

EXAMGOOD

QUESTION & ANSWER

Exam Good provides update free of charge in one year!

Accurate study guides
High passing rate!

<http://www.examgood.com>

Exam : **70-467**

Title : Designing Business
Intelligence Solutions with
Microsoft SQL Server 2012

Version : Demo

1.DRAG DROP

You are designing a SQL Server Reporting Services (SSRS) solution. A report project must access multiple SQL Azure databases. Each database is on a different host. The databases have identical schema and security configurations.

You have the following requirements:

- The report must support subscriptions.
- Users must be able to select the host when running the report.

What should you do? To answer, drag the appropriate phrase or phrases from the list to the correct location or locations in the answer area. (Answer choices may be used once, more than once, or not all.)

SQL Azure data.	Create a <input style="width: 150px; height: 15px;" type="text"/>
SQL Azure hosts.	Create a report parameter that displays available values of <input style="width: 150px; height: 15px;" type="text"/>
a shared dataset.	Create <input style="width: 150px; height: 15px;" type="text"/>
stored credentials.	Configure the data source to use <input style="width: 150px; height: 15px;" type="text"/>
integrated security.	
data source in the report.	
an expression-based connection string.	
shared data source in the report	

Answer:

SQL Azure data.	Create a <input style="width: 150px; height: 15px; background-color: #fff9c4;" type="text" value="data source in the report."/>
a shared dataset.	Create a report parameter that displays available values of <input style="width: 150px; height: 15px; background-color: #fff9c4;" type="text" value="SQL Azure hosts."/>
integrated security.	Create <input style="width: 150px; height: 15px; background-color: #fff9c4;" type="text" value="an expression-based connection string."/>
shared data source in the report	Configure the data source to use <input style="width: 150px; height: 15px; background-color: #fff9c4;" type="text" value="stored credentials."/>

Explanation:

* To include data in a report, you must first create *data connections*, also known as *data sources*, and then create datasets.

* A data connection includes the data source type, connection information, and the type of credentials to use. There are two types of data sources: embedded and shared. An embedded data source is defined in the report and used only by that report (fits this scenario). A shared data source is defined independently from a report and can be used by multiple reports.

* Built-in data extensions include the following data connection types:

- Microsoft SQL Server
- Microsoft SQL Server Analysis Services
- Microsoft SharePoint List
- Windows Azure SQL Database

Etc.

* Expression-based connection strings are evaluated at run time. For example, you can specify the data

source as a parameter, include the parameter reference in the connection string, and allow the user to choose a data source for the report.

* Credentials You provide the credentials that are needed to access the data. The data source owner must have granted you the appropriate permissions to access both the data source and the specific data on the data source.

Reference: Data Connections, Data Sources, and Connection Strings (SSRS)


2.DRAG DROP

You are designing a SQL Server Reporting Services (SSRS) solution. An existing report aggregates data from a SQL Server database in a chart. You need to use the chart in a new report and ensure that other users can use the chart in their reports.

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.)

In Power View, open the report that contains the chart.	
In Report Designer, insert the report part into a new report.	
In Report Designer, open the report that contains the chart.	
In Report Builder, insert the report part into a new report.	
In Power View, insert the report part into a new report.	
Select the chart for publication as a report part and publish the report.	

Answer:

In Power View, open the report that contains the chart.		In Report Designer, open the report that contains the chart.
In Report Designer, insert the report part into a new report.		Select the chart for publication as a report part and publish the report.
		In Report Builder, insert the report part into a new report.
In Power View, insert the report part into a new report.		

Explanation:

* In Report Designer, after you create tables, charts, and other report items in a project, you can publish them as report parts to a report server or SharePoint site integrated with a report server so that you and others can reuse them in other reports.

* . By using Report Builder, you can customize and update reports that were created in SQL Server Data Tools (SSDT) Report Designer.

* In Report Builder, IT pros and power users can create powerful operational reports, and reusable report parts and shared datasets.

Incorrect:

* (incorrect) Power View, a feature of SQL Server 2012 Reporting Services Add-in for Microsoft SharePoint Server 2010 Enterprise Edition, is an interactive data exploration, visualization, and presentation experience. It provides intuitive ad-hoc reporting for business users such as data analysts, business decision makers, and information workers. They can easily create and interact with views of data from data models based on PowerPivot workbooks published in a PowerPivot Gallery, or tabular models deployed to SQL Server 2012 Analysis Services (SSAS) instances. Power View is a browser-based Silverlight application launched from SharePoint Server 2010 that enables users to present and share insights with others in their organization through interactive presentations.

Reference: Getting Started with Report Builder

Reference: Report Parts in Report Designer (SSRS)

3.You are designing a subscription strategy for a SQL Server Reporting Services (SSRS) report. You have an application that populates a table with user-specific subscription schedules and report formats. You need to ensure that users can receive reports by email according to their preferences.

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

- A. Create a standard subscription for each record in the table.
- B. Create a data-driven subscription for each record in the schedule table.
- C. Create one data-driven subscription. Schedule the subscription to frequently retrieve user preferences.
- D. Create a standard subscription for each subscription schedule.

Answer: C

4.You are modifying a SQL Server Reporting Services (SSRS) report for a SQL Server Analysis Services (SSAS) cube. The report defines a report parameter of data type Date/Time with which users can filter the report by a single date. The parameter value cannot be directly used to filter the Multidimensional Expressions (MDX) query for the dataset. You need to ensure that the report displays data filtered by the user-entered value. You must achieve this goal by using the least amount of development effort.

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

- A. Edit the dataset query parameter. Change the Value property of the report parameter to an expression that uses the same format as the date dimension member key value.
- B. Edit the dataset query parameter. Change the Name property of the dataset query parameter so that it points to a name value for each date dimension member.
- C. Edit the dataset query parameter. Create a subcube subquery that uses the StrToSet MDX function and accepts the report parameter value.
- D. Change the dataset query to Transact-SQL (T-SQL). Use the OPENROWSET function to query the cube. Output the cube results to the T-SQL query and use a Convert function to change the report parameter value into the same format as the date dimension member.

Answer: A

5.You administer a SQL Server Reporting Services (SSRS) instance in native mode.

You need to assign a predefined role that meets the following requirements:

- Members of the role must be able to update shared data sources.
- Members of the role must not be able to consume reports or manage subscriptions.
- The role must provide only the minimum permissions required.

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

- A. The Content Manager role
- B. The Read and Process role
- C. The Publisher role
- D. The Browser role

Answer: C

Explanation:

Name:
 Publisher

Description:
 May publish reports and linked reports to the Report Server.

Select the tasks that members of this role can perform:

<input type="checkbox"/>	Task ^A	Description
<input type="checkbox"/>	Consume reports	Reads report definitions
<input checked="" type="checkbox"/>	Create linked reports	Create linked reports and publish them to a report server folder.
<input type="checkbox"/>	Manage all subscriptions	View, modify, and delete any subscription regardless of who owns the subscription.
<input checked="" type="checkbox"/>	Manage data sources	Create and delete shared data source items; and modify data source properties.
<input checked="" type="checkbox"/>	Manage folders	Create, view and delete folders; and view and modify folder properties.
<input type="checkbox"/>	Manage individual subscriptions	Each user can create, view, modify and delete subscriptions that he or she owns.
<input checked="" type="checkbox"/>	Manage models	Create, view, and delete models; and view and modify model properties.
<input type="checkbox"/>	Manage report history	Create, view, and delete report history snapshots; and modify report history properties.
<input checked="" type="checkbox"/>	Manage reports	Create, and delete reports; and modify report properties.
<input checked="" type="checkbox"/>	Manage resources	Create, modify and delete resources, and modify resource properties.
<input type="checkbox"/>	Set security for individual items	View and modify security settings for reports, folders, resources, and shared data sources.
<input type="checkbox"/>	View data sources	View shared data source items in the folder hierarchy; and view data source properties.
<input type="checkbox"/>	View folders	View folder items in the folder hierarchy; and view folder properties.
<input type="checkbox"/>	View models	View models in the folder hierarchy, use models as data sources for a report, and run queries against the model to retrieve data.
<input type="checkbox"/>	View reports	View reports and linked reports in the folder hierarchy; view report history snapshots and report properties.
<input type="checkbox"/>	View resources	View resources in the folder hierarchy; and view resource properties.

6.You are designing a strategy for an enterprise reporting solution that uses SQL Server Reporting Services (SSRS).

Many of the SSRS reports will use common utilities and functions, including the following:

- Report utility functions and business logic in code
 - Standardized report formatting properties such as fonts and colors for report branding
- Formatting may change and new functions may be added as the reporting solution evolves.

You need to create a strategy for deploying the formatting and code across the entire enterprise reporting solution. You must also ensure that reports can be easily updated to reflect formatting and function changes.


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

- A. Create a report as a template. Apply standardized formatting to the template. Store code in the Code section of the template.
- B. Build a web service that retrieves formatting properties and runs code. Call the web service through a report dataset.
- C. Store the formatting properties and code in database objects. Use stored procedures to populate a default value for report parameters and map each parameter to a corresponding formatting property.
- D. Create an assembly that contains formatting properties and code. Deploy the assembly on the Reporting Server and reference the assembly from each report.

Answer: D

7.DRAG DROP

You are designing a dataset for a SQL Server Reporting Services (SSRS) report. The report includes the report items displayed in the following graphic.




The dataset is sourced from a commonly used stored procedure in an inventory data mart hosted in a SQL Azure database.


It returns the status for all products across all storage sites. The report must display data for the storage site that is selected by the Site report parameter. You cannot change the stored procedure code.

You need to filter the dataset to use only data specific to the selected site.

How should you configure the filter? To answer, drag the appropriate expression or expressions to the correct location or locations in the answer area. (Answer choices may be used once, more than once, or not all.)



Answer:



Explanation:

To set a filter on the dataset

1. Open a report in Design view.
2. Right-click a dataset in the Report Data pane and then click Dataset Properties. The Dataset Properties dialog box opens.
3. Click Filters. This displays the current list of filter equations. By default, the list is empty.
4. Click Add. A new blank filter equation appears.
5. In Expression, type or select the expression for the field to filter. To edit the expression, click the expression (fx) button.

Box 1: Here we use the Fields expression.

6. From the drop-down box, select the data type that matches the type of data in the expression you created in step 5.
7. In the Operator box, select the operator that you want the filter to use to compare the values in the Expression box and the Value box. The operator you choose determines the number of values that are

used from the next step.

Box 2: we test for equality.

8. In the Value box, type the expression or value against which you want the filter to evaluate the value in Expression.

Box 3: we compare to the value of the Parameter named Site.

9. Click OK.

Reference: How to: Add a Filter (Reporting Services)

8 HOTSPOT


You are designing a SQL Server Integration Services (SSIS) package configuration strategy.

The package configuration must meet the following requirements:

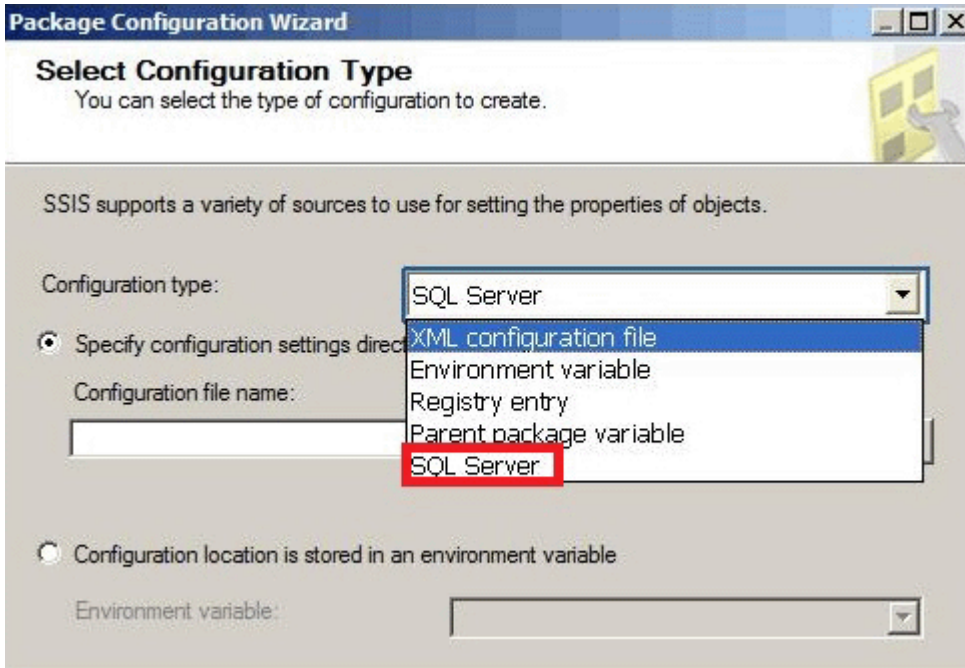
- Include multiple properties in a configuration.
- Support several packages with different configuration settings.

You need to select the appropriate configuration.

Which configuration type should you use? To answer, select the appropriate option from the drop-down list in the dialog box.



Answer:



Explanation:

Package Configuration Types

The following table describes the package configuration types.

* SQL Server table

A table in a SQL Server database contains the configuration. The table can include multiple configurations.

* XML configuration file

An XML file contains the configurations. The XML file can include multiple configurations.

* Environment variable

An environment variable contains the configuration.

* Registry entry

A Registry entry contains the configuration.

* Parent package variable

A variable in the package contains the configuration. This configuration type is typically used to update properties in child packages.

Reference: Package Configurations

9.You are designing a SQL Server Integration Services (SSIS) solution. The solution will contain an SSIS project that includes several SSIS packages. Each SSIS package will define the same connection managers and variables.

You have the following requirements:

- Ensure that the deployment model supports changing the content of connection strings by using parameters at execution time.
- Ensure that the deployment model automatically starts from calls to the catalog.start_execution stored procedure in the SSISDB database.
- Maximize performance at execution time.
- Minimize development effort.

You need to design a solution that meets the requirements.

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

- A. Use a project deployment model. Modify connection manager properties to use project parameters. Ensure that the SSISDB database is created.
- B. Use a project deployment model. Configure connections in an XML configuration file referenced by an environment variable that corresponds to the SQL Server environment of each SSIS package.
- C. Use a package deployment model. Use a SQL Server package configuration with a common filter. Change the contents of the SSIS Configurations table at runtime.
- D. Use a package deployment model. Save each SSIS package to a file share that can be accessed from all environments.

Answer: A

10.DRAG DROP

You are creating a SQL Server Integration Services (SSIS) package to populate a fact table from a source table. The fact table and source table are located in a SQL Azure database. The source table has a price field and a tax field. The OLE DB source uses the data access mode of Table.

You have the following requirements:

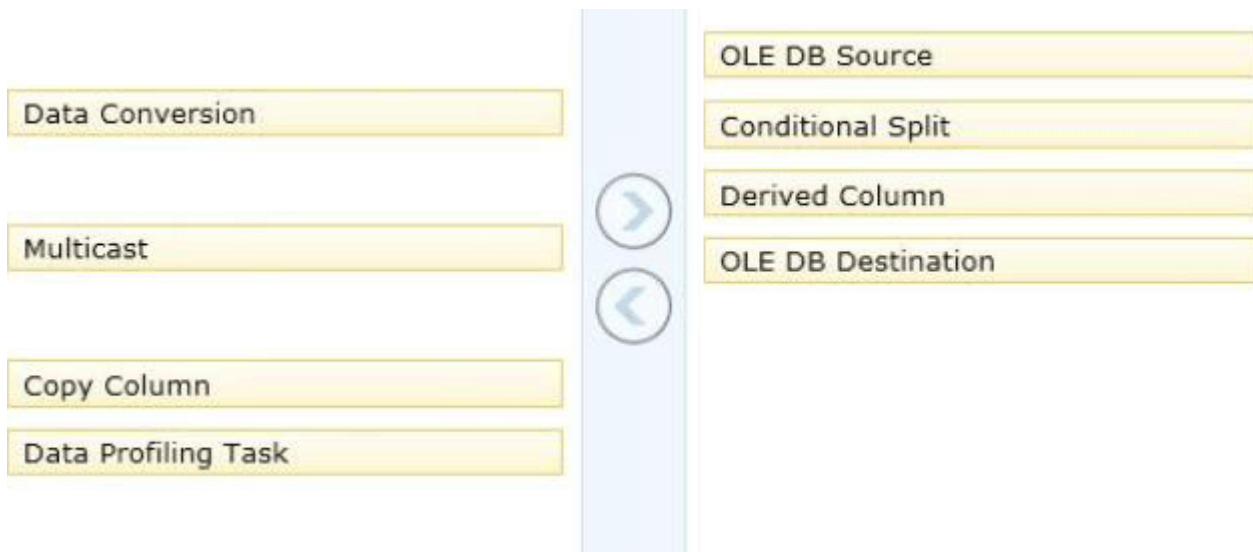
- The fact table must populate a column named TotalCost that computes the sum of the price and tax columns.
- Before the sum is calculated, any records that have a price of zero must be discarded.

You need to create the SSIS package in SQL Server Data Tools.

In what sequence should you order four of the listed components for the data flow task? (To answer, move the appropriate components from the list of components to the answer area and arrange them in the correct order.)

The image shows a list of SSIS components in a drag-and-drop interface. The components listed from top to bottom are: Conditional Split, Data Conversion, OLE DB Destination, Multicast, Derived Column, Copy Column, Data Profiling Task, and OLE DB Source. To the right of the list is a vertical bar containing two circular arrows: the top one points to the right and the bottom one points to the left, indicating that the components can be reordered.

Answer:



Explanation:

* You configure a Data Flow task by adding components to the Data Flow tab. SSIS supports three types of data flow components:

Sources: Where the data comes from

Transformations: How you can modify the data

Destinations: Where you want to put the data

* Creating a data flow includes the following steps:

/ Adding one or more sources to extract data from files and databases, and add connection managers to connect to the sources.

/ Adding the transformations that meet the business requirements of the package. A data flow is not required to include transformations.

Some transformations require a connection manager. For example, the Lookup transformation uses a connection manager to connect to the database that contains the lookup data.

/ Connecting data flow components by connecting the output of sources and transformations to the input of transformations and destinations.

/ Adding one or more destinations to load data into data stores such as files and databases, and adding connection managers to connect to the data sources.

/ Configuring error outputs on components to handle problems.

At run time, row-level errors may occur when data flow components convert data, perform a lookup, or evaluate expressions. For example, a data column with a string value cannot be converted to an integer, or an expression tries to divide by zero. Both operations cause errors, and the rows that contain the errors can be processed separately using an error flow.

/ Include annotations to make the data flow self-documenting.

* The capabilities of transformations vary broadly. Transformations can perform tasks such as updating, summarizing, cleaning, merging, and distributing data. You can modify values in columns, look up values in tables, clean data, and aggregate column values.

* The Data Flow task encapsulates the data flow engine that moves data between sources and destinations, and lets the user transform, clean, and modify data as it is moved. Addition of a Data Flow task to a package control flow makes it possible for the package to extract, transform, and load data. A data flow consists of at least one data flow component, but it is typically a set of connected data flow

components: sources that extract data; transformations that modify, route, or summarize data; and destinations that load data.

11.DRAG DROP

You are designing a SQL Server Integration Services (SSIS) package to execute 12 Transact-SQL (T-SQL) statements on a SQL Azure database.

The T-SQL statements may be executed in any order. The T-SQL statements have unpredictable execution times.

You have the following requirements:

- The package must maximize parallel processing of the T-SQL statements.
- After all the T-SQL statements have completed, a Send Mail task must notify administrators.

You need to design the SSIS package.

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.)

Add a Send Mail task to the control flow. Add a precedence constraint for Completion to the final Execute SQL task and link it to the Send Mail task.

Add a Sequence container to the control flow.

Add a Send Mail task to the control flow. Add a precedence constraint for Completion to the Sequence container and link it to the Send Mail task.

Create precedence constraints for Completion between all the Execute SQL tasks.

Add 12 Execute SQL tasks to the control flow and configure the tasks.

Add 12 Execute SQL tasks to the Sequence container and configure the tasks.

Answer:

Add a Send Mail task to the control flow. Add a precedence constraint for Completion to the final Execute SQL task and link it to the Send Mail task.

Add 12 Execute SQL tasks to the control flow and configure the tasks.

Create precedence constraints for Completion between all the Execute SQL tasks.

Add 12 Execute SQL tasks to the Sequence container and configure the tasks.

Add a Sequence container to the control flow.

Add 12 Execute SQL tasks to the Sequence container and configure the tasks.

Add a Send Mail task to the control flow. Add a precedence constraint for Completion to the Sequence container and link it to the Send Mail task.

Explanation:

The Sequence container defines a control flow that is a subset of the package control flow. Sequence containers group the package into multiple separate control flows, each containing one or more tasks and containers that run within the overall package control flow.

Reference: Sequence Container

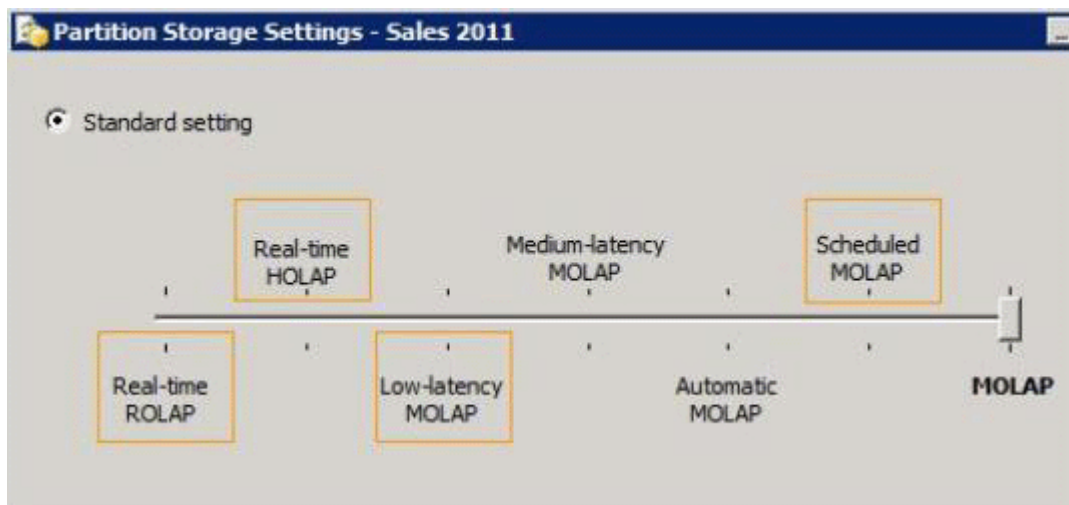
12 HOTSPOT

You are configuring the partition storage settings for a SQL Server Analysis Services (SSAS) cube. The partition storage must meet the following requirements:

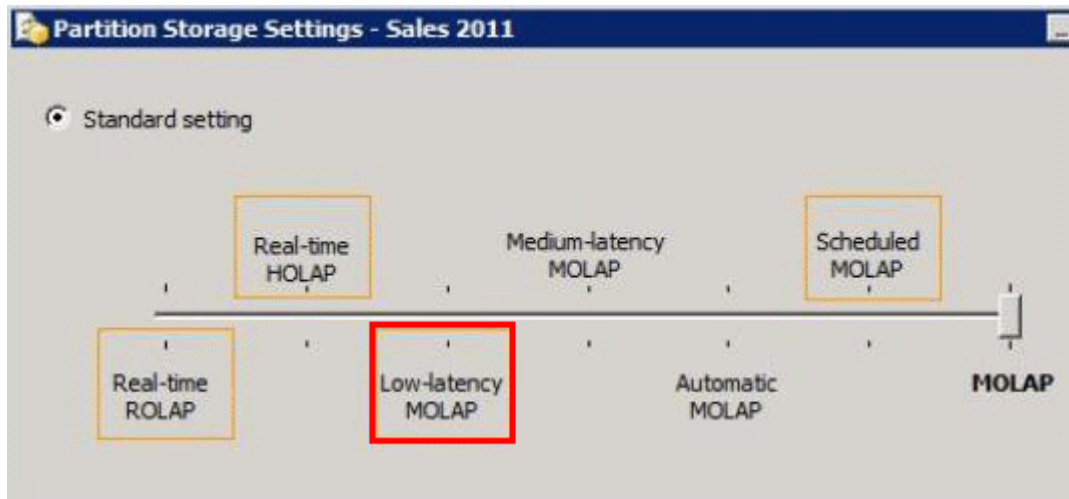
- Optimize the storage of source data and aggregations in the cube.
- Use proactive caching.
- Drop cached data that is more than 30 minutes old.
- Update the cache when data changes, with a silence interval of 10 seconds.

You need to select the partition storage setting.

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



Answer:



Explanation:

<http://msdn.microsoft.com/en-us/library/ms175646.aspx>

Low Latency MOLAP

Detail data and aggregations are stored in multidimensional format. The server listens for notifications of changes to the data and switches to real-time ROLAP while MOLAP objects are reprocessed in a cache. A silence interval of at least **10 seconds** is required before updating the cache. There is an override

interval of 10 minutes if the silence interval is not attained. Processing occurs automatically as data changes with a target latency of **30 minutes** after the first change.

This setting would typically be used for a data source with frequent updates when query performance is somewhat more important than always providing the most current data. This setting automatically processes MOLAP objects whenever required after the latency interval. Performance is slower while the MOLAP objects are being reprocessed.

13 HOTSPOT


A SQL Server Analysis Services (SSAS) cube contains billions of rows of data and is rapidly increasing in size. The cube consists of a single measure group and a single partition. The cube is currently processed by using the Process Full process option.

You have the following requirements to reduce the cube processing time:


- Partition the measure group by month.
- Create a staging table that contains only data which is more recent than the last time the cube was processed.
- Do not include data updates or deletions in the staging table.
- Insert records from the staging table into the appropriate partition.

You need to change the process option to meet the requirements.

Which process option should you choose? To answer, select the appropriate option from the drop-down list in the dialog box.



Answer:

**Explanation:**

ProcessAdd

ProcessAdd applies only to dimensions and partitions.

ProcessAdd is a new processing option for dimensions that did not exist in Analysis Services 2000. It essentially optimizes ProcessUpdate for the scenario where only new members are added. ProcessAdd never deletes or updates existing members. It only adds new members. The user can restrict the dimension table so that ProcessAdd reads only the new rows.

ProcessAdd for partitions is the equivalent of incremental partition processing in Analysis Services 2000. The user typically specifies an alternate fact table or a filter condition pointing to the new rows.

ProcessAdd internally creates a temporary partition, processes it with the specified fact data, and merges it into the target partition.

Reference: Analysis Services 2005 Processing Architecture

14.DRAG DROP

You administer a SQL Server Analysis Services (SSAS) instance. You need to capture a continuous log of detailed event and subevent durations and custom trace events from queries executed in the SSAS instance.

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.)

Write a Multidimensional Expressions (MDX) script to query the DISCOVER_TRACES dynamic management view (DMV).

Write an XMLA script to log the extended events of the trace.

Launch SQL Server Profiler and connect to the instance.

Execute the script.

Launch SQL Server Management Studio and connect to the instance.

Configure the trace to save to a SQL Server database table.

➤
➤

Answer:

Write a Multidimensional Expressions (MDX) script to query the DISCOVER_TRACES dynamic management view (DMV).

Write an XMLA script to log the extended events of the trace.

Launch SQL Server Profiler and connect to the instance.

Execute the script.

Launch SQL Server Management Studio and connect to the instance.

Configure the trace to save to a SQL Server database table.

Explanation:

* Auditing an instance of SQL Server or a SQL Server database involves tracking and logging events that occur on the system. The SQL Server Audit object collects a single instance of server- or database-level actions and groups of actions to monitor. The audit is at the SQL Server instance level. You can have multiple audits per SQL Server instance. The Server Audit Specification object belongs to an audit. You can create one server audit specification per audit, because both are created at the SQL Server instance scope.

* Trace events can be started and captured using SQL Server Profiler, , or can be started from an XMLA command as SQL Server Extended Events and later analyzed.

* Extended Event tracing is enabled using a similar XMLA create object script.

Reference: Analysis Services Trace Events

15.DRAG DROP

You plan to deploy a SQL Server Integration Services (SSIS) project by using the project deployment model. You need to monitor control flow tasks to determine whether any of them are running longer than usual.

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.)

Write a query against the **catalog.operation_messages** view. Add a calculation to the query to compare durations to the **catalog.executables** view.

Execute the query.

Write a query against the **catalog.execution_component_phases** view. Add a calculation to the query to compare durations to the **catalog.executables** view.

Connect to the **SSISDB** database.

Connect to the **msdb** database.

Write a query against the **catalog.execution_component_phases** view. Add a calculation to the query to compare durations to the **catalog.executions** view.

Write a query against the **catalog.operation_messages** view. Add a calculation to the query to compare durations to the **catalog.executions** view.



Answer:

Write a query against the **catalog.operation_messages** view. Add a calculation to the query to compare durations to the **catalog.executables** view.

Connect to the **SSISDB** database.

Write a query against the **catalog.execution_component_phases** view. Add a calculation to the query to compare durations to the **catalog.executables** view.

Execute the query.



Connect to the **msdb** database.

Write a query against the **catalog.execution_component_phases** view. Add a calculation to the query to compare durations to the **catalog.executions** view.

Write a query against the **catalog.operation_messages** view. Add a calculation to the query to compare durations to the **catalog.executions** view.

Explanation:

* execution_component_phases

Displays the time spent by a data flow component in each execution phase.

* The following example uses the catalog.execution_component_phases view to find the total amount of time that a specific package has spent executing in all phases (**active_time**), and the total elapsed time for the package (**total_time**).

use SSISDB

```
select package_name, task_name, subcomponent_name, execution_path,
SUM(DATEDIFF(ms,start_time,end_time)) as active_time,
DATEDIFF(ms,min(start_time), max(end_time)) as total_time
```

```
from catalog.execution_component_phases
where execution_id = 1841
group by package_name, task_name, subcomponent_name, execution_path
order by package_name, task_name, subcomponent_name, execution_path
* catalog.executables
```

This view displays a row for each executable in the specified execution.

An executable is a task or container that you add to the control flow of a package.

*(incorrect) catalog.executions (SSISDB Database)

Displays the instances of package execution in the Integration Services catalog. Packages that are executed with the Execute Package task run in the same instance of execution as the parent package.

This view displays a row for each instance of execution in the catalog.

* (incorrect) catalog.operation_messages

Displays messages that are logged during operations in the Integration Services catalog.

This view displays a row for each message that is logged during an operation in the catalog. The message can be generated by the server, by the package execution process, or by the execution engine.

Reference: catalog.execution_component_phases

Reference: catalog.executables

16.You are designing a partitioning strategy for a large fact table in a data warehouse. Tens of millions of new records are loaded into the data warehouse weekly, outside of business hours. Most queries are generated by reports and by cube processing. Data is frequently queried at the day level and occasionally at the month level. You need to partition the table to maximize the performance of queries.

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

- A. Partition the fact table by month, and compress each partition.
- B. Partition the fact table by week.
- C. Partition the fact table by year.
- D. Partition the fact table by day, and compress each partition.

Answer: D

17.You are designing an extract, transform, load (ETL) process for loading data from a SQL Server database into a large fact table in a data warehouse each day with the prior day's sales data.

The ETL process for the fact table must meet the following requirements:

- Load new data in the shortest possible time.
- Remove data that is more than 36 months old.
- Ensure that data loads correctly.
- Minimize record locking.
- Minimize impact on the transaction log.

You need to design an ETL process that meets the requirements.

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

- A. Partition the destination fact table by date. Insert new data directly into the fact table and delete old data directly from the fact table.
- B. Partition the destination fact table by date. Use partition switching and staging tables both to remove old data and to load new data.
- C. Partition the destination fact table by customer. Use partition switching both to remove old data and to

load new data into each partition.

D. Partition the destination fact table by date. Use partition switching and a staging table to remove old data. Insert new data directly into the fact table.

Answer: B

18.DRAG DROP

You are administering a SQL Server Analysts Services (SSAS) database on a server. The database hosts a financial cube based on a SQL Azure database. You need to grant write access to the financial cube for all users in the group USA\PowerUsers.

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.)

In SQL Server Management Studio (SSMS), connect to the SSAS instance on the server.	
Create a new role for the database.	
Add the group USA\PowerUsers as a SQL Server login to the server.	
In SQL Server Management Studio (SSMS), connect to the Database Engine instance on the server.	
Add the USA\PowerUsers group to the role. Set the cube access for the role to Select .	
Add the USA\PowerUsers group to the role. Set the cube access for the role to Read/Write .	
Add the USA\PowerUsers group to the role. Set the cube access for the role to Read and Process .	

Answer:

		In SQL Server Management Studio (SSMS), connect to the SSAS instance on the server.
		Create a new role for the database.
		Add the USA\PowerUsers group to the role. Set the cube access for the role to Read and Process .
Add the group USA\PowerUsers as a SQL Server login to the server.		
In SQL Server Management Studio (SSMS), connect to the Database Engine instance on the server.		
Add the USA\PowerUsers group to the role. Set the cube access for the role to Select .		
Add the USA\PowerUsers group to the role. Set the cube access for the role to Read/Write .		

Explanation:

* A member of the server role for Microsoft SQL Server Analysis Services, or a member of a database role that has Full Control (Administrator) permissions in a particular database, can create a database role that only has permission to process specified objects within the database. Giving a database role permission to process a database object lets an administrator delegate the task of processing certain objects, without also granting extraneous permissions to the user who is performing the processing.

* To give a database role permission to process a cube

1. In SQL Server Management Studio, connect to the instance of Analysis Services, expand Roles for the appropriate database in Object Explorer, and then double-click a database role (or right-click Roles and select New Role to create a new database role). If this is a new role, make sure that you enter a name for the role in the Role name box.
2. Click Cubes in the Select a Page pane, locate the cube in the Cube list, and then select the Process check box for the cube.
3. Click the OK button.

* There is no write permissions on a cube.

Reference: Grant Process Permissions on an Analysis Services Multidimensional Database

19.DRAG DROP

You are validating whether a SQL Server Integration Services (SSIS) package named Master.dtsx in the SSIS catalog is executing correctly. You need to display the number of rows in each buffer passed between each data flow component of the package.

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.)

Execute a SQL statement with a package name of Master.dtsx against the **catalog.executions** view and return its execution ID.

Run the Master.dtsx package with the logging level set to **Performance**.

Execute a SQL statement with the execution ID equal to the previously retrieved execution ID against the **catalog.execution_data_statistics** view and return the **rows_sent** column values for all the rows.

Run the Master.dtsx package with the logging level set to **Verbose**.

Execute a SQL statement with a package name of Master.dtsx against the **msdb..sysssislog** table and return its execution ID.

Answer:

Execute a SQL statement with a package name of Master.dtsx against the **catalog.executions** view and return its execution ID.

Run the Master.dtsx package with the logging level set to **Performance**.

Run the Master.dtsx package with the logging level set to **Verbose**.

Execute a SQL statement with a package name of Master.dtsx against the **msdb..sysssislog** table and return its execution ID.

Execute a SQL statement with the execution ID equal to the previously retrieved execution ID against the **catalog.execution_data_statistics** view and return the **rows_sent** column values for all the rows.

20. You are creating a Multidimensional Expressions (MDX) calculation for Projected Revenue in a cube. For Customer A, Projected Revenue is defined as 150 percent of the Total Sales for the customer. For all other customers, Projected Revenue is defined as 110 percent of the Total Sales for the customer. You need to calculate the Projected Revenue as efficiently as possible.

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

- A.

```
CREATE MEMBER CurrentCube.[Measures].[Projected Revenue]
AS [Measures].[Total Sales];
SCOPE ([Customer].[Customer Name].MEMBERS, [Measures].[Projected Revenue]);
    [Measures].[Total Sales] * 1.1;
    IF [Customer].[Customer Name].CurrentMember.Name = "Customer A"
        THEN [Measures].[Total Sales] * 1.5
    END IF;
END SCOPE;
```
- B.

```
CREATE MEMBER CurrentCube.[Measures].[Projected Revenue]
AS CASE WHEN [Customer].[Customer Name].CurrentMember.Name = "Customer A"
THEN [Measures].[Total Sales] * 1.5
ELSE [Measures].[Total Sales] * 1.1 END
```
- C.

```
CREATE MEMBER CurrentCube.[Measures].[Projected Revenue]
AS [Measures].[Total Sales] * 1.1;
SCOPE ([Customer].[Customer Name].&[Customer A],[Measures].[Projected Revenue]);
    THIS = [Measures].[Total Sales] * 1.5;
END SCOPE;
```
- D.

```
CREATE MEMBER CurrentCube.[Measures].[Projected Revenue]
AS [Measures].[Total Sales] * 1.1;
SCOPE ([Customer].[Customer Name].MEMBERS, [Measures].[Projected Revenue]);
    [Customer].[Customer Name].&[Customer A] = [Measures].[Total Sales] * 1.5;
END SCOPE;
```

- A. Option A
B. Option B
C. Option C

D. Option D

Answer: C