple facts. In relational database terminology, this is referred to as a one-to-many relationship. However, it is frequently useful to join a single fact to multiple dimension members. For example, a bank customer might have multiple accounts (checking, saving, credit card, and investment accounts), and an account can also have joint or multiple owners. The Customer dimension constructed from such relationships would then have multiple members that relate to a single account transaction.

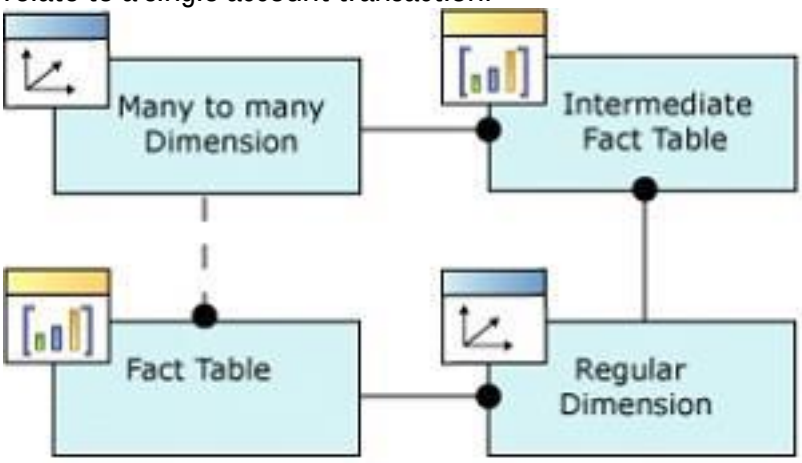

https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models-olap-logical-cube-objects/dimension-relationships

NO.5 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.
https://docs.microsoft.com/en-us/sql/database-engine/configure-windows/buffer-pool-extension

NO.6 Drag and Drop Question

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.

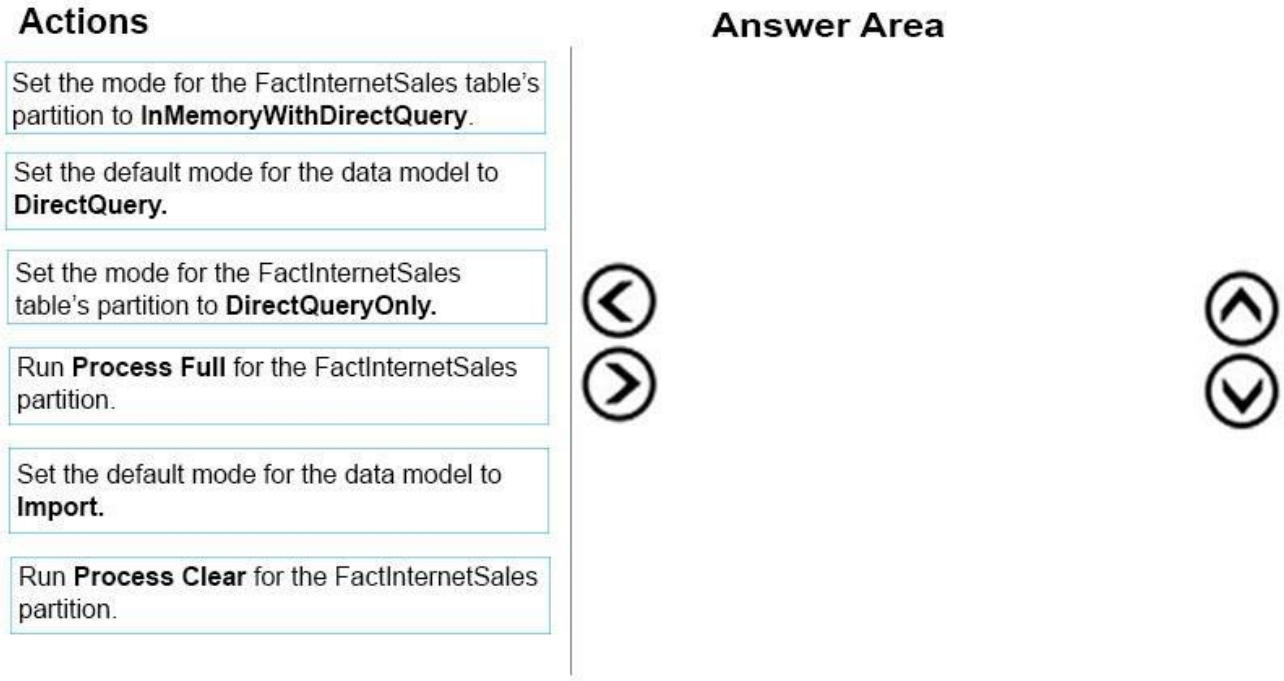

Answer:

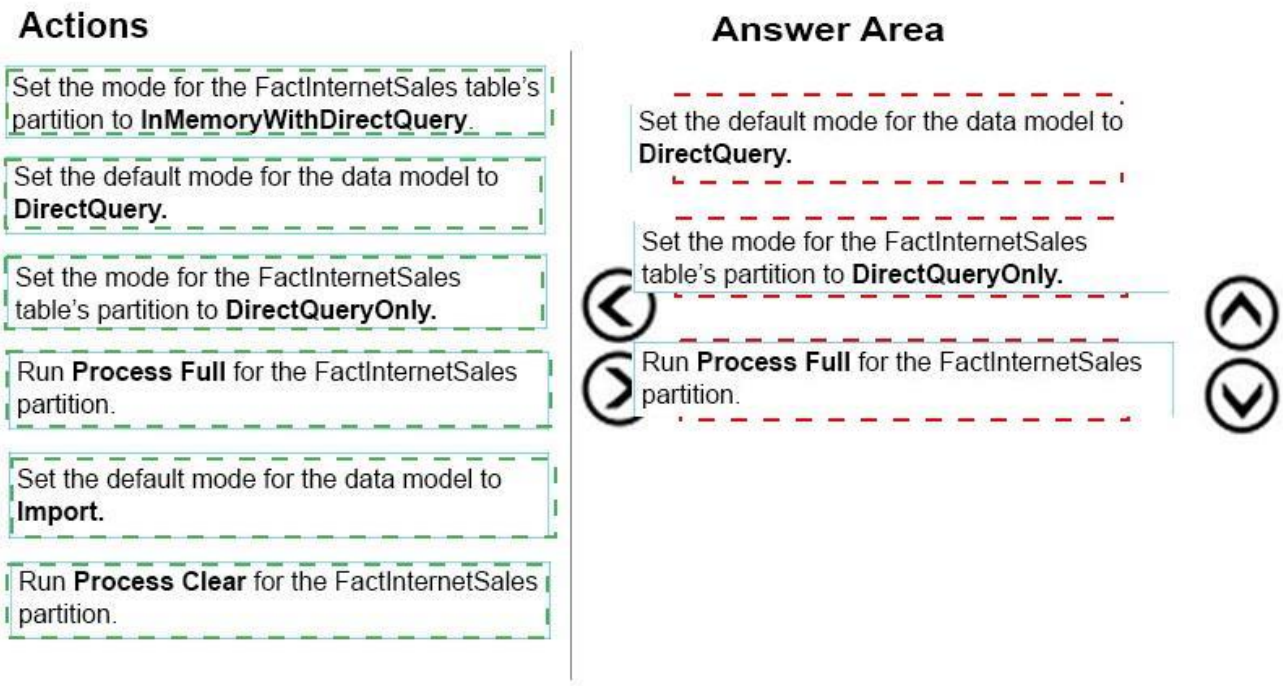

Explanation:
Step 1: Set the default mode for the data model to DirectQuery. You discover that the project has been deployed with the Direct Query Mode option set to OFF.
Step 2: Set the mode for the FactInternetSales table's partition to DirectQueryOnly. Initially, even DirectQuery models are always created in memory. The default query mode for the workspace database is also set toDirectQuery 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