

IT 认证电子书



质 量 更 高 服 务 更 好

半年免费升级服务

<http://www.itrenzheng.com>

Exam : 70-460

Title : Transition Your MCITP:
Business Intelligence
Developer 2008 to MCSE:
Business Intelligence

Version : Demo

1.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:

- 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 **(default)**.

- Open the cube for editing and then select the **Partitions** tab.
- On the Partitions tab, select the partition to edit.
- On the Properties window, change the **ProcessingMode** property to **LazyAggregations**.

Explanation:

* 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 Processing Mode = Lazy Aggregations 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 Process Update (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 Processing Mode property from Regular to Lazy Aggregations. 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.

2.A production SQL Server Analysis Services (SSAS) cube is processed daily. The users query facts by using a hierarchy named Geography from a dimension named Geography. The Dim Geography table in the data source view is used as the source of the Geography dimension.

The table has the following structure.

```
CREATE TABLE [dbo].[DimGeography](
    [DimensionKey] [int] IDENTITY(1,1) NOT NULL,
    [CityKey] [int] NOT NULL,
    [CityName] [varchar](50) NOT NULL,
    [StateProvinceKey] [int] NOT NULL,
    [StateProvinceName] [varchar](50) NOT NULL,
    [CountryKey] [int] NOT NULL,
    [CountryName] [varchar](50) NOT NULL
) ON [PRIMARY]
```

The Geography dimension has three attribute hierarchies:

- City
- State-Province
- Country

The attributes have the following relationships defined: City > State-Province > Country. Each attribute has a key and a name sourced from the related key and name columns in the Dim Geography table.

During processing, you receive the following error message: "Errors in the OLAP storage engine: A duplicate attribute key has been found when processing: Table: 'dbo_DimGeography, Column: State Province Key', Value: '23'. The attribute is 'State-Province'."

You verify that the data is accurate.

You need to ensure that the dimension processes successfully.

What should you do?


- A. Remove the duplicate data from the Dim Geography table.
- B. Delete the Geography hierarchy.
- C. Relate the State-Province and Country attributes directly to the City attribute.
- D. Remove the State-Province attribute.

Answer: C

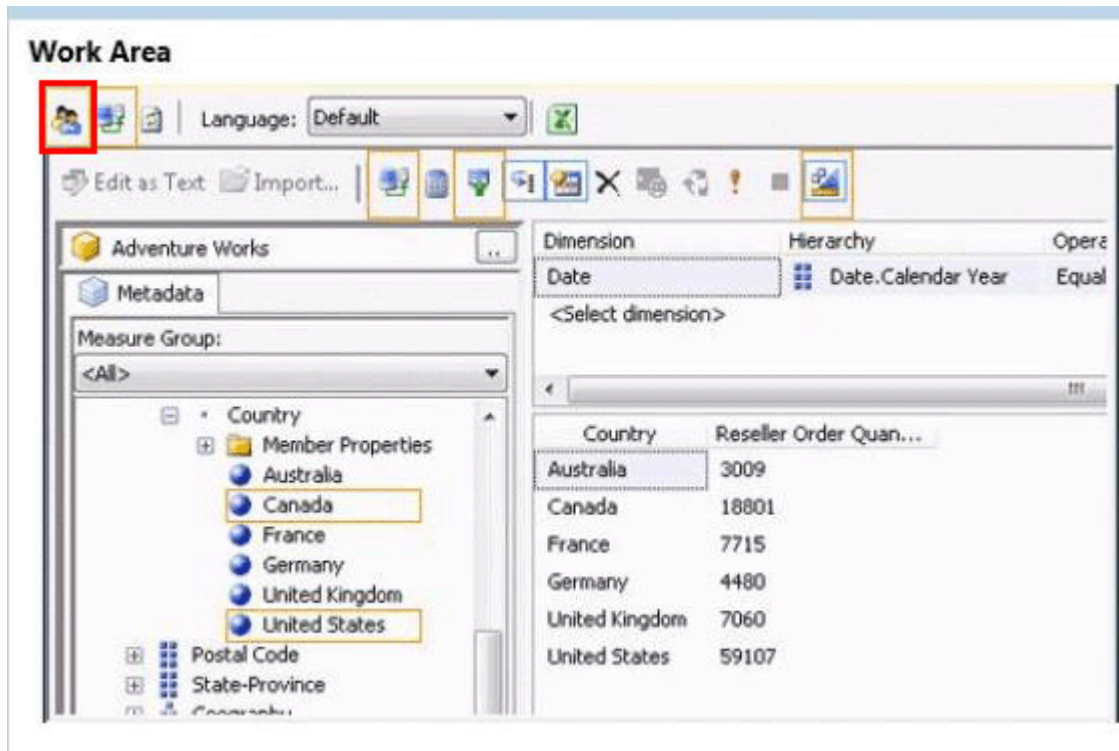
3.HOTSPOT

A SQL Server Analysis Services (SSAS) cube has roles to define dimension data security. A role named USA allows users to browse data pertaining to the United States. A role named Canada allows users to browse data pertaining to Canada. A user can browse sales data pertaining to the United States but cannot browse sales data pertaining to Canada. You validate that the user belongs to the USA and Canada roles. You need to reproduce the issue in SQL Server Management Studio (SSMS).

Which option should you select? To answer, select the appropriate action in the answer area.



Answer:



4.DRAG DROP

You are developing a SQL Server Analysis Services (SSAS) multidimensional project that is configured to source data from a SQL Azure database. You plan to use multiple servers to process different partitions simultaneously. You create and configure a new data source. You need to create a new partition and configure SQL Server Analysis Services (SSAS) to use a remote server to process data contained within the partition.

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, select the **Partitions** tab, and then click **New Partition**.

Click **Enable Proactive caching** and then select the **HOLAP** storage mode.

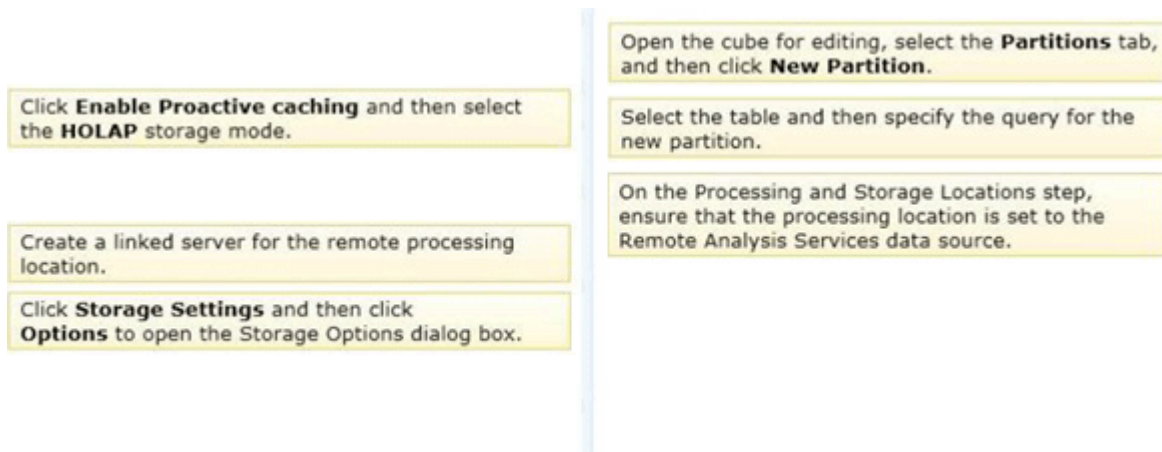
Select the table and then specify the query for the new partition.

Create a linked server for the remote processing location.

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

On the Processing and Storage Locations step, ensure that the processing location is set to the Remote Analysis Services data source.

Answer:



5.You are designing a SQL Server Reporting Services (SSRS) report based on a SQL Server Analysis Services (SSAS) cube. The cube is used to measure sales growth by salesperson. The cube contains a Key Performance Indicator (KPI) to show if a salesperson's sales are off target, slightly off target, or on target. You need to add a report item that visually displays the KPI status value as a red, yellow, or green flag.

Which report item should you add?

- A. a Gauge that uses the Linear type
- B. an Indicator
- C. a Sparkline
- D. a Data Bar
- E. a Gauge that uses the Radial type

Answer: B

6.You are designing a SQL Server Reporting Services (SSRS) report that sources data from a Microsoft Azure SQL Database database. The report must display the value and status of a Key Performance Indicator (KPI).

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

- A. Image
- B. Gauge
- C. Indicator
- D. Data Bar
- E. Sparkline

Answer: B, C

7.You install SQL Server Reporting Services (SSRS). You need to restore a copy of the symmetric key. Which command should you run?

- A. Rskeymgmt -d
- B. Rskeymgmt -a -f %temp%\rs.key -p Password1
- C. Rskeymgmt -i
- D. Rskeymgmt -e -f %temp%\rs.key -p Password1

Answer: B

8.You are creating a new report in SQL Server Report Builder. You add a Windows Azure SQL Database

data source. Then you add a dataset that has four fields named Year, Country, Category, and Sales. You must design a matrix as shown in the following table.

	CY 2003	CY2004
Accessories	293,710	407,050
Bikes	9,359,103	9,162,325
Clothing	138,248	201,525
Australia	3,033,784	2,563,884
Canada	535,784	673,628
France	1,026,325	922,179
Germany	1,058,406	1,076,891
United Kingdom	1,298,249	1,210,286
United States	2,838,512	3,324,031
Total	9,791,060	9,770,900

The category rows (the first three rows as shown in the diagram) must present total sales amount by category. The country rows (the next six rows as shown in the diagram) must present total sales amount by country. The total row must present the total sales for each year. You add a matrix to the report. You add a grouping of the Country field on the rows and a grouping of the Year field on the columns. You need to add the categories on the rows of the matrix.

Which Row Group option should you select when you add the group?

- A. Adjacent Below
- B. Child Group
- C. Parent Group
- D. Adjacent Above

Answer: A

9.You are designing a SQL Server Reporting Services (SSRS) report to display product names and their year-to-date (YTD) sales quantity. YTD sales quantity values are classified in three bands: High Sales, Medium Sales, and Low Sales. You add a table to the report. Then you define two columns based on the fields named Product Name and YTD Sales Quantity. You need to set the color of the product text to red, yellow, or blue, depending on the value of the YTD sales quantity values.

What should you do?

- A. Use an expression for the Text Decoration property of the text box.
- B. Use an expression for the Color property of the text box.
- C. Use an expression for the Font property of the text box.
- D. Use an expression for the Style property of the text box.
- E. Add an indicator to the table.

Answer: B

10.You are developing a SQL Server Reporting Services (SSRS) report that renders in HTML. The report includes a dataset with fields named Description, Price, and Color. The report layout includes a table that displays product details and also includes columns named Description, Price, and Color. You need to modify the report so that users can sort products by the Price column.

What should you do?

- A. Add a group for the Price field.
- B. In the Expression dialog box for the Price text box. enter the =SortByFields!Price, Value expression.

C. Set the SortExpression property to =Fields!Price.Value for the Price text box.

D. Set the SortExpression property to =Fields!Price.Name for the Price text box.

Answer: C

11.DRAG DROP

You are developing a SQL Server Reporting Services (SSRS) report to display a list of employees. The report will be embedded into a Microsoft SharePoint Server Web Part Page of the company intranet site. The report consists of a single table.

The design of the table is shown in the following diagram.

Employee	Email Address
[EmployeeName]	[EmailAddress]

You need to configure the Email Address detail text box to create a new email message. The email message must be addressed to the email address that was clicked by the user.

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.)

Open the Text Box Properties window of the **EmailAddress** detail text box.

In the box, enter the =Fields!EmailAddress.Value expression.

Select the **Action** page.

In the box, enter the =UrlEncode("mailto:" & Fields!EmailAddress.Value)expression.

In the box, enter the ="mailto:" & Fields!EmailAddress.Value expression.

Select the **General** page.

Select the **Go to URL** option.

Select the **Go to Bookmark** option.

Answer:

In the box, enter the =Fields!EmailAddress.Value expression.

In the box, enter the =UrlEncode("mailto:" & Fields!EmailAddress.Value)expression.

Select the **General** page.

Select the **Go to Bookmark** option.

Open the Text Box Properties window of the **EmailAddress** detail text box.

Select the **Action** page.

Select the **Go to URL** option.

In the box, enter the ="mailto:" & Fields!EmailAddress.Value expression.

Explanation:

- * Select Action on the Text Box Properties dialog box to enable hyperlink options for the text box.
- * Go to URL

Choose this option to define a link to a Web page. Type or select the URL of a Web page or an expression that evaluates to the URL of a Web page. Click the Expression (fx) button to change the expression. This expression can include a field that contains a URL.

12.Note: This question is part of a series of questions that use the same set of answer choices. An answer choice may be correct for more than one question in the series. You are creating a SQL Server Analysis Services (SSAS) cube. You need to create a time dimension. It must be linked to a measure group named Sales at the day granularity level. It must also be linked to a measure group named Salary at the month granularity level.

What should you do?

- A. Use role playing dimensions.
- B. Use the Business Intelligence Wizard to define dimension intelligence.
- C. Add a measure that uses the Count aggregate function to an existing measure group.
- D. Add a measure that uses the Distinct Count aggregate function to an existing measure group.
- E. Add a measure group that has one measure that uses the Distinct Count aggregate function.
- F. Add a calculated measure based on an expression that counts members filtered by the Exists and NonEmpty functions.
- G. Add a hidden measure that uses the Sum aggregate function. Add a calculated measure aggregating the measure along the time dimension.
- H. Create several dimensions. Add each dimension to the cube.
- I. Create a dimension. Then add a cube dimension and link it several times to the measure group.
- J. Create a dimension. Create regular relationships between the cube dimension and the measure group. Configure the relationships to use different dimension attributes.
- K. Create a dimension with one attribute hierarchy. Set the IsAggregatable property to False and then set the Default Member property. Use a regular relationship between the dimension and measure group.
- L. Create a dimension with one attribute hierarchy. Set the IsAggregatable property to False and then set the Default Member property. Use a many-to-many relationship to link the dimension to the measure group.
- M. Create a dimension with one attribute hierarchy. Set the IsAggregatable property to False and then set the Default Member property. Use a many-to-many relationship to link the dimension to the measure group.
- N. Create a dimension with one attribute hierarchy. Set the Value Column property, set the IsAggregatable property to False, and then set the Default Member property. Configure the cube dimension so that it does not have a relationship with the measure group. Add a calculated measure that uses the Member Value attribute property.
- O. Create a new named calculation in the data source view to calculate a rolling sum. Add a measure that uses the Max aggregate function based on the named calculation.

Answer: K

13.Note: This question is part of a series of questions that use the same set of answer choices. An answer choice may be correct for more than one question in the series. You are creating a SQL Server Analysis

Services (SSAS) multidimensional database.

Users need a time dimension for:

- Dates
- Delivery dates
- Ship dates

You need to implement the minimum number of required SSAS objects.

What should you do?

- A. Use role playing dimensions.
- B. Use the Business Intelligence Wizard to define dimension intelligence.
- C. Add a measure that uses the Count aggregate function to an existing measure group.
- D. Add a measure that uses the Distinct Count aggregate function to an existing measure group.
- E. Add a measure that uses the LastNonEmpty aggregate function. Use a regular relationship between the time dimension and the measure group.
- F. Add a measure group that has one measure that uses the Distinct Count aggregate function.
- G. Add a calculated measure based on an expression that counts members filtered by the Exists and NonEmpty functions.
- H. Add a hidden measure that uses the Sum aggregate function. Add a calculated measure aggregating the measure along the time dimension.
- I. Create several dimensions. Add each dimension to the cube.
- J. Create a dimension. Then add a cube dimension and link it several times to the measure group.
- K. Create a dimension. Create regular relationships between the cube dimension and the measure group. Configure the relationships to use different dimension attributes.
- L. Create a dimension with one attribute hierarchy. Set the IsAggregatable property to False and then set the Default Member property. Use a regular relationship between the dimension and measure group.
- M. Create a dimension with one attribute hierarchy. Set the IsAggregatable property to False and then set the Default Member property. Use a many-to-many relationship to link the dimension to the measure group.
- N. Create a dimension with one attribute hierarchy. Set the Value Column property, set the IsAggregatable property to False, and then set the Default Member property. Configure the cube dimension so that it does not have a relationship with the measure group. Add a calculated measure that uses the Member Value attribute property.
- O. Create a new named calculation in the data source view to calculate a rolling sum. Add a measure that uses the Max aggregate function based on the named calculation.

Answer: A

14.Note: This question is part of a series of questions that use the same set of answer choices. An answer choice may be correct for more than one question in the series. You are designing a SQL Server Analysis Services (SSAS) cube. You need to create a measure to count unique customers.

What should you do?

- A. Use role playing dimensions.
- B. Use the Business Intelligence Wizard to define dimension intelligence.
- C. Add a measure that uses the Count aggregate function to an existing measure group.
- D. Add a measure that uses the Distinct Count aggregate function to an existing measure group.
- E. Add a measure group that has one measure that uses the Distinct Count aggregate function.

- F. Add a calculated measure based on an expression that counts members filtered by the Exists and NonEmpty functions.
- G. Add a hidden measure that uses the Sum aggregate function. Add a calculated measure aggregating the measure along the time dimension.
- H. Create several dimensions. Add each dimension to the cube.
- I. Create a dimension. Then add a cube dimension and link it several times to the measure group.
- J. Create a dimension. Create regular relationships between the cube dimension and the measure group. Configure the relationships to use different dimension attributes.
- K. Create a dimension with one attribute hierarchy. Set the IsAggregatable property to False and then set the Default Member property. Use a regular relationship between the dimension and measure group.
- L. Create a dimension with one attribute hierarchy. Set the IsAggregatable property to False and then set the Default Member property. Use a many-to-many relationship to link the dimension to the measure group.
- M. Create a dimension with one attribute hierarchy. Set the IsAggregatable property to False and then set the Default Member property. Use a many-to-many relationship to link the dimension to the measure group.
- N. Create a dimension with one attribute hierarchy. Set the Value Column property, set the IsAggregatable property to False, and then set the Default Member property. Configure the cube dimension so that it does not have a relationship with the measure group. Add a calculated measure that uses the Member Value attribute property.
- O. Create a new named calculation in the data source view to calculate a rolling sum. Add a measure that uses the Max aggregate function based on the named calculation.

Answer: E

15. You are developing a SQL Server Analysis Services (SSAS) tabular project. A model contains tables and columns that must not be visible to the user. The columns and tables cannot be removed because they are used in calculations. The calculations are used to calculate the budget and forecast for the current quarter. You need to hide the tables and columns.

What should you do?

- A. Before adding the forecast calculations to the model, right-click the applicable tables and columns and select the Hide option.
- B. After adding the budget calculations to the model, in the Properties window for the applicable tables and columns, set the Visible property to True.
- C. Before adding the forecast calculations to the model, right-click the applicable tables and columns and select the Hide from Client Tools option.
- D. After adding the budget calculations to the model, in the Properties window for the applicable tables and columns, set the Enabled property to False

Answer: C