

Microsoft

Exam 70-467

Designing Business Intelligence Solutions with Microsoft SQL Server 2012

Verson: Demo

[Total Questions: 10]

Topic break down

Topic	No. of Questions
Topic 2: Contoso, Ltd Case A	1
Topic 3: Data Architect	1
Topic 5: Contoso, Ltd Case B	1
Topic 7: Mix Questions	7

Topic 2, Contoso, Ltd Case A

General Background

You are the SQL Server Administrator for Contoso, Ltd. You have been tasked with upgrading all existing SQL Server instances to SQL Server 2012.

Technical Background

The corporate environment includes an Active Directory Domain Services (AD DS) domain named contoso.com. The forest and domain levels are set to Windows Server 2008. All default containers are used for computer and user accounts. All servers run Windows Server 2008 R2 Service Pack 1 (SP1). All client computers run Windows 7 Professional SP1. All servers and client computers are members of the contoso.com domain.


The current SQL Server environment consists of a single instance failover cluster of SQL Server 2008 R2 Analysis Services (SSAS). The virtual server name of the cluster is SSASCluster. The cluster includes two nodes: Node1 and Node2. Node1 is currently the active node. In anticipation of the upgrade, the prerequisites and shared components have been upgraded on both nodes of the cluster, and each node was rebooted during a weekly maintenance window.

A single-server deployment of SQL Server 2008 R2 Reporting Services (SSRS) in native mode is installed on a server named SSRS01. The Reporting Server service is configured to use a domain service account. SSRS01 hosts reports that access the SSAS databases for sales data as well as modeling data for the Research team. SSRS01 contains 94 reports used by the organization. These reports are generated continually during business hours. Users report that report subscriptions on SSRS01 are not being delivered. You run the reports on demand from Report Manager and find that the reports render as expected.

A new server named SSRS02 has been joined to the domain, SSRS02 will host a single-server deployment of SSRS so that snapshots of critical reports are accessible during the upgrade.

The server configuration is shown in the exhibit. (Click the Exhibit button.)

Server Configuration



The production system includes three SSAS databases that are described in the following table.

Database name	Size
Customer Sales	350 MB
Manufacturing	1.2 GB
Research	620 MB

All SSAS databases are backed up once a day, and backups are stored offsite.

Business Requirements

After the upgrade users must be able to perform the following tasks:

- ✍ Ad-hoc analysis of data in the SSAS databases by using the Microsoft Excel PivotTable client.
- ✍ Daily operational analysis by executing a custom application that uses ADOMD.NET and existing Multidimensional Expressions (MDX) queries.

The detailed data must be stored in the model.

Technical Requirements

You need to minimize downtime during the SSASCluster upgrade. The upgrade must minimize user intervention and administrative effort.

The upgrade to SQL Server 2012 must maximize the use of all existing servers, require the least amount of administrative effort, and ensure that the SSAS databases are operational as soon as possible.

You must implement the highest level of domain security for client computers connecting to SSRS01. The SSRS instance on SSRS01 must use Kerberos delegation to connect to the SSAS databases. Email notification for SSRS01 has not been previously configured. Email notification must be configured to use the SMTP server SMTP01 with a From address of reports@contoso.com. Report distribution must be secured by using SSL and must be limited to the contoso.com domain.

You have the following requirements for SSRS02:

- ✍ Replicate the SSRS01 configuration.
- ✍ Ensure that all current reports are available on SSRS02.
- ✍ Minimize the performance impact on SSR501.

In preparation for the upgrade, the SSRS-related components have been installed on the new SSRS02 server by using the Reporting Services file-only installation mode. The Reporting Services databases have been restored from SSRS01 and configured appropriately.

You must design a strategy to recover the SSRS instance on SSRS01 in the event of a system failure. The strategy must ensure that SSRS can be recovered in the minimal amount of time and that reports are available as soon as possible. Only functional components must be recovered.

SSRS02 is the recovery server and is running the same version of SSRS as SSRS01. A full backup of the SSRS databases on SSRS01 is performed nightly. The report server configuration files, custom assemblies, and extensions on SSRS02 are manually synchronized with SSRS01.

Prior to implementing the upgrade to SQL Server 2012, you must back up all existing

SSAS databases.

Databases on SSRS01 is performed nightly. The report server configuration files, custom assemblies, and extensions on SSRS02 are manually synchronized with SSRS01.

Prior to implementing the upgrade to SQL Server 2012, you must back up all existing SSAS databases. The backup must include only the partitioning, metadata, and aggregations to minimize the processing time required when restoring the databases. You must minimize processing time and the amount of disk space used by the backups.

Before upgrading SSAS on the SSASCluster, all existing databases must be moved to a temporary staging server named SSAS01 that hosts a default instance of SQL Server 2012 Analysis Services. This server will be used for testing client applications connecting to SSAS 2012, and as a disaster recovery platform during the upgrade. You must move the databases by using the least amount of administrative effort and minimize downtime. All SSAS databases other than the Research database must be converted to tabular BI Semantic Models (BISMs) as part of the upgrade to SSAS 2012. The Research team must have access to the Research database for modeling throughout the upgrade. To facilitate this, you detach the Research database and attach it to SSAS01.

While testing the Research database on SSAS01, you increase the compatibility level to 1100. You then discover a compatibility issue with the application. You must roll back the compatibility level of the database to 1050 and retest.

After completing the upgrade, you must do the following:

1. Design a role and assign an MDX expression to the Allowed member set property of the Customer dimension to allow sales representatives to browse only members of the Customer dimension that are located in their sales regions. Use the sales representatives' logins and minimize impact on performance.
2. Deploy a data model to allow the ad-hoc analysis of data. The data model must be cached and source data from an OData feed.

Question No : 1 - (Topic 2)

You need to configure security for the SSRS instance on SSRS01 to connect to SSAS and minimize downtime.

What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Register a service principal name for the Report Server service.
- B. Register a service principal name for the Analysis Services service.
- C. Restart the IIS service.
- D. Configure SSRS01 to use the Negotiate authentication type.
- E. Configure SSRS01 to use the Custom authentication type.

Answer: A,D

Explanation: A (not B): If you are deploying Reporting Services in a network that uses the Kerberos protocol for mutual authentication, you must create a Service Principal Name (SPN) for the Report Server service if you configure it to run as a domain user account.

D (not E):

* See step 6 below.

To register an SPN for a Report Server service running as a domain user

- ✍ Install Reporting Services and configure the Report Server service to run as a domain user account. Note that users will not be able to connect to the report server until you complete the following steps.
- ✍ Log on to the domain controller as domain administrator.
- ✍ Open a Command Prompt window.
- ✍ Copy the following command, replacing placeholder values with actual values that are valid for your network:
 - ✍ `Setspn -a http/<computer-name>.<domain-name>:<port><domain-user-account>`
 - ✍ Run the command.
 - ✍ Open the `RsReportServer.config` file and locate the `<AuthenticationTypes>` section. Add `<RSWindowsNegotiate/>` as the first entry in this section to enable NTLM.

* `RSWindowsNegotiate`. If you initially set the Windows service account for the report server to `NetworkService` or `LocalSystem` in Reporting Services Configuration Manager, `RSWindowsNegotiate` is added to the `RSReportServer.config` file as the default setting. With this setting, the report server can accept requests from client applications requesting Kerberos or NTLM authentication. If Kerberos is requested and the authentication fails, the report server switches to NTLM authentication and prompts the user for credentials unless the network is configured to manage authentication transparently.

Using `RSWindowsNegotiate` is your best option because it provides the greatest flexibility for multiple clients in an intranet environment.

Not C: IIS is not mention in this scenario.

Note:

* From scenario:

/ A single-server deployment of SQL Server 2008 R2 Reporting Services (SSRS) in native mode is installed on a server named `SSRS01`. The Reporting Server service is configured to use a domain service account.

Reference: Register a Service Principal Name (SPN) for a Report Server

Topic 3, Data Architect

General Background

You are a Data Architect for a company that uses SQL Server 2012 Enterprise edition.

You have been tasked with designing a data warehouse that uses the company's financial database as the data source. From the data warehouse, you will develop a cube to simplify the creation of accurate financial reports and related data analysis.

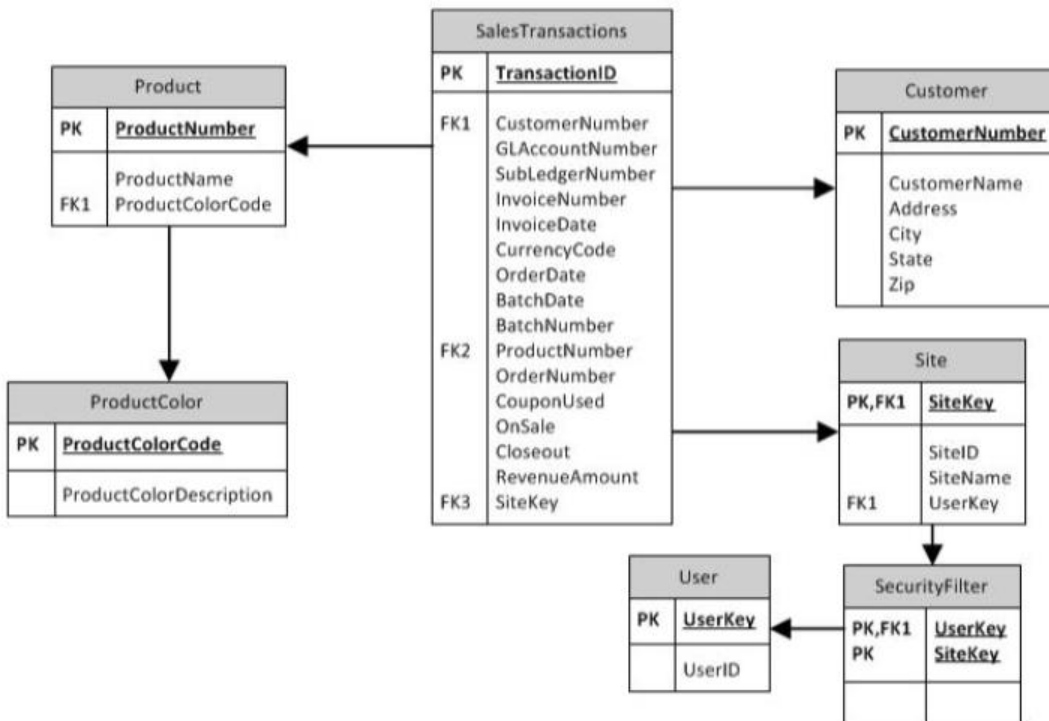
Background

You will utilize the following three servers:

- ✍ ServerA runs SQL Server Database Engine. ServerA is a production server and also hosts the financial database.
- ✍ ServerB runs SQL Server Database Engine, SQL Server Analysis Services (SSAS) in multidimensional mode, SQL Server Integration Services (SSIS), and SQL Server Reporting Services (SSRS).
- ✍ ServerC runs SSAS in multidimensional mode.
- ✍ The financial database is used by a third-party application and the table structures cannot be modified.

The relevant tables in the financial database are shown in the exhibit. (Click the Exhibit button.)

Financial Database tables



The SalesTransactions table is 500 GB and is anticipated to grow to 2 TB. The table is partitioned by month. It contains only the last five years of financial data. The CouponUsed, OnSale, and Closeout columns contain only the values Yes or No. Each of the other tables is less than 10 MB and has only one partition.

The SecurityFilter table specifies the sites to which each user has access.

Business Requirements

The extract, transform, load (ETL) process that updates the data warehouse must run daily between 8:00 P.M. and 5:00 A.M. so that it doesn't impact the performance of ServerA during business hours. The cube data must be available by 8:00 A.M.

The cube must meet the following business requirements:

- ✍ Ensure that reports display the most current information available.
- ✍ Allow fast access to support ad-hoc reports and data analysis.

Business Analysts will access the data warehouse tables directly, and will access the cube by using SSRS, Microsoft Excel, and Microsoft SharePoint Server 2010 PerformancePoint Services. These tools will access only the cube and not the data warehouse.

Technical Requirements

SSIS solutions must be deployed by using the project deployment model.

You must develop the data warehouse and store the cube on ServerB. When the number of concurrent SSAS users on ServerB reaches a specific number, you must scale out SSAS to ServerC and meet following requirements:

- ✍ Maintain copies of the cube on ServerB and ServerC.
- ✍ Ensure that the cube is always available on both servers.
- ✍ Minimize query response time.

The cube must meet the following technical requirements:

- ✍ The cube must be processed by using an SSIS package.
- ✍ The cube must contain the prior day's data up to 8:00 P.M. but does not need to contain same-day data.
- ✍ The cube must include aggregation designs when it is initially deployed.
- ✍ A product dimension must be added to the cube. It will contain a hierarchy comprised of product name and product color.

Because of the large size of the SalesTransactions table, the cube must store only aggregations—the data warehouse must store the detailed data. Both the data warehouse and the cube must minimize disk space usage.

As the cube size increases, you must plan to scale out to additional servers to minimize processing time.

The data warehouse must use a star schema design. The table design must be as denormalized as possible. The history of changes to the Customer table must be tracked in the data warehouse. The cube must use the data warehouse as its only data source.

Security settings on the data warehouse and the cube must ensure that queries against the SalesTransactions table return only records from the sites to which the current user has access.

The ETL process must consist of multiple SSIS packages developed in a single project by

using the least amount of effort. The SSIS packages must use a database connection string that is set at execution time to connect to the financial database. All data in the data warehouse must be loaded by the SSIS packages.

You must create a Package Activity report that meets the following requirements:

- ✍ Track SSIS package execution data (including package name, status, start time, end time, duration, and rows processed).
- ✍ Use the least amount of development effort.

Question No : 2 - (Topic 3)

You need to implement the aggregation designs for the cube.

What should you do?

- A. Use the CREATE CACHE statement.
- B. Use the Aggregation Design Wizard.
- C. Create relational indexes on the source tables.
- D. Use the Usage-Based Optimization Wizard.

Answer: B

Topic 5, Contoso, Ltd Case B

General Background

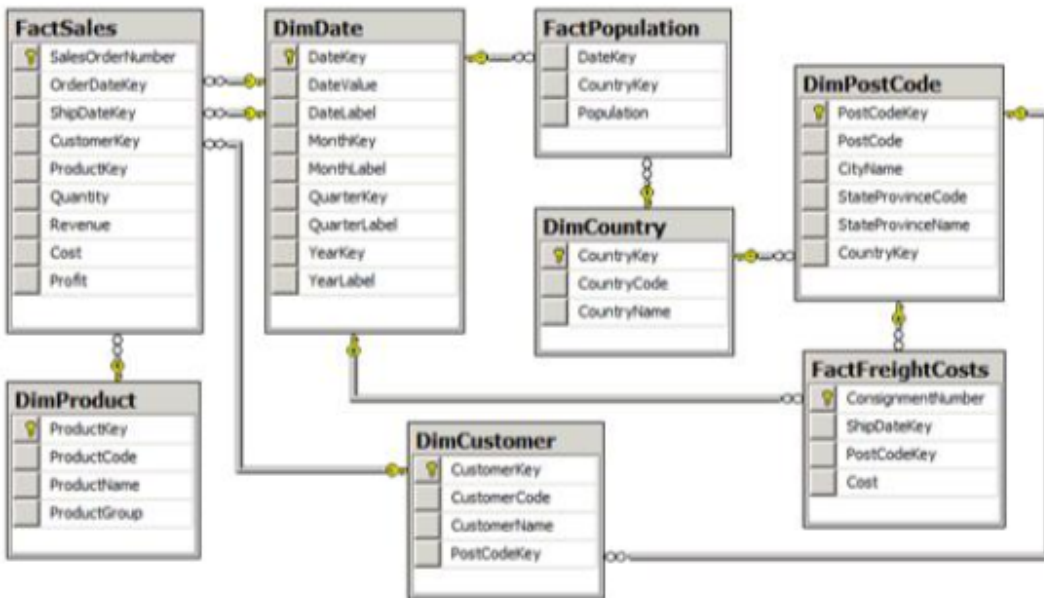
You are the business intelligence (BI) solutions architect for Contoso, Ltd, an online retailer.

You produce solutions by using SQL Server 2012 Business Intelligence edition and Microsoft SharePoint Server 2010 Service Pack 1 (SP1) Enterprise edition.

A SharePoint farm has been installed and configured for intranet access only. An Internet-facing web server hosts the company's public e-commerce website. Anonymous access is not configured on the Internet-facing web server.

Data Warehouse

The data warehouse is deployed on a SQL Server 2012 relational database instance. The data warehouse is structured as shown in the following diagram.



The following Transact-SQL (T-SQL) script is used to create the FactSales and FactPopulation tables:

```
CREATE TABLE [dbo].[FactSales]
(
    [SalesOrderNumber] NCHAR(10) PRIMARY KEY
    , [OrderDateKey] INT NOT NULL
    , [ShipDateKey] INT NOT NULL
    , [CustomerKey] INT NOT NULL
    , [ProductKey] INT NOT NULL
    , [Quantity] INT NOT NULL
    , [Revenue] SMALLMONEY NOT NULL
    , [Cost] SMALLMONEY NOT NULL
    , [Profit] AS ([Revenue]-[Cost])
);
GO

CREATE TABLE [dbo].[FactPopulation]
(
    [DateKey] INT NOT NULL
    , [CountryKey] INT NOT NULL
    , [Population] INT NOT NULL
);
GO
```

The FactPopulation table is loaded each year with data from a Windows Azure Marketplace commercial dataset. The table contains a snapshot of the population values for all countries of the world for each year. The world population for the last year loaded exceeds 6.8 billion people.

ETL Process

SQL Server Integration Services (SSIS) is used to load data into the data warehouse. All SSIS projects are developed by using the project deployment model.

A package named StageFactSales loads data into a data warehouse staging table. The package sources its data from numerous CSV files exported from a mainframe system. The CSV file names begin with the letters GLSD followed by a unique numeric identifier that never exceeds six digits. The data content of each CSV file is identically formatted.

A package named LoadFactFreightCosts sources data from a Windows Azure SQL Database database that has data integrity problems. The package may retrieve duplicate rows from the database.

The package variables of all packages have the RaiseChangedEvent property set to true. A package-level event handler for the OnVariableValueChanged event consists of an Execute SQL task that logs the System::VariableName and System::VariableValue variables.

Data Models

SQL Server Analysis Services (SSAS) is used to host the Corporate BI multidimensional database. The Corporate BI database contains a single data source view named Data Warehouse. The Data Warehouse data source view consists of all data warehouse tables. All data source view tables have been converted to named queries.

The Corporate BI database contains a single cube named Sales Analysis and three database dimensions: Date, Customer and Product. The dimension usage for the Sales Analysis cube is as shown in the following image.

Measure Groups			
Dimensions	Sales	Freight Costs	Population
Date (Order Date)	Date		Year
Date (Ship Date)	Date	Date	
Customer	Customer	Post Code	Country
Product	Product		

The Customer dimension contains a single multi-level hierarchy named Geography. The structure of the Geography hierarchy is shown in the following image.



The Sales Analysis cube's calculation script defines one calculated measure named Sales Per Capita. The calculated measure expression divides the Revenue measure by the Population measure and multiplies the result by 1,000. This calculation represents revenue per 1,000 people.

The Sales Analysis cube produces correct Sales Per Capita results for each country of the world; however, the Grand Total for all countries is incorrect, as shown in the following image (rows 2-239 have been hidden).

	A	B	C
1	Row Labels [dropdown]	Revenue	Sales Per Capita
240	Western Sahara	253	0.46
241	Yemen	12,345	0.52
242	Zambia	1,700	0.13
243	Zimbabwe	16,000	1.25
244	Grand Total	46,030,298	-26.76

A role named Analysts grants Read permission for the Sales Analysis cube to all sales and marketing analysts in the company.

SQL Server Reporting Services (SSRS) is configured in SharePoint integrated mode. All reports are based on shared data sources.

Corporate logo images used in reports were originally configured as data-bound images sourced from a SQL Server relational database table. The image data has been exported to JPG files. The image files are hosted on the Internet-facing web server. All reports have been modified to reference the corporate logo images by using the fully qualified URLs of the image files. A red X currently appears in place of the corporate logo in reports.

Users configure data alerts on certain reports. Users can view a report named Sales

Profitability on demand; however, notification email messages are no longer being sent when Sales Profitability report data satisfies alert definition rules. The alert schedule settings for the Sales Profitability report are configured as shown in the following image.

Schedule settings

Recurrence pattern:

Hourly - every 1 hour(s)

Advanced

Start alert on: 1/1/2012 2:00:00 AM

Stop alert on: 12/31/2020

Send message only if alert results change

Business Requirements

Data Models

Users must be able to:

- Provide context to measures and filter measures by using all related data warehouse dimensions.
- Analyze measures by order date or ship date.

Additionally, users must be able to add a measure named Sales to the report canvas by clicking only once in the Power View field list. The Sales measure must allow users to analyze the sum of the values in the Revenue column of the FactSales data warehouse table. Users must be able to change the aggregation function of the Sales measure.

Analysis and Reporting

A sales manager has requested the following query results from the Sales Analysis cube for the 2012 fiscal year:

- Australian postal codes and sales in descending order of sales.
- Australian states and the ratio of sales achieved by the 10 highest customer sales made for each city in that state.

Technical Requirements

ETL Processes

If an SSIS package variable value changes, the package must log the variable name and the new variable value to a custom log table.

The StageFactSales package must load the contents of all files that match the file name pattern. The source file name must also be stored in a column of the data warehouse staging table.

In the design of the LoadFactSales package, if a lookup of the dimension surrogate key value for the product code fails, the row details must be emailed to the data steward and written as an error message to the SSIS catalog log by using the public API.

You must configure the LoadFactFreightCosts package to remove duplicate rows, by using the least development effort.

Data Models

Users of the Sales Analysis cube frequently filter on the current month's data. You must ensure that queries to the Sales Analysis cube default to the current month in the Order Date dimension for all users.

You must develop and deploy a tabular project for the exclusive use as a Power View reporting data source. The model must be based on the data warehouse. Model table names must exclude the Dim or Fact prefixes. All measures in the model must format values to display zero decimal places.

Analysis and Reporting

Reports must be developed that combine the SSIS catalog log messages with the package variable value changes.

Question No : 3 - (Topic 5)

You need to develop the tabular project to support the date analysis requirements.

What should you do?

A.

- Create one date table named Date.
- Create an active relationship between the DateKey column of the Date table and the OrderDateKey column of the Sales table.
- Create an inactive relationship between the DateKey column of the Date table and the ShipDateKey column of the Sales table.

B.

- Create two date tables, one named Order Date and one named Ship Date.
- Create an active relationship between the DateKey column of the Order Date table and the OrderDateKey column of the Sales table.
- Create an inactive relationship between the DateKey column of the Ship Date table and the ShipDateKey column of the Sales table.

C.

- Create one date table named Date.
- Create an active relationship between the DateKey column of the Date table and the ShipDateKey column of the Sales table.
- Create an inactive relationship between the DateKey column of the Date table and the

OrderDateKey column of the Sales table.

D.

- Create two date tables, one named Order Date and one named Ship Date.
- Create an active relationship between the DateKey column of the Order Date table and the OrderDateKey column of the Sales table.
- Create an active relationship between the DateKey column of the Ship Date table and the ShipDateKey column of the Sales table.

Answer: C

Topic 7, Mix Questions

Question No : 4 DRAG DROP - (Topic 7)

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 **Basic**.

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 **catalog.event_messages** view 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 Verbose .
Run the Master.dtsx package with the logging level set to Basic .	
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.	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 Verbose .	
Execute a SQL statement with a package name of Master.dtsx against the catalog.event_messages view 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.

Explanation: Box 1:

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

Box 2:

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

Box 3:

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.

Note:

* You are going to become very very familiar indeed with [catalog].[executions]. It is a view that provides a record of all package executions on the server and, most importantly, it contains [execution_id] – the identifier for each execution and the field to which all other objects herein will be related.

Question No : 5 - (Topic 7)

The Sales database includes a table named Products that contains 42 columns, and two tables that record all product sales transactions. The database does not store aggregated data. The database is continually maintained to provide optimal indexing. The database server is sized appropriately.

The Marketing department requests a delimited text file that includes seven fields containing product information from the Products table.

You have the following requirements:

- ✍ Display the ProductID in the first column. Organize the remaining columns alphabetically by column heading.
- ✍ Sort rows in ProductID order.
- ✍ Aggregate product units sold.
- ✍ Minimize the time required to return the results.

You create a SQL Server Integration Services (SSIS) package and add the appropriate connection managers and a Data Flow task.

You need to develop the Data Flow task.

What should you add to the Data Flow task? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. A table-based OLE DB source; a Script component; and a Flat File destination.
- B. Only a SQL command-based OLE DB source and a Flat File destination.
- C. A table-based OLE DB source; Lookup, Aggregate, and Sort transformations; and a Flat File destination.
- D. A SQL command-based OLE DB source; Aggregate and Sort transformations; and a Flat File destination.

Answer: D

Question No : 6 HOTSPOT - (Topic 7)

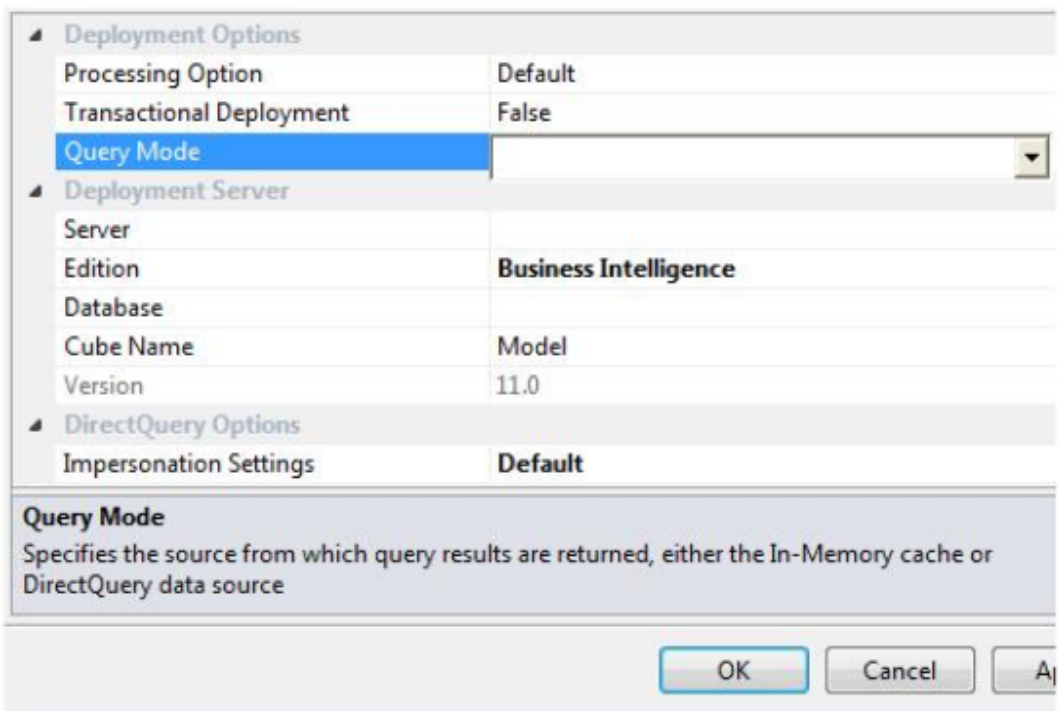
You are developing a SQL Server Analysis Services (SSAS) tabular project. The DirectQuery Mode property of the data model has been set to On. The current partition for the table is configured as the DirectQuery partition.

Data is loaded into a data mart throughout the day and must be available in the data model as soon as possible. The model must support querying by clients that issue both Data Analysis Expressions (DAX) queries and Multidimensional Expressions (MDX) queries. Clients issuing DAX queries must receive real-time query results.

You need to select the appropriate Query Mode property for deployment.

Which query mode should you select?

To answer, select the appropriate query mode from the drop-down list in the dialog box.



Deployment Options	
Processing Option	Default
Transactional Deployment	False
Query Mode	
Deployment Server	
Server	DirectQuery
Edition	DirectQuery with In-Memory
Database	In-Memory
Cube Name	In-Memory with DirectQuery
Version	Model
Version	11.0
DirectQuery Options	
Impersonation Settings	Default

Query Mode
Specifies the source from which query results are returned, either the In-Memory cache or DirectQuery data source

OK Cancel Ap

Answer:

Deployment Options	
Processing Option	Default
Transactional Deployment	False
Query Mode	
Deployment Server	
Server	DirectQuery
Edition	DirectQuery with In-Memory
Database	In-Memory
Cube Name	In-Memory with DirectQuery
Version	Model
Version	11.0
DirectQuery Options	
Impersonation Settings	Default

Query Mode
Specifies the source from which query results are returned, either the In-Memory cache or DirectQuery data source

OK Cancel Ap

Question No : 7 - (Topic 7)

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

Question No : 8 - (Topic 7)

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

Question No : 9 DRAG DROP - (Topic 7)

You administer a SQL Server Reporting Services (SSRS) infrastructure. Data alerts are provisioned.

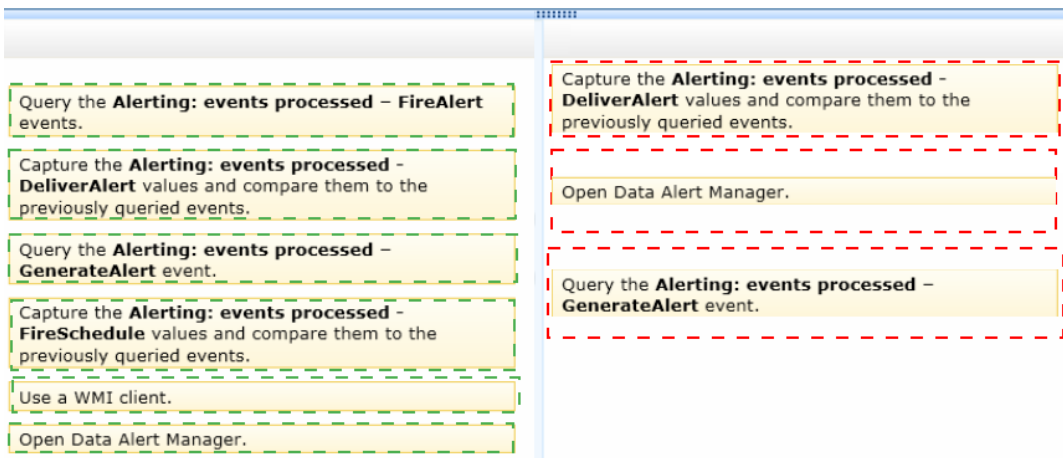
You need to monitor the number of report data feed processing events during every automated execution.

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.)

The screenshot shows a drag-and-drop interface with a list of actions on the left and an empty answer area on the right. The actions are:

- Query the **Alerting: events processed - FireAlert** events.
- Capture the **Alerting: events processed - DeliverAlert** values and compare them to the previously queried events.
- Query the **Alerting: events processed - GenerateAlert** event.
- Capture the **Alerting: events processed - FireSchedule** values and compare them to the previously queried events.
- Use a WMI client.
- Open Data Alert Manager.

Answer:



Explanation:

Box 1:

Capture the **Alerting: events processed - DeliverAlert** values and compare them to the previously queried events.

Box 2:

Open Data Alert Manager.

Box 3:

Query the **Alerting: events processed - GenerateAlert** event.

Note:

* DeliverAlert

The runtime creates the data alert message and sends it to all recipients by email.

* GenerateAlert

The alerting runtime processes the report data feed, applies the rules specified in the data alert definition, determines whether to create an instance of the data alert, and if needed creates an instance of the data alert.

* Data Alert Manager lists alert definitions and error information that help information workers and alerting administrators understand why the failure occurred. Some common reasons for failure include:

The report data feed changed and columns that are used in the data alert definition rules are no longer included in the data feed.

Permission to view the report was revoked.

The data type in the underlying data source changed and the alert definition is no longer valid.

Question No : 10 DRAG DROP - (Topic 7)

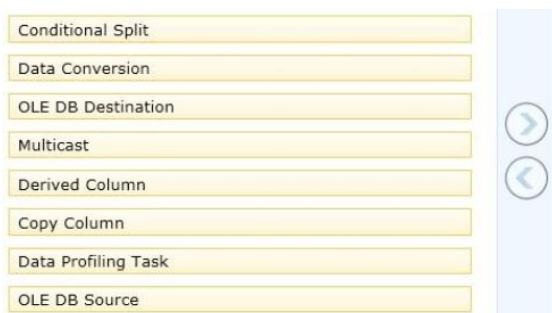
You are creating a SQL Server Integration Services (SSIS) package to populate a fact table from a source table. The fact table and source table are located in a SQL Azure database. The source table has a price field and a tax field. The OLE DB source uses the data access mode of Table.

You have the following requirements:

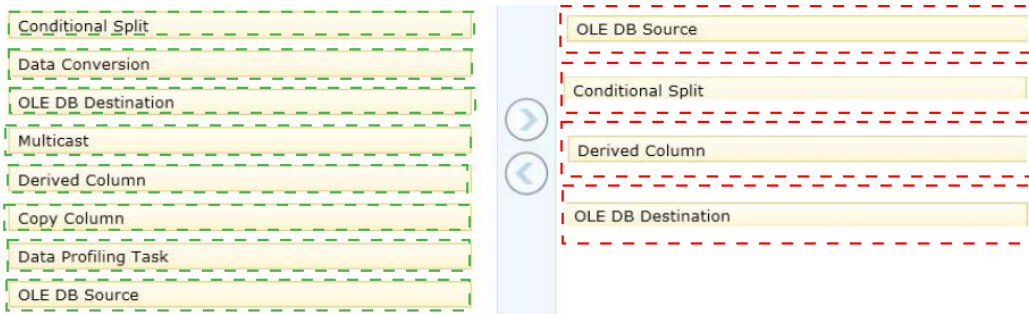
- ✎ The fact table must populate a column named TotalCost that computes the sum of the price and tax columns.
- ✎ Before the sum is calculated, any records that have a price of zero must be discarded.

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.)



Answer:



Explanation:

Box 1: Ole DB Source

Box 2: Conditional Split

Disregard lines with a 0 sum.

Box 3: Derived column

Box 4: Ole DB Destination

Note:

* 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.