

100% Money Back Guarantee

Vendor: Microsoft

Exam Code: 70-452

Exam Name: PRO:MS SQL Server 2008, Designing a Business Intelligence

Version: Demo

Question No : 1

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You plan to create a SQL Server 2008 Reporting Services (SSRS) solution that contains five sales dashboard reports.

Users must be able to manipulate the reports' parameters to analyze data.

You need to ensure that the following requirements are met:

Which two tasks should you perform?

(Each correct answer presents part of the solution. Choose two.)

- A. Filter data by using expressions.
- B. Specify the default values for each parameter.
- C. Create an available values list for each parameter.
- D. Create report parameters by using query parameters to filter data at the data source.

Answer: A,B

Explanation:

Question No : 2

You design a SQL Server 2008 Reporting Services (SSRS) solution. You create a report by using Microsoft Visual Studio .NET 2008.

The report contains the following components:

You need to ensure that a summary of sales transactions is displayed for each customer after the customer details.

Which component should you add to the report?

- A. List
- B. Table

- C. Matrix
- D. Subreport

Answer: D

Explanation:

[http://msdn.microsoft.com/en-us/library/ms160348\(SQL.100\).aspx](http://msdn.microsoft.com/en-us/library/ms160348(SQL.100).aspx)

How to: Add a Subreport and Parameters (Reporting Services)

Add subreports to a report when you want to create a main report that is a container for multiple related reports. A subreport is a reference to another report. To relate the reports through data values (for example, to have multiple reports show data for the same customer), you must design a parameterized report (for example, a report that shows the details for a specific customer) as the subreport. When you add a subreport to the main report, you can specify parameters to pass to the subreport.

You can also add subreports to dynamic rows or columns in a table or matrix. When the main report is processed, the subreport is processed for each row. In this case, consider whether you can achieve the desired effect by using data regions or nested data regions.

Question No : 3

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The solution includes a SQL Server 2008 Analysis Services (SSAS) database. The database contains a data mining structure that uses a SQL Server 2008 table as a data source. A table named OrderDetails contains detailed information on product sales. The OrderDetails table includes a column named Markup.

You build a data mining model by using the Microsoft Decision Trees algorithm. You classify Markup as discretized content.

The algorithm produces a large number of branches for Markup and results in low confidence ratings on predictable columns.

You need to verify whether the Markup values include inaccurate data.

What should you do?

- A. Modify the content type of Markup as Continuous.
- B. Create a data mining dimension in the SSAS database from OrderDetails.
- C. Create a data profile by using SQL Server 2008 Integration Services (SSIS).
- D. Create a cube in SSAS. Use OrderDetails as a measure group. Recreate the data mining structure and mining model from the cube data.

Answer: C

Explanation:

Discretized The column has continuous values that are grouped into buckets. Each bucket is considered to have a specific order and to contain discrete values. Possible values for discretization method are automatic, equal areas, or clusters. Automatic means that SSAS determines which method to use. Equal areas results in the input data being divided into partitions of equal size. This method works best with data with regularly distributed values. Clusters means that SSAS samples the data to produce a result that accounts for “clumps” of data values. Because of this sampling, Clusters can be used only with numeric input columns. You can use the date, double, long, or text data type with the Discretized content type.

Microsoft Decision Trees Algorithm

Microsoft Decision Trees is probably the most commonly used algorithm, in part because of its flexibility— decision trees work with both discrete and continuous attributes—and also because of the richness of its included viewers. It’s quite easy to understand the output via these viewers. This algorithm is used to both view and to predict. It is also used (usually in conjunction with the Microsoft Clustering algorithm) to find deviant values. The Microsoft Decision Trees algorithm processes input data by splitting it into recursive (related) subsets. In the default viewer, the output is shown as a recursive tree structure. If you are using discrete data, the algorithm identifies the particular inputs that are most closely correlated with particular predictable values, producing a result that shows which columns are most strongly predictive of a selected attribute. If you are using continuous data, the algorithm uses standard linear regression to determine where the splits in the decision tree occur. Clicking a node displays detailed information in the Mining Legend window. You can configure the view using the various drop-down lists at the top of the viewer, such as Tree, Default Expansion, and so on. Finally, if you’ve enabled drillthrough on your model, you can display the drillthrough information—either columns from the model or (new to SQL Server 2008) columns from the mining structure, whether or not they are included in this model.

Data Profiling

The control flow Data Profiling task relates to business problems that are particularly prominent in BI projects: how to deal with huge quantities of data and what to do when this

data originates from disparate sources.

Understanding source data quality in BI projects—when scoping, early in prototyping, and during package development—is critical when estimating the work involved in building the ETL processes to populate the OLAP cubes and data mining structures. It's common to underestimate the amount of work involved in cleaning the source data before it is loaded into the SSAS destination structures. The Data Profiling task helps you to understand the scope of the source-data cleanup involved in your projects. Specifically, this cleanup involves deciding which methods to use to clean up your data. Methods can include the use of advanced package transformations (such as fuzzy logic) or more staging areas (relational tables) so that fewer in-memory transformations are necessary during the transformation processes. Other considerations include total number of tasks in a single package, or overall package size.

Question No : 4

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The solution contains a SQL Server 2008 Analysis Services (SSAS) database. A measure group in the database contains log entries of manufacturing events. These events include accidents, machine failures, production capacity metrics, and other activities.

You need to implement a data mining model that meets the following requirements:

Which algorithm should the data mining model use?

- A. the Microsoft Time Series algorithm
- B. the Microsoft Decision Trees algorithm
- C. the Microsoft Linear Regression algorithm
- D. the Microsoft Logistic Regression algorithm

Answer: A

Explanation:

Microsoft Time Series Algorithm

Microsoft Time Series is used to impact a common business problem, accurate forecasting. This algorithm is often used to predict future values, such as rates of sale for a particular product. Most often the inputs are continuous values. To use this algorithm, your source data must contain at one column marked as Key Time. Any predictable columns must be of

type Continuous. You can select one or more inputs as predictable columns when using this algorithm.

Time series source data can also contain an optional Key Sequence column.

Function The ARTxp algorithm has proved to be very good at short-term prediction. The ARIMA algorithm is much better at longer-term prediction. By default, the Microsoft Time Series algorithm blends the results of the two algorithms to produce the best prediction for both the short and long term.

Microsoft Decision Trees Algorithm

Microsoft Decision Trees is probably the most commonly used algorithm, in part because of its flexibility— decision trees work with both discrete and continuous attributes—and also because of the richness of its included viewers. It's quite easy to understand the output via these viewers. This algorithm is used to both view and to predict. It is also used (usually in conjunction with the Microsoft Clustering algorithm) to find deviant values. The Microsoft Decision Trees algorithm processes input data by splitting it into recursive (related) subsets. In the default viewer, the output is shown as a recursive tree structure. If you are using discrete data, the algorithm identifies the particular inputs that are most closely correlated with particular predictable values, producing a result that shows which columns are most strongly predictive of a selected attribute. If you are using continuous data, the algorithm uses standard linear regression to determine where the splits in the decision tree occur. Clicking a node displays detailed information in the Mining Legend window. You can configure the view using the various drop-down lists at the top of the viewer, such as Tree, Default Expansion, and so on. Finally, if you've enabled drillthrough on your model, you can display the drillthrough information—either columns from the model or (new to SQL Server 2008) columns from the mining structure, whether or not they are included in this model.

Microsoft Linear Regression Algorithm

Microsoft Linear Regression is a variation of the Microsoft Decision Trees algorithm, and works like classic linear regression—it fits the best possible straight line through a series of points (the sources being at least two columns of continuous data). This algorithm calculates all possible relationships between the attribute values and produces more complete results than other (non-data mining) methods of applying linear regression. In addition to a key column, you can use only columns of the continuous numeric data type. Another way to understand this is that it disables splits. You use this algorithm to be able to visualize the relationship between two continuous attributes. For example, in a retail scenario, you might want to create physical placement locations in a retail store and rate of sale for items. The algorithm result is similar to that produced by any other linear regression method in that it produces a trend line. Unlike most other methods of calculating linear regression, the Microsoft Linear Regression algorithm in SSAS calculates all possible

relationships between all input dataset values to produce its results. This differs from other methods of calculating linear regression, which generally use progressive splitting techniques between the source inputs

Question No : 5

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The solution includes a SQL Server 2008 Analysis Services (SSAS) database. A cube in the database contains a large dimension named Customers. The database uses a data source that is located on a remote server.

Each day, an application adds millions of fact rows and thousands of new customers. Currently, a full process of the cube takes several hours.

You need to ensure that queries return the most recent customer data with the minimum amount of latency.

Which cube storage model should you use?

- A. hybrid online analytical processing (HOLAP)
- B. relational online analytical processing (ROLAP)
- C. multidimensional online analytical processing (MOLAP)
- D. automatic multidimensional online analytical processing (automatic MOLAP)

Answer: A

Explanation:

Relational OLAP

Relational OLAP (ROLAP) stores the cube structure in a multidimensional database. The leaf-level measures are left in the relational data mart that serves as the source of the cube. The preprocessed aggregates are also stored in a relational database table. When a decision maker requests the value of a measure for a certain set of dimension members, the ROLAP system first checks to determine whether the dimension members specify an aggregate or a leaf-level value. If an aggregate is specified, the value is selected from the relational table. If a leaf-level value is specified, the value is selected from the data mart. Also, because the ROLAP architecture retrieves leaf-level values directly from the data mart, the leaf-level values returned by the ROLAP system are always as up-to-date as the

data mart itself. In other words, the ROLAP system does not add latency to leaf-level data. The disadvantage of a ROLAP system is that the retrieval of the aggregate and leaf-level values is slower than the other OLAP architectures.

Multidimensional OLAP

Multidimensional OLAP (MOLAP) also stores the cube structure in a multidimensional database. However, both the preprocessed aggregate values and a copy of the leaf-level values are placed in the multidimensional database as well. Because of this, all data requests are answered from the multidimensional database, making MOLAP systems extremely responsive. Additional time is required when loading a MOLAP system because all the leaf-level data is copied into the multidimensional database. Because of this, times occur when the leaf-level data returned by the MOLAP system is not in sync with the leaf-level data in the data mart itself. A MOLAP system, therefore, does add latency to the leaf-level data. The MOLAP architecture also requires more disk space to store the copy of the leaf-level values in the multidimensional database. However, because MOLAP is extremely efficient at storing values, the additional space required is usually not significant.

Hybrid OLAP

Hybrid OLAP (HOLAP) combines ROLAP and MOLAP storage. This is why we end up with the word “hybrid” in the name. HOLAP tries to take advantage of the strengths of each of the other two architectures while minimizing their weaknesses. HOLAP stores the cube structure and the preprocessed aggregates in a multidimensional database. This provides the fast retrieval of aggregates present in MOLAP structures. HOLAP leaves the leaf-level data in the relational data mart that serves as the source of the cube. This leads to longer retrieval times when accessing the leaf-level values. However, HOLAP does not need to take time to copy the leaf-level data from the data mart. As soon as the data is updated in the data mart, it is available to the decision maker. Therefore, HOLAP does not add latency to the leaf-level data. In essence, HOLAP sacrifices retrieval speed on leaf-level data to prevent adding latency to leaf-level data and to speed the data load.

Question No : 6

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The solution includes a SQL Server 2008 Analysis Services (SSAS) database. The database contains a cube named Financials. The cube contains objects as shown in the exhibit.

Measure Groups	
Dimensions	Product Costs Sale Details
Cost Center (Dim Cost Cen...)	Cost Center Key Cost Center Key
Manager (Dim Manager)	Dim Cost Center - Cost Center Key Cost Center Key

A calculated member named Gross Margin references both Sales Details and Product Costs.

You need to ensure that the solution meets the following requirements:

What should you do?

- A. Add dimension-level security and enable the Visual Totals option.
- B. Add cell-level security that has read permissions on the Gross Margin measure
- C. Add cell-level security that has read contingent permissions on the Gross Margin measure.
- D. Change the permissions on the Managers dimension level from Read to Read/Write.

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms174927.aspx>

User Access Security Architecture

Microsoft SQL Server Analysis Services relies on Microsoft Windows to authenticate users. By default, only authenticated users who have rights within Analysis Services can establish a connection to Analysis Services. After a user connects to Analysis Services, the permissions that user has within Analysis Services are determined by the rights that are assigned to the Analysis Services roles to which that user belongs, either directly or through membership in a Windows role.

Dimension-Level Security

A database role can specify whether its members have permission to view or update dimension members in specified database dimensions. Moreover, within each dimension to which a database role has been granted rights, the role can be granted permission to view or update specific dimension members only instead of all dimension members. If a database role is not granted permissions to view or update a particular dimension and some or all the dimension's members, members of the database role have no permission to view the dimension or any of its members.

Note Dimension permissions that are granted to a database role apply to the cube dimensions based on the database dimension, unless different permissions are explicitly

granted within the cube that uses the database dimension.

Cube-Level Security

A database role can specify whether its members have read or read/write permission to one or more cubes in a database. If a database role is not granted permissions to read or read/write at least one cube, members of the database role have no permission to view any cubes in the database, despite any rights those members may have through the role to view dimension members.

Cell-Level Security

A database role can specify whether its members have read, read contingent, or read/write permissions on some or all cells within a cube. If a database role is not granted permissions on cells within a cube, members of the database role have no permission to view any cube data. If a database role is denied permission to view certain dimensions based on dimension security, cell level security cannot expand the rights of the database role members to include cell members from that dimension. On the other hand, if a database role is granted permission to view members of a dimension, cell-level security can be used to limit the cell members from the dimension that the database role members can view.

Question No : 7

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The solution includes a SQL Server 2008 Reporting Services (SSRS) infrastructure in a scale-out deployment. All reports use a SQL Server 2008 relational database as the data source. You implement row-level security.

You need to ensure that all reports display only the expected data based on the user who is viewing the report.

What should you do?

- A. Store the credential of a user in the data source.
- B. Configure the infrastructure to support Kerberos authentication.
- C. Configure the infrastructure to support anonymous authentication by using a custom

authentication extension.

D. Ensure that all report queries add a filter that uses the User.UserID value as a hiddenparameter.

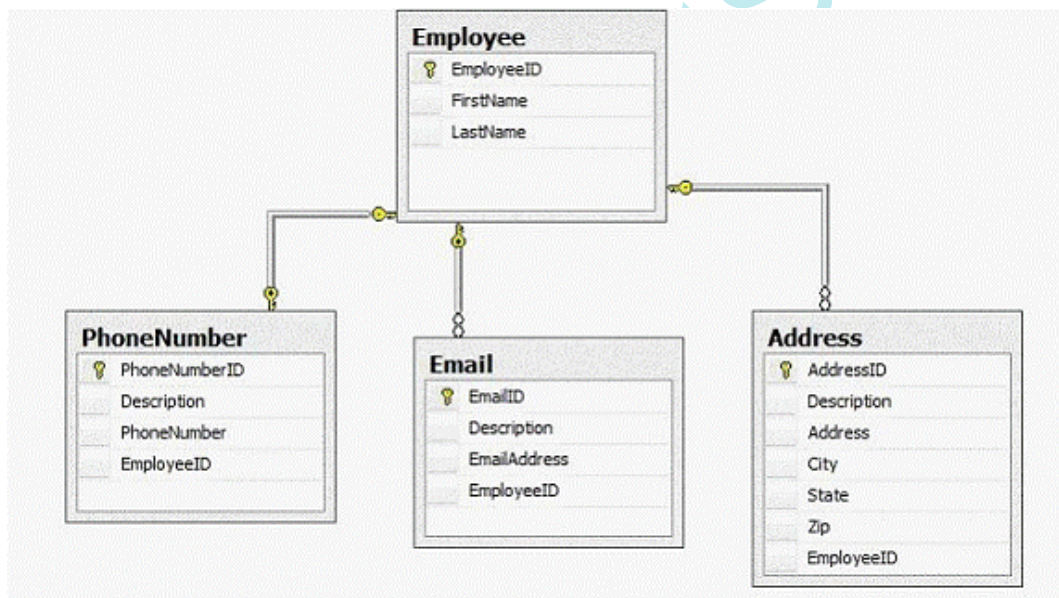
Answer: B

Explanation:

Question No : 8

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You need to load data into your online transaction processing (OLTP) database once a week by using data from a flat file. The file contains all the details about new employees who joined your company last week. The data must be loaded into the tables shown in the exhibit. (Click the Exhibit button.) Employee.EmployeeID is an identity.



A SQL Server 2008 Integration Services (SSIS) package contains one data flow for each of the destination tables. In the Employee Data Flow, an OLE DB Command transformation executes a stored procedure that loads the Employee record and returns the EmployeeID value.

You need to accomplish the following tasks:

What should you do?

- A. Use a Lookup Transformation in each of the child table data flows to find the EmployeeID based on first name and last name.
- B. Store the EmployeeID values in SSIS variables and use the variables to populate the FK columns in each of the child tables.
- C. After the Employee table is loaded, write the data to a Raw File Destination and use the raw file as a source for each of the subsequent Data Flows.
- D. After the Employee table is loaded, write the data to a Flat File Destination and use the flat file as a source for each of the subsequent Data Flows.

Answer: C

Explanation:

<http://technet.microsoft.com/en-us/library/ms141661.aspx>

Raw File Destination

The Raw File destination writes raw data to a file. Because the format of the data is native to the destination, the data requires no translation and little parsing. This means that the Raw File destination can write data more quickly than other destinations such as the Flat File and the OLE DB destinations. You can configure the Raw File destination in the following ways: Specify an access mode which is either the name of the file or a variable that contains the name of the file to which the Raw File destination writes.

Indicate whether the Raw File destination appends data to an existing file that has the same name or creates a new file.

The Raw File destination is frequently used to write intermediary results of partly processed data between package executions. Storing raw data means that the data can be read quickly by a Raw File source and then further transformed before it is loaded into its final destination. For example, a package might run several times, and each time write raw data to files. Later, a different package can use the Raw File source to read from each file, use a Union All transformation to merge the data into one data set, and then apply additional transformations that summarize the data before loading the data into its final destination such as a SQL Server table.

Raw File Source The Raw File source lets us utilize data that was previously written to a raw data file by a Raw File destination. The raw file format is the native format for Integration Services. Because of this, raw files can be written to disk and read from disk rapidly. One of the goals of Integration Services is to improve processing efficiency by moving data from the original source to the ultimate destination without making any stops in between. However, on some occasions, the data must be staged to disk as part of an Extract, Transform, and Load process. When this is necessary, the raw file format provides the most efficient means of accomplishing this task.

Question No : 9

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You create a SQL Server 2008 Integration Services (SSIS) package to perform an extract, transform, and load (ETL) process to load data to a DimCustomer dimension table that contains 1 million rows.

Your data flow uses the following components:

What should you do?

- A. Modify the UPDATE statement in the OLE DB Command transform to use the PAGLOCK table hint.
- B. Modify the UPDATE statement in the OLE DB Command transform to use the TABLOCK table hint.
- C. Stage the data in the data flow. Replace the OLE DB Command transform in the dataflow with an Execute SQL task in the control flow.
- D. Stage the data in the data flow. Replace the UPDATE statement in the OLE DB Command transform with a DELETE statement followed by an INSERT statement.

Answer: C

Explanation:

Data Flow

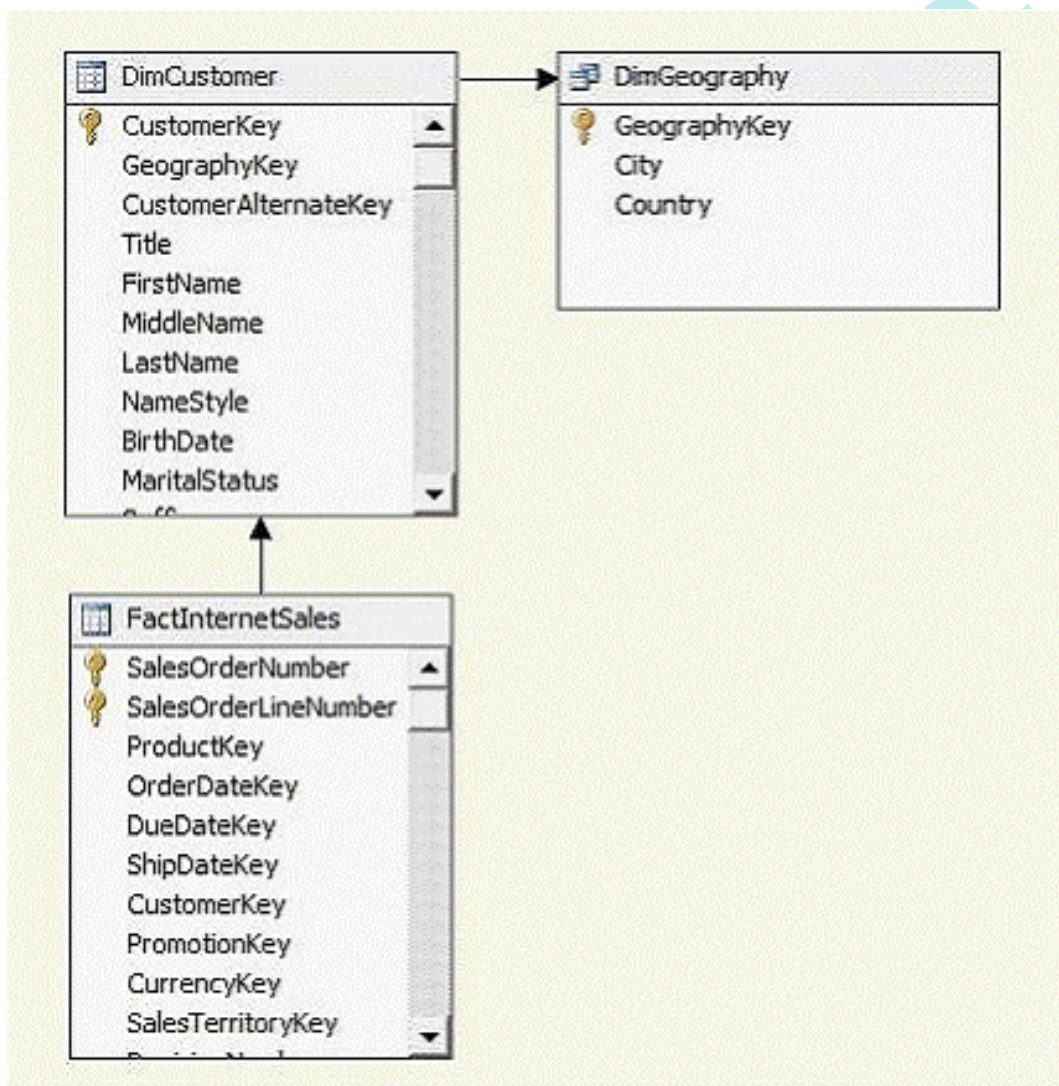
Once we set the precedence constraints for the control flow tasks in the package, we can define each of the data flows. This is done on the Data Flow Designer tab. Each data flow task that was added to the control flow has its own layout on the Data Flow Designer tab. We can switch between different data flows using the Data Flow Task drop-down list located at the top of the Data Flow tab. The Data Flow Toolbox contains three types of items: data flow sources, data flow transformations, and data flow destinations. However, on some occasions, the data must be staged to disk as part of an Extract, Transform, and Load process. When this is necessary, the Raw File format provides the most efficient means of accomplishing this task. Execute SQL Task The Execute SQL task enables us to execute SQL statements or stored procedures. The contents of variables can be used for input, output, or input/output parameters and the return value. We can also save the result set from the SQL statements or stored procedure in a package variable. This result set could be a single value, a multirow/multicolumn result set, or an XML document.

<http://www.sqlservercentral.com/blogs/dknight/archive/2008/12/29/ssis-avoid-ole-db-command.aspx> SSIS – Avoid OLE DB Command he OLE DB Command runs insert, update or delete statements for each row, while a Bulk Insert in this instance. That means every

single row that goes through your package would have an insert statement run when it gets to an OLE DB Command. So if you know you are dealing with more than just a couple hundred rows per run then I would highly suggest using a staging table vs. the OLE DB Command.

Question No : 10

You design a SQL Server 2008 Analysis Services (SSAS) solution. The data source view has tables as shown in the exhibit. (Click the Exhibit button.)



The FactInternetSales measure will be queried frequently based on the city and country of the customer.

You need to design a cube that will provide optimal performance for queries.

Which design should you choose?

A.

- Create two dimensions named Customer and Geography from the DimCustomer table and the DimGeography table, respectively.
- Create a materialized reference relationship between the Geography dimension and the FactInternetSales measure by using the Customer dimension as an intermediate dimension.

B.

- Create two dimensions named Customer and Geography from the DimCustomer table and the DimGeography table, respectively.
- Create an unmaterialized reference relationship between the Geography dimension and the FactInternetSales measure by using the Customer dimension as an intermediate dimension.

C.

- Create a dimension named Customer by joining the DimGeography and DimCustomer tables.
- Add an attribute relationship from CustomerKey to City and from City to Country. Create a regular relationship in the cube between the Customer dimension and the FactInternetSales measure.

D.

- Create a dimension named Customer by joining the DimGeography and DimCustomer tables.
- Add an attribute relationship from CustomerKey to City and from CustomerKey to Country.
- Create a regular relationship in the cube between the Customer dimension and the FactInternetSales measure.

Answer: C

Explanation:

Question No : 11

You design a Business Intelligence (BI) solution by using SQL Server 2008.

Employees use a Windows Forms application based on Microsoft .NET Framework 3.5. SQL Server is not installed on the employees' computers.

You write a report by using Report Definition Language (RDL).

You need to ensure that if the employees are disconnected from the corporate network, the application renders the report.

What should you do?

- A. Configure the application to use an SSRS Web service by using the Render method.
- B. Configure the application to use an SSRS Web service by using the RenderStream method.
- C. Embed ReportViewer in the application and configure ReportViewer to render reports by using the local processing mode.
- D. Embed ReportViewer in the application and configure ReportViewer to render reports by using the remote processing mode.

Answer: C

Explanation:

Embedding Custom ReportViewer Controls

Microsoft provides two controls in Visual Studio 2008 that allow you to embed SSRS reports (or link to an existing SSRS report hosted on an SSRS instance) in your custom Windows Forms or Web Forms applications. Alternatively, you can also design some types of reports from within Visual Studio and then host them in your custom applications. The two report processing modes that this control supports are remote processing mode and local processing mode.

Remote processing mode allows you to include a reference to a report that has already been deployed to a report server instance. In remote processing mode, the ReportViewer control encapsulates the URL access method we covered in the previous section. It uses the SSRS Web service to communicate with the report server. Referencing deployed reports is preferred for BI solutions because the overhead of rendering and processing the often large BI reports is handled by the SSRS server instance or instances. Also, you can choose to scale report hosting to multiple SSRS servers if scaling is needed for your solution. Another advantage to this mode is that all installed rendering and data extensions are available to be used by the referenced report. Local processing mode allows you to run a report from a computer that does not have SSRS installed on it.

Local reports are defined differently within Visual Studio itself, using a visual design interface that looks much like the one in BIDS for SSRS. The output file is in a slightly different format for these reports if they're created locally in Visual Studio. It's an *.rdlc file rather than an *.rdl file, which is created when using a Report Server Project template in BIDS. The *.rdlc file is defined as an embedded resource in the Visual Studio project. When displaying *.rdlc files to a user, data retrieval and processing is handled by the hosting application, and the report rendering (translating it to an output format such as

HTML or PDF) is handled by the ReportViewer control. No server-based instance of SSRS is involved, which makes it very useful when you need to deploy reports to users that are only occasionally connected to the network and thus wouldn't have regular access to the SSRS server. Only PDF, Excel, and image-rendering extensions are supported in local processing mode. If you use local processing mode with some relational data as your data source, a new report design area opens up. As mentioned, the metadata file generated has the *.rdlc extension. When working in local processing mode in Visual Studio 2008, you're limited to working with the old-style data containers—that is, table, matrix, or list.

The new combined-style Tablix container is not available in this report design mode in Visual Studio 2008. Both versions of this control include a smart tag that helps you to configure the associated required properties for each of the usage modes. Also, the ReportViewer control is freely redistributable, which is useful if you're considering using either version as part of a commercial application.

Question No : 12

You design a SQL Server 2008 Reporting Services (SSRS) solution. The solution contains a report. The report includes information that is grouped into hierarchical levels.

You need to ensure that the solution meets the following requirements:

Which feature should the report use?

- A. filter
- B. drilldown
- C. drillthrough
- D. a document map

Answer: B

Explanation: Explanation/Reference:

<http://technet.microsoft.com/en-us/library/dd207141.aspx>

Drillthrough, Drilldown, Subreports, and Nested Data Regions (Report Builder 3.0 and SSRS)

You can organize data in a variety of ways to show the relationship of the general to the detailed. You can put all the data in the report, but set it to be hidden until a user clicks to

reveal details; this is a drilldown action. You can display the data in a data region, such as a table or chart, which is nested inside another data region, such as a table or matrix. You can display the data in a subreport that is completely contained within a main report. Or, you can put the detail data in drillthrough reports, separate reports that are displayed when a user clicks a link.

Question No : 13

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You plan to develop SQL Server 2008 Reporting Services (SSRS) reports. Several reports will contain identical data regions.

You need to minimize the amount of maintenance when changes are made to the data regions.

What should you do?

- A. Grant the Create Linked Reports role to all users.
- B. Create each data region as a report. Embed the reports by using the subreport control.
- C. Create a report template for each data region. Use the report template to create each report.
- D. Create a shared data source in the SSRS project. Use the new shared data source for all reports.

Answer: B

Explanation:

Question No : 14

You are designing a SQL Server 2008 Reporting Services (SSRS) solution. You have a report that has several parameters that are populated when users execute the report.

You need to ensure that the solution meets the following requirements:

Which feature should you use?

- A. My Reports
- B. Linked Reports
- C. Standard Subscription
- D. Data-Driven Subscription

Answer: B

Explanation:

With a linked report, our report is deployed to one folder. It is then pointed to by links placed elsewhere within the Report Catalog. To the user, the links look just like a report. Because of these links, the report appears to be in many places. The sales department sees it in their folder. The personnel department sees it in their folder.

The fact of the matter is the report is only deployed to one location, so it is easy to administer and maintain. An execution snapshot is another way to create a cached report instance. Up to this point, we have discussed situations where cached report instances are created as the result of a user action. A user requests a report, and a copy of that report's intermediate format is placed in the report cache. With execution snapshots, a cached report instance is created automatically. Not all users can change execution snapshots. To change the execution snapshot properties for a report, you must have rights to the Manage Reports task. Of the four predefined security roles, the Content Manager, My Reports, and Publisher roles have rights to this task.

(McGraw-Hill - Delivering Business Intelligence with Microsoft SQL Server 2008 (2009))

<http://msdn.microsoft.com/en-us/library/bb630404.aspx>

A linked report is a report server item that provides an access point to an existing report. Conceptually, it is similar to a program shortcut that you use to run a program or open a file. A linked report is derived from an existing report and retains the original's report definition. A linked report always inherits report layout and data source properties of the original report. All other properties and settings can be different from those of the original report, including security, parameters, location, subscriptions, and schedules.

You can create a linked report on the report server when you want to create additional versions of an existing report. For example, you could use a single regional sales report to create region-specific reports for all of your sales territories. Although linked reports are typically based on parameterized reports, a parameterized report is not required. You can create linked reports whenever you want to deploy an existing report with different settings.

Question No : 15

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You have developed SQL Server 2008 Reporting Services (SSRS) reports that are deployed on an SSRS instance.

You plan to develop a new application to view the reports. The application will be developed by using Microsoft ASP.NET 3.5.

You need to ensure that the application can perform the following tasks:

What should you do?

- A. Configure the ASP.NET application to use the SSRS Web service.
- B. Configure the ASP.NET application to use URL access along with the Command parameter.
- C. Embed a ReportViewer control in the ASP.NET application. Configure the control to use the local processing mode.
- D. Embed a ReportViewer control in the ASP.NET application. Configure the control to use the remote processing mode.

Answer: A

Explanation:

Report Server Web Service

The Report Server Web service is the core engine for all on-demand report and model processing requests that are initiated by a user or application in real time, including most requests that are directed to and from Report Manager. It includes more than 70 public methods for you to access SSRS functionality programmatically. The Report Manager Web site accesses these Web services to provide report rendering and other functionality. Also, other integrated applications, such as the Report Center in Office SharePoint Server 2007, call SSRS Web services to serve up deployed reports to authorized end users. The Report Server Web service performs end-to-end processing for reports that run on demand. To support interactive processing, the Web service authenticates the user and checks the authorization rules prior to handing a request. The Web service supports the default Windows security extension and custom authentication extensions. The Web service is also the primary programmatic interface for custom applications that integrate with Report Server, although its use is not required. If you plan to develop a custom interface for your reports, rather than using the provided Web site or some other integrated application (such as Office SharePoint Server 2007), you'll want to explore the SQL Server Books Online topic "Reporting Services Web Services Class Library." There you can examine specific Web methods.

Question No : 16

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You design a SQL Server 2008 Reporting Services (SSRS) report that meets the following requirements:

You need to design the report to minimize the impact on bandwidth.

What should you do?

- A.** Create a standard report that contains all sales orders. Implement report filtering based on the month.
- B.** Create a standard report that contains all sales orders. Implement grouping for the monthly summaries.
- C.** Create a standard report that contains the monthly summaries. Create a subreport for the sales orders for any given month.
- D.** Create a standard report that contains the monthly summaries. Create a drillthrough report for the sales orders for any given month.

Answer: D

Explanation:

Drillthrough Action Defines a dataset to be returned as a drillthrough to a more detailed level. Creating Drillthrough Actions For the most part, Drillthrough Actions have the same properties as Actions. Drillthrough Actions do not have Target Type or Target Object properties. In their place, the Drillthrough Action has the following:

- Drillthrough Columns Defines the objects to be included in the drillthrough dataset.
- Default A flag showing whether this is the default Drillthrough Action.
- Maximum Rows The maximum number of rows to be included in the drillthrough dataset.

<http://technet.microsoft.com/en-us/library/ff519554.aspx>

Drillthrough Reports(Report Builder 3.0 and SSRS)

A drillthrough report is a report that a user opens by clicking a link within another report. Drillthrough reports commonly contain details about an item that is contained in an original summary report. The data in the drillthrough report is not retrieved until the user clicks the link in the main report that opens the drillthrough report. If the data for the main report and the drillthrough report must be retrieved at the same time, consider using a subreport

Question No : 17

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You create a sales report by using SQL Server 2008 Reporting Services (SSRS). The report is used by managers in a specific country.

Each manager prints multiple copies of the report that contains the previous day's sales for each of their sales executives.

You need to ensure that the report uses the minimum number of round trips to the database server.

What should you do?

- A. Query the database for both Country and Sales Executive.
- B. Implement report filtering for both Country and Sales Executive.
- C. Implement report filtering for Country and query the data source for Sales Executive.
- D. Implement report filtering for Sales Executive and query the data source for Country.

Answer: D

Explanation:

<http://technet.microsoft.com/en-us/library/dd239395.aspx>

Choosing When to Set a Filter Specify filters for report items when you cannot filter data at the source. For example, use report filters when the data source does not support query parameters, or you must run stored procedures and cannot modify the query, or a parameterized report snapshot displays customized data for different users.

You can filter report data before or after it is retrieved for a report dataset. To filter data before it is retrieved, change the query for each dataset. When you filter data in the query, you filter data at the data source, which reduces the amount data that must be retrieved and processed in a report. To filter data after it is retrieved, create filter expressions in the report. You can set filter expressions for a dataset, a data region, or a group, including detail groups. You can also include parameters in filter expressions, providing a way to filter data for specific values or for specific users, for example, filtering on a value that

identifies the user viewing the report.

Question No : 18

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You plan to create a SQL Server 2008 Reporting Services (SSRS) report. The report must display the list of orders placed through the Internet.

You need to ensure that the following requirements are met:

Which type of report should you create?

- A. Linked
- B. Ad Hoc
- C. Cached
- D. Snapshot

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/bb630404.aspx#Snapshot>

A report snapshot is a report that contains layout information and query results that were retrieved at a specific point in time. Unlike on-demand reports, which get up-to-date query results when you select the report, report snapshots are processed on a schedule and then saved to a report server. When you select a report snapshot for viewing, the report server retrieves the stored report from the report server database and shows the data and layout that were current for the report at the time the snapshot was created.

Report snapshots are not saved in a particular rendering format. Instead, report snapshots are rendered in a final viewing format (such as HTML) only when a user or an application requests it. Deferred rendering makes a snapshot portable. The report can be rendered in the correct format for the requesting device or Web browser. Report snapshots serve three purposes:

- Report history. By creating a series of report snapshots, you can build a history of a report that shows how data changes over time.
- Consistency. Use report snapshots when you want to provide consistent results for multiple users who must work with identical sets of data. With volatile data, an on-demand

report can produce different results from one minute to the next. A report snapshot, by contrast, allows you to make valid comparisons against other reports or analytical tools that contain data from the same point in time.

- Performance. By scheduling large reports to run during off-peak hours, you can reduce processing impact on the report server during core business hours.

Question No : 19

You are creating a SQL Server 2008 Reporting Services (SSRS) solution for a company that has offices in different countries. The company has a data server for each country.

Sales data for each country is persisted in the respective data server for the country. Report developers have only Read access to all data servers. All data servers have the same schema for the database.

You design an SSRS solution to view sales data.

You need to ensure that users are able to easily switch between sales data for different countries.

What should you do?

- A. Implement a single shared data source.
- B. Implement multiple shared data sources.
- C. Implement an embedded data source that has a static connection string.
- D. Implement an embedded data source that has an expression-based connection string.

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms156450.aspx>

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. For example, suppose a multinational firm has data servers in several countries. With an expression-based connection string, a user who is running a sales report can select a data source for a particular country before running the report. Design the report using a static connection string. A static connection string refers to a connection string that is not set through an

expression (for example, when you follow the steps for creating a report-specific or shared data source, you are defining a static connection string). Using a static connection string allows you to connect to the data source in Report Designer so that you can get the query results you need to create the report.

When defining the data source connection, do not use a shared data source. You cannot use a data source expression in a shared data source. You must define an embedded data source for the report. Specify credentials separately from the connection string. You can use stored credentials, prompted credentials, or integrated security.

Add a report parameter to specify a data source. For parameter values, you can either provide a static list of available values (in this case, the available values should be data sources you can use with the report) or define a query that retrieves a list of data sources at run time.

Be sure that the list of data sources shares the same database schema. All report design begins with schema information. If there is a mismatch between the schema used to define the report and the actual schema used by the report at run time, the report might not run. Before publishing the report, replace the static connection string with an expression. Wait until you are finished designing the report before you replace the static connection string with an expression. Once you use an expression, you cannot execute the query in Report Designer. Furthermore, the field list in the Report Data pane and the Parameters list will not update automatically.

Question No : 20

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The solution will contain a total of 100 different reports created by using Report Definition Language (RDL).

Each report must meet the following requirements:

The business rules for all reports that determine the calculations change frequently.

You need to design a solution that meets the requirements. You need to perform this action by using the minimum amount of development and maintenance effort.

What should you do?

- A. Create hidden parameters in each report.
- B. Create internal parameters in each report.
- C. Implement the function in the <Code> element of each report.
- D. Implement the function in a custom assembly. Reference the assembly in each report.

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms159238.aspx>

Including References to Code from Custom Assemblies

To use custom assemblies in a report, you must first create the assembly, make it available to Report Designer, add a reference to the assembly in the report, and then use an expression in the report to refer to the methods contained in that assembly. When the report is deployed to the report server, you must also deploy the custom assembly to the report server.

To refer to custom code in an expression, you must call the member of a class within the assembly. How you do this depends on whether the method is static or instance-based. Static methods within a custom assembly are available globally within the report. You can access static methods in expressions by specifying the namespace, class, and method name. The following example calls the method ToGBP, which converts the value of the StandardCost value from dollar to pounds sterling:

Copy

```
=CurrencyConversion.DollarCurrencyConversion.ToGBP(Fields!StandardCost.Value)
```

Instance-based methods are available through a globally defined Code member. You access these by referring to the Code member, followed by the instance and method name. The following example calls the instance method ToEUR, which converts the value of StandardCost from dollar to euro: Copy

```
=Code.m_myDollarConversion.ToEUR(Fields!StandardCost.Value)
```

Note

In Report Designer, a custom assembly is loaded once and is not unloaded until you close Visual Studio. If you preview a report, make changes to a custom assembly used in the report, and then preview the report again, the changes will not appear in the second preview. To reload the assembly, close and reopen Visual Studio and then preview the report.

Question No : 21

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You create a SQL Server 2008 Reporting Services (SSRS) solution. The solution contains a report named Sales Details that displays sales information of all the employees.

You create an SSRS report named Sales Summary that displays the total monthly sales of each employee.

Users who view the Sales Summary report occasionally require the monthly sales details for a particular employee.

You need to ensure that the users can click a value in the month column of the Sales Summary report to open and render the Sales Details report.

What should you do?

- A. Use a subreport.
- B. Use a bookmark link.
- C. Use the drilldown functionality.
- D. Use a drillthrough report link.

Answer: D

Explanation:

<http://technet.microsoft.com/en-us/library/ff519554.aspx>

Drillthrough Reports (Report Builder 3.0 and SSRS)

A drillthrough report is a report that a user opens by clicking a link within another report. Drillthrough reports commonly contain details about an item that is contained in an original summary report. The data in the drillthrough report is not retrieved until the user clicks the link in the main report that opens the drillthrough report. If the data for the main report and the drillthrough report must be retrieved at the same time, consider using a subreport

Question No : 22

You design a Business Intelligence (BI) solution by using SQL Server 2008. You create a SQL Server 2008 Reporting Services (SSRS) report.

The report contains summary information in two sections named Agencies and States.

The Agency summary section contains two matrices and the State summary section contains a table. The information about each section is grouped together.

You need to design the report to meet the following requirements:

What should you do?

- A. Select the Keep together on one page option on all report items.
- B. Select a line component between the report items for the Agency and State summary sections.
- C. Select all the report items for each section in a list report item and enable the Add A Page Break Before option on the list report item.
- D. Select all the report items for each section in a rectangle report item and enable the Add A Page Break Before option on the rectangle report item.

Answer: D

Explanation:

[http://technet.microsoft.com/en-us/library/ms155915\(SQL.100\).aspx](http://technet.microsoft.com/en-us/library/ms155915(SQL.100).aspx)

Add a rectangle to your report when you want a graphical element to separate areas of the report, emphasize areas of a report, or provide a background for one or more report items. Rectangles are also used as containers to help control the way data regions render in a report. You can customize the appearance of a rectangle by editing rectangle properties such as the background and border colors.

Question No : 23

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You plan to create a SQL Server 2008 Reporting Services (SSRS) solution. Developers generate random reports against a data source that contains 200 tables. Power users generate random reports against four of the 200 tables.

You need to design a strategy for the SSRS solution to meet the following requirements:

Which strategy should you use?

A.

- Create two Report Builder models.
- Include the four frequently used tables in the first model and all the tables in the second model.

B.

- Create a Report Builder model by using all the tables.
- Create a perspective within the model to use only the four frequently used tables.

C.

- Create a Report Builder model by using all the tables.
- Create two folders.
- Place the four frequently used tables in the first folder and the remaining tables in the second folder.

D.

- Create two Data Source Views.
- Include all the tables in one Data Source View and the four frequently used tables in the other Data Source View.
- Create two Report Builder models so that each model uses one of the Data SourceViews.

Answer: B

Explanation:

Creating a Report Model

Like reports, Report Models are created in the Business Intelligence Development Studio and then deployed to a report server. Unlike reports, Report Models can have security rights assigned to different pieces of their structure to provide the fine-grained security that is often required in ad hoc reporting situations. We use the Report Model Wizard to create the Report Model, and then do some manual tweaking to make it more usable. We then deploy the Report Model to the report server.

Finally, we set security within the model itself.

A perspective is a subset of the information in the model. Usually, a perspective coincides with a particular job or work area within an organization. If a plus sign is to the left of the model, the model contains one or more perspectives. Click the plus sign to view the perspectives. If you select one of these perspectives as the data source for your report, only the entities in that perspective will be available to your report. Because perspectives reduce the number of entities you have to look through to find the data you need on your report, it is usually a good idea to choose a perspective, rather than using the entire Report Model.

<http://msdn.microsoft.com/en-us/library/ms345316.aspx>

For models that contain many subject areas, for example, Sales, Manufacturing, and Supply data, it might be helpful to Report Builder users if you create perspectives of the model.

A perspective is a sub-set of a model. Creating perspectives can make navigating through the contents of the model easier for your model users.

Question No : 24

You design a Business Intelligence (BI) solution by using SQL Server 2008. Users have no specialized knowledge of the Transact-SQL (T-SQL) language.

You plan to create a reporting solution by using SQL Server 2008 Reporting Services (SSRS). You have data stored in a SQL Server 2008 relational data warehouse.

You need to ensure that users are able to create the necessary reports by using data from the data warehouse.

What should you do?

A.

- Create a shared data source that points to the data warehouse.
- Instruct the users to use Report Designer in Business Intelligence Development Studio (BIDS) to create the reports by using the shared data source.

B.

- Create a Report Model from the data warehouse.
- Instruct the users to use Report Builder 2.0 to create the reports by using the Report Model.

C.

- Create a shared data source that point to the data warehouse.
- Instruct the users to use Report Builder 2.0 to create the reports by using the shared data source.

D.

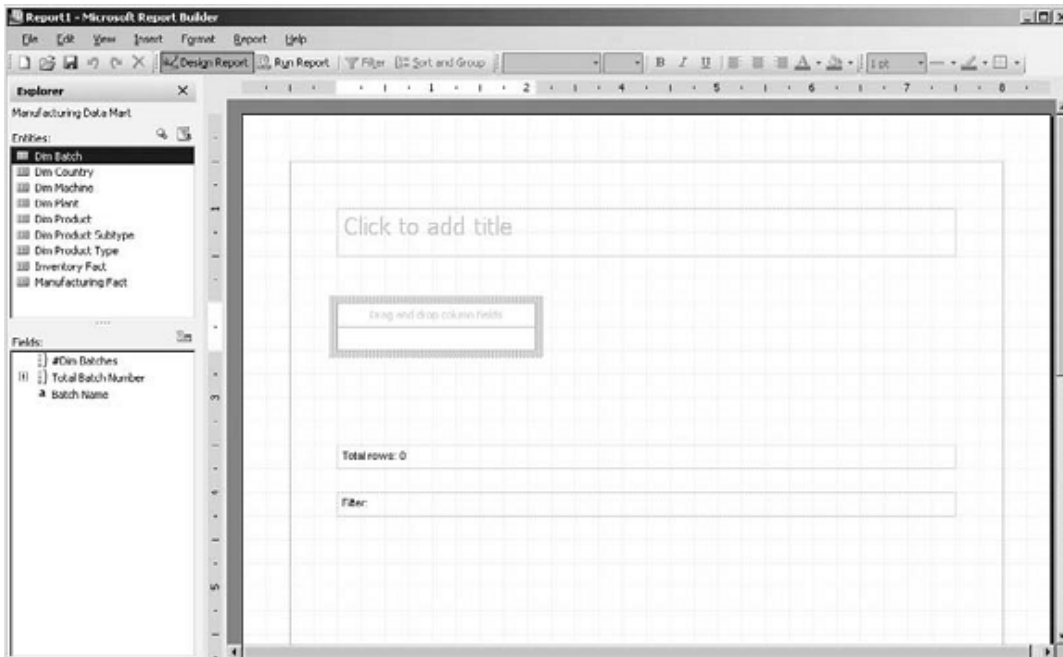
- Create a Report Model from the data warehouse.
- Instruct the users to use Report Designer in Business Intelligence Development Studio (BIDS) to create the reports by using the Report Model.

Answer: B

Explanation:

If you are creating a new report, select the Report Model or the perspective that should serve as the data source, along with the report layout, and click OK. If you want to edit an existing report, click the Open button on the toolbar. You can then navigate the report server folder structure to find the Report Builder report you want to edit. You cannot use the Report Builder to edit reports that were created or edited using the Report Designer in Visual Studio 2008 or the Business Intelligence Development Studio. Once a Report Model

has been selected for your data source, or an existing report has been chosen for editing, the main Report Builder opens. When creating a new report, the Report Builder appears similar to Figure.



[http://msdn.microsoft.com/en-us/library/ms159750\(SQL.100\).aspx](http://msdn.microsoft.com/en-us/library/ms159750(SQL.100).aspx)

Report Builder allows business users to create their own reports based on a user-friendly report model created in Model Designer. Fully integrated with Microsoft SQL Server Reporting Services, Report Builder leverages the full reporting platform to bring ad hoc reporting to all users.

Report Builder Users create reports with the Report Builder tool. The Report Builder interface is built on top of familiar Microsoft Office paradigms such as Excel and PowerPoint. Users start with report layout templates containing pre-defined data regions to build combinations of tables, matrices and charts. They navigate the reporting model to select report items and set constraints to filter the report data. The reporting model contains all of the necessary information for the Report Builder to automatically generate the source query and retrieve the requested data.

The Report Builder also allows users to:

Question No : 25

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You create a SQL Server 2008 Reporting Services (SSRS) solution. Twenty users edit the content of the reports. The users belong to an Active Directory group named Developers.

At times, the reports are published with incomplete information.

You need to design an approval process.

What should you do?

- A. Restrict the Developers group to the Browser role in SSRS.
- B. Add the Developers group to the Content Manager role in SSRS.
- C. Deploy the reports to a Microsoft Office SharePoint Server 2007 environment.
- D. Create a shared schedule for the reports. Set the snapshot execution option for all the reports by using the shared schedule.

Answer: C

Explanation:

Question No : 26

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You create a SQL Server 2008 Reporting Services (SSRS) solution. The solution has a report named SalesDetails that contains a parameter named EmployeeID.

You have the following constraints:

You need to ensure that the constraints are met before you deliver the report to the employees.

What should you do?

- A. Create a data-driven subscription.
- B. Create a SharePoint Report Center site.
- C. Create a subscription for each employee.
- D. Create a report model for each employee.

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms159150.aspx>

A data-driven subscription provides a way to use dynamic subscription data that is retrieved from an external data source at run time. A data-driven subscription can also use static text and default values that you specify when the subscription is defined.

You can use data-driven subscriptions to do the following:

- Distribute a report to a fluctuating list of subscribers. For example, you can use data-driven subscriptions to distribute a report throughout a large organization where subscribers vary from one month to the next, or use other criteria that determines group membership from an existing set of users.
- Filter the report output using report parameter values that are retrieved at run time.
- Vary report output formats and delivery options for each report delivery.

A data-driven subscription is composed of multiple parts. The fixed aspects of a data-driven subscription are defined when you create the subscription, and these include the following:

- The report for which the subscription is defined (a subscription is always associated with a single report).
- The delivery extension used to distribute the report. You can specify report server e-mail delivery, file share delivery, the null delivery provider used for preloading the cache, or a custom delivery extension. You cannot specify multiple delivery extensions within a single subscription.
- The subscriber data source. You must specify a connection string to the data source that contains subscriber data when you define the subscription. The subscriber data source cannot be specified dynamically at run time.
- The query that you use to select subscriber data must be specified when you define the subscription. You cannot change the query at run time.

Dynamic values used in a data-driven subscription are obtained when the subscription is processed. Examples of variable data that you might use in a subscription include the subscriber name, e-mail address, preferred report output format, or any value that is valid for a report parameter. To use dynamic values in a data-driven subscription, you define a mapping between the fields that are returned in the query to specific delivery options and to report parameters. Variable data is retrieved from a subscriber data source each processed.

Question No : 27

You are a SQL Server 2008 Analysis Services (SSAS) data mining architect.

The customer table contains the following column names:

You plan to build a data mining model by using Microsoft Decision Trees algorithm for the customer table.

You need to identify the data column-model parameter pairs to predict possible buyers.

Which model should you select?

www.Pass4itSure.com

A.

Data Column	Model Parameter Type
Customer_key	Input
Name	Ignore
Education Level	Input
Age	Input
IsBuyer	Key

B.

Data Column	Model Parameter Type
Customer_key	Input
Name	Ignore
Education Level	Input, Predict
Age	Input, Predict
IsBuyer	Key

C.

Data Column	Model Parameter Type
Customer_key	Predict
Name	Key
Education Level	Input
Age	Input
IsBuyer	Input

D.

Data Column	Model Parameter Type
Customer_key	Key
Name	Ignore
Education Level	Input
Age	Input
IsBuyer	Predict

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

Answer: D

Explanation:

<http://technet.microsoft.com/en-us/library/ms175423.aspx>

Changing Mining Column Usage

You can change which columns are included in a mining model and how each column is used, such as input, key, or predictable, by using the cells for that model column in the grid on the Mining Models tab. Each cell corresponds to a column in the mining structure. For key columns, you can set the cell to Key or Ignore. For input and output columns, you can set the cell to the following values:

- Ignore
- Input
- Predict
- PredictOnly

If you set a cell to Ignore, the column is removed from the mining model, but that column can still be used by other mining models in the structure.

Question No : 28

You design a SQL Server 2008 Analysis Services (SSAS) solution.

You have the following requirements for a single data mining model:

You need to design the model to meet the requirements.

What should you include in the design?

- A. Onecolumn of the Long data type and the Discrete content type
- B. Onecolumn of the Long data type and the Continuous content type
- C. Twocolumns, each of the Long data type and the Discrete content type
- D. Twocolumns, each of the Long data type and the Discretized content type

Answer: D

Explanation:

We do this by choosing Discretized as the Content Type, and then selecting a method that groups continuous values into a discrete number of buckets (that is, ages 11–15, ages

16–20, and so on). Often, it is much easier to do analysis on discretized values than on continuous values. For the purposes of a given analysis, the buying habits of 16–20-year-olds may be similar enough so that we can study them as a group in one discretized bucket. (McGraw-Hill - Delivering Business Intelligence with Microsoft SQL Server 2008 (2009)) Discretized The column has continuous values that are grouped into buckets. Each bucket is considered to have a specific order and to contain discrete values. You saw an example of this in Figure 12-2 using the Age column in the Targeted Mining sample. Note that you'll also set the DiscretizationMethod and (optionally) the DiscretizationBucketCount properties if you mark your column as Discretized. In our sample, we've set the bucket size to 10 and DiscretizationMethod to Automatic. Possible values for discretization method are automatic, equal areas, or clusters. Automatic means that SSAS determines which method to use. Equal areas results in the input data being divided into partitions of equal size. This method works best with data with regularly distributed values. Clusters means that SSAS samples the data to produce a result that accounts for "clumps" of data values. Because of this sampling, Clusters can be used only with numeric input columns. You can use the date, double, long, or text data type with the Discretized content type.

Question No : 29

You design a SQL Server 2008 Analysis Services (SSAS) solution.

The solution includes a mining structure that is created by using the default options and a mining model that uses the Microsoft Clustering algorithm.

You need to ensure that users can access source data by querying the mining model.

What should you do?

- A. Modify the mining structure to include a filter.
- B. Modify the mining structure to enable drillthrough.
- C. Include a task in the solution to process the mining model.
- D. Include a task in the solution to delete all cached data from the mining model.

Answer: B

Explanation:

Microsoft Clustering algorithm

Function The Microsoft Clustering algorithm builds clusters of entities as it processes the training data set. It looks at the values of each attribute for the entities in the cluster. By entering the attribute value we want, we can have the clusters color-coded according to the concentration of our desired value.

Tasks

The main purpose of the Microsoft Clustering algorithm is

- Segmentation

It can also be used for

- Regression
- Classification

(McGraw-Hill - Delivering Business Intelligence with Microsoft SQL Server 2008 (2009))

Microsoft Clustering Algorithm

As its name indicates, the Microsoft Clustering algorithm focuses on showing you meaningful groupings in your source data. Unlike Nave Bayes, which requires discrete content inputs and considers all input attributes of equal weight, Microsoft Clustering allows for more flexibility in input types and grouping methodologies. You can use more content type as input and you can configure the method used to create the groups or clusters. Microsoft Clustering separates your data into intelligent groupings. As we mentioned in the previous paragraph, you can use Continuous, Discrete, and most other content types. You can optionally supply a predictive value, by marking it as predict only. Be aware that Microsoft Clustering is generally not used for prediction—you use it to find natural groupings in your data.

(Smart Business Intelligence Solutions with Microsoft SQL Server® 2008, Copyright © 2009 by Kevin Goff and Lynn Langit)

Drillthrough Action Defines a dataset to be returned as a drillthrough to a more detailed level. Creating Drillthrough Actions

For the most part, Drillthrough Actions have the same properties as Actions. Drillthrough Actions do not have Target Type or Target Object properties. In their place, the Drillthrough Action has the following:

- Drillthrough Columns Defines the objects to be included in the drillthrough dataset.

Question No : 30

You design a Business Intelligence (BI) solution by using SQL Server 2008.

A data warehouse named CustomerDW contains a Fact table named FactCustomer. The FactCustomer table contains two columns named CustomerKey and CustomerSales.

You create a data mining model named CustomerModel by using SQL Server 2008 Analysis Services (SSAS).

A report that is developed by using SQL Server 2008 Reporting Services (SSRS) lists the top 50 customers based on the sales amount. The report extracts data from a SQL Server relational database.

You add a column named UpSell to the report.

You need to ensure that the UpSell column displays the probability values of the expensive products that customers are likely to purchase.

Which Data Mining Extensions (DMX) query should you use?

- A.** SELECT PredictProbability(t.[UpSell]) as [UpSell],
[CustomerKey], m.[CustomerSales]
From [CustomerModel] m
PREDICTION JOIN OPENQUERY([CustomerDW],
'SELECT
[CustomerKey], [CustomerSales] From FactCustomer
ORDER BY [CustomerSales]
') AS t
ON m.[CustomerKey] = t.[CustomerKey]
- B.** SELECT PredictProbability(m.[UpSell]) as [UpSell],
[CustomerKey], t.[CustomerSales]
From [CustomerModel] m
PREDICTION JOIN
OPENQUERY([CustomerDW],
'SELECT TOP 50
[CustomerKey], [CustomerSales]
FROM FactCustomer
ORDER BY [CustomerSales]
') AS t
- C.** SELECT PredictProbability(m.[UpSell]) as [UpSell],
[CustomerKey], t.[CustomerSales]
From [CustomerModel] m
PREDICTION JOIN OPENQUERY([CustomerDW],
'SELECT TOP 50
[CustomerKey],[CustomerSales]
From FactCustomer

```

ORDER BY [CustomerSales]
') AS t
ON m.[CustomerKey] = t.[CustomerKey]
D. SELECT Probability(m.[UpSell]) as [UpSell],
[CustomerKey], t.[CustomerSales]
From [CustomerModel] m
PREDICTION JOIN OPENQUERY([CustomerDW],
'SELECT
[CustomerKey], [CustomerSales]
From FactCustomer
ORDER BY [CustomerSales]
') AS t
ON m.[CustomerKey] = t.[CustomerKey]

```

Answer: C

Explanation:

<http://msdn.microsoft.com/en-us/library/ms132031.aspx>

SELECT FROM <model> PREDICTION JOIN (DMX)

SQL Server 2008 R2 Other Versions

Uses a mining model to predict the states of columns in an external data source. The PREDICTION JOIN

statement matches each case from the source query to the model.

SELECT [FLATTENED] [TOP <n>] <select expression list>

FROM <model> | <sub select> [NATURAL] PREDICTION JOIN

<source data query> [ON <join mapping list>]

[WHERE <condition expression>]

[ORDER BY <expression> [DESC|ASC]]

DMX Prediction Queries

The Data Mining Extensions language is modeled after the SQL query language. You probably recognized the SQL-like syntax of SELECT...FROM...JOIN...ON from the code generated in the preceding example. Note the use of the DMX Predict function and the PREDICTION JOIN keyword.

Question No : 31

You design a Business Intelligence solution by using SQL Server 2008.

The solution includes a SQL Server 2008 Analysis Services (SSAS) database. You design a data mining structure. The structure is used by a data mining model that predicts the expected growth for a particular market location. The model also populates a data mining dimension named Market Location.

The database includes a cube that contains a calculated member named Predicted Revenue. The calculated member uses predictions from the data mining model.

You have the following business requirements:

The view shown in the following exhibit must be displayed to consultants. (Click the Exhibit button.)

Market Location ▼	Predicted Revenue
Bay Area	#N/A
Europe HQ	#N/A
Germany	#N/A
Gulf Coast	#N/A
Microsoft Dynamics	#N/A
Midwest	#N/A
Northeast	#N/A
Pacific Northwest	#N/A
Pacific Southwest	#N/A
Portugal	#N/A
Rocky Mountains	#N/A
Southeast	#N/A
Southwest	#N/A
Spain	#N/A
Support Services	#N/A
UK	#N/A
Unassigned	#N/A
Grand Total	#N/A

The view shown in the following exhibit must be displayed to managers. (Click the Exhibit button.)

Market Location ▼	Predicted Revenue
Bay Area	16,190,550
Europe HQ	
Germany	12,000
Gulf Coast	24,857,499
Microsoft Dynamics	483,327
Midwest	21,325,118
Northeast	15,804,479
Pacific Northwest	34,562,679
Pacific Southwest	34,511,559
Portugal	1,480,058
Rocky Mountains	31,032,424
Southeast	26,686,317
Southwest	56,455,624
Spain	2,277,760
Support Services	43,200
UK	8,692,785
Unassigned	-6,650
Grand Total	274,408,728

You need to design a solution that meets the business requirements.

What should you do?

- A. Implement cell-level security on the cube.
- B. Implement drillthrough security on the cube.
- C. Implement dimension-level security on the Market Location dimension.
- D. Create a new reference dimension that joins Windows user names and their allowed market locations. Implement dimension-level security on the new reference dimension.

Answer: A

Explanation:

Question No : 32

You design a Business Intelligence (BI) solution by using SQL Server 2008. The solution includes a SQL Server 2008 Analysis Services (SSAS) database.

A measure group in the database contains transaction details. The transaction details include the price, volume of shares, trade type, and several other attributes of each transaction.

You need to implement a data mining model that will estimate the future prices based on the existing transaction data.

Which algorithm should the data mining model use?

- A. the Microsoft Clustering algorithm
- B. the Microsoft Association algorithm
- C. the Microsoft Naive Bayes algorithm
- D. the Microsoft Neural Network algorithm

Answer: D

Explanation:

Microsoft Neural Network Algorithm

Microsoft Neural Network is by far the most powerful and complex algorithm. This algorithm creates classification and regression mining models by constructing a Multilayer Perceptron network of neurons. Similar to the Microsoft Decision Trees algorithm, the Microsoft Neural Network algorithm calculates probabilities for each possible state of the input attribute when given each state of the predictable attribute. You can later use these probabilities to predict an outcome of the predicted attribute, based on the input attributes. It is recommended for use when other algorithms fail to produce meaningful results, such as those measured by a lift chart output. We often use Microsoft Neural Network as a kind of a last resort, when dealing with large and complex datasets that fail to produce meaningful results when processed using other algorithms. This algorithm can accept a data type of Discrete or Continuous as input

(Smart Business Intelligence Solutions with Microsoft SQL Server® 2008, Copyright © 2009 by Kevin Goff and Lynn Langit)

Microsoft Neural Network

Neural networks were developed in the 1960s to model the way human neurons function. Microsoft has created the Microsoft Neural Network algorithm so we can use neural networks for such mundane activities as predicting product sales. Of course, predicting product sales might not seem so mundane if your future employment is dependent on being correct.

Function The Microsoft Neural Network algorithm creates a web of nodes that connect inputs derived from attribute values to a final output. The combination function determines how to combine the inputs coming into the node.

Certain inputs might get more weight than others when it comes to affecting the output from this node.

The second function in each node is the activation function. The activation function takes input from the combination function and comes up with the output from this node to be sent to the next node in the network.

Tasks

The main purposes of the Microsoft Neural Network algorithm are

- Classification
- Regression

(McGraw-Hill - Delivering Business Intelligence with Microsoft SQL Server 2008 (2009))

<http://msdn.microsoft.com/en-us/library/ms174941.aspx>

Microsoft Neural Network Algorithm

In SQL Server Analysis Services, the Microsoft Neural Network algorithm combines each possible state of the input attribute with each possible state of the predictable attribute, and uses the training data to calculate probabilities. You can later use these probabilities for classification or regression, and to predict an outcome of the predicted attribute, based on the input attributes.

A mining model that is constructed with the Microsoft Neural Network algorithm can contain multiple networks, depending on the number of columns that are used for both input and prediction, or that are used only for prediction. The number of networks that a single mining model contains depends on the number of states that are contained by the input columns and predictable columns that the mining model uses.

Creating Predictions

After the model has been processed, you can use the network and the weights stored within each node to make predictions. A neural network model supports regression, association, and classification analysis. Therefore, the meaning of each prediction might be different. You can also query the model itself, to review the correlations that were found and retrieve related statistics. For examples of how to create queries against a neural network model, see [Querying a Neural Network Model \(Analysis Services- Data Mining\)](#).

Remarks

Does not support drillthrough or data mining dimensions. This is because the structure of the nodes in the mining model does not necessarily correspond directly to the underlying data. Does not support the creation of models in Predictive Model Markup Language (PMML) format. Supports the use of OLAP mining models.

Does not support the creation of data mining dimension

Question No : 33

You design a Business Intelligence (BI) solution by using SQL Server 2008.

Your company processes all transaction data in a Point of Sale (POS) application.

Based on the transaction data, you design a solution to predict the type of products a customer tends to purchase on a single visit.

You need to identify the appropriate algorithm to design the solution.

Which algorithm should you use?

- A. Clustering
- B. Native Bayes
- C. Association Rules
- D. Time Series

Answer: C

Explanation:

Microsoft Association Rules

The Microsoft Association Rules algorithm deals with items that have been formed into sets within the data. In the example here, we look at sets of products that were purchased by the same customer. The Business Intelligence Development Studio provides us with three viewers for examining these sets.

Question No : 34

You design a Business Intelligence (BI) solution by using SQL Server 2008.

Your solution includes a data mining structure that uses SQL Server 2008 Analysis Services (SSAS) as its data source. The measure groups use 100 percent multidimensional online analytical processing (MOLAP) storage.

You need to provide detailed information on the training and test data to ensure the accuracy of the mining model. You also need to minimize the time required to create the training and test data.

Which two tasks should you perform? (Each correct answer presents part of the solution. Choose two.)

- A. Perform cross-validation queries to the test and training data.
- B. Create a new mining structure that has a holdout value.
- C. Create a SQL Sever 2008 Integration Services (SSIS) package that partitions test and training datasets and merges case and nested tables.
- D. Use a Sort Data Flow transformation.
- E. Use an ORDER BY clause in the Data Flow source query. Define a SortKeyPosition ordinal key for the appropriate output column.

Answer: A,B

Explanation:

Cross Validation

The cross validation tool was added specifically to address requests from enterprise customers. Keep in mind that cross validation does not require separate training and testing datasets. You can use testing data, but you won't always need to. This elimination of the need for holdout (testing) data can make cross validation more convenient to use for data mining model validation. Cross validation works by automatically separating the source data into partitions of equal size. It then performs iterative testing against each of the partitions and shows the results in a detailed output grid. Cross validation works according to the value specified in the Fold Count parameter on the Cross Validation tab of the Mining Accuracy Chart tab in BIDS. The default value for this parameter is 10, which equates to 10 sets. If you're using temporary mining models to cross validate in Excel 2007, 10 is the maximum number of allowable folds. If you're using BIDS, the maximum number is 256. Of course, a greater number of folds equates to more processing overhead.

You can also implement cross validation using newly introduced stored procedures. A reason to use the new cross-validation capability is that it's a quick way to perform validation using multiple mining models as source inputs.

Note Cross validation cannot be used to validate models built using the Time Series or Sequence Clustering algorithms. This is logical if you think about it because both of these algorithms depend on sequences and if the data was partitioned for testing, the validity of the sequence would be violated

Question No : 35

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The solution has been deployed by using default settings on a SQL Server 2008 Analysis

Services (SSAS) instance. The solution has a large cube that processes 10 million fact rows.

You frequently encounter out-of-memory exceptions when the cube is processed.

You need to recommend a solution to resolve the out-of-memory exceptions when the cube is processed. You want to achieve this task by using the minimum amount of development effort.

What should you do?

- A. Reduce the number of aggregations.
- B. Partition the cube. Process the cube based on each partition.
- C. Increase the physical memory available to the SSAS instance by modifying the Memory\TotalMemoryLimit server property.
- D. Increase the physical memory available to the SSAS instance by modifying the OLAP\Process \BufferMemoryLimit server property.

Answer: D

Explanation:

Question No : 36

You design a Business Intelligence (BI) solution by using SQL Server 2008.

You deploy a SQL Server 2008 Analysis Services (SSAS) cube. The cube contains a measure group that uses table binding. The measure group contains 200 million rows.

A job that processes the measure group fails. The log shows an out-of-memory error. The job uses the Process Update option.

You need to resolve the error. You need to perform this action without increasing the available physical memory for the SSAS instance.

What should you do?

- A. Change the job to process the cube.

- B. Change the job to process the measure group with the Process Full option.
- C. Increase the number of partitions in the measure group.
- D. Increase the number of aggregations in the measure group.

Answer: C

Explanation:

Question No : 37

You design a Business Intelligence (BI) solution by using SQL Server 2008. Your solution includes relational and analysis services.

The solution has a cube that is queried by more than 650 users. During peak hours, more than 100 active connections are open on the cube at any given time. Users connect to and query the cube by using custom-built applications.

You need to view the connection details and the application name that is used to connect to the cube of all users.

What should you do?

- A. Use the Resource Governor.
- B. Use the Database Tuning Advisor.
- C. Use the Analysis Services performance counters.
- D. Prepare a report by using a dynamic management view.

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms188754.aspx>

Dynamic Management Views and Functions (Transact-SQL)

Dynamic management views and functions return server state information that can be used to monitor the health of a server instance, diagnose problems, and tune performance.

Important Dynamic management views and functions return internal, implementation-specific state data. Their schemas and the data they return may change in future releases of SQL Server. Therefore, dynamic management views and functions in future releases may not be compatible with the dynamic management views and functions in this release. For example, in future releases of SQL Server, Microsoft may augment the definition of any dynamic management view by adding columns to the end of the column list. We

recommend against using the syntax `SELECT * FROM dynamic_management_view_name` in production code because the number of columns returned might change and break your application.

There are two types of dynamic management views and functions:

- Server-scoped dynamic management views and functions. These require `VIEW SERVER STATE` permission on the server.
- Database-scoped dynamic management views and functions. These require `VIEW DATABASE STATE` permission on the database.

Querying Dynamic Management Views

Dynamic management views can be referenced in Transact-SQL statements by using two-part, three-part, or four-part names. Dynamic management functions on the other hand can be referenced in Transact-SQL statements by using either two-part or three-part names.

Dynamic management views and functions cannot be referenced in Transact-SQL statements by using one-part names.

All dynamic management views and functions exist in the `sys` schema and follow this naming convention `dm_*`. When you use a dynamic management view or function, you must prefix the name of the view or function by using the `sys` schema. For example, to query the `dm_os_wait_stats` dynamic management view, run the following query:

```
SELECT wait_type, wait_time_ms  
FROM sys.dm_os_wait_stats;  
GO
```

Question No : 38

You design a Business Intelligence (BI) solution by using SQL Server 2008. A SQL Server 2008 Reporting Services (SSRS) instance contains a report. Employees execute the report by using Report Manager.

Each employee has his own specific set of parameters to execute the report. Data for the report is updated once daily. For each employee, the report takes more than five minutes to execute.

You discover that data retrieval takes most of the time during report execution.

You need to reduce the execution time of the report.

What should you do?

- A. Create a report execution snapshot.
- B. Create a data-driven subscription that uses the NULL delivery method.
- C. Create a data-driven subscription that uses the file share delivery method.
- D. Create a standard subscription that uses the file share delivery method.

Answer: B

Explanation:

Data-Driven Subscriptions

A better name for a data-driven subscription might be “mass mailing.” The data-driven subscription enables us to take a report and e-mail it to a number of people on a mailing list. The mailing list can be queried from any valid Reporting Services data source. The mailing list can contain fields in addition to the recipient’s e-mail address, which are used to control the content of the e-mail sent to each recipient.

<http://msdn.microsoft.com/en-us/library/ms159762.aspx>

Standard and Data-Driven Subscriptions

Reporting Services supports two kinds of subscriptions: standard and data-driven.

Standard subscriptions are created and managed by individual users. A standard subscription consists of static values that cannot be varied during subscription processing. For each standard subscription, there is exactly one set of report presentation options, delivery options, and report parameters.

Data-driven subscriptions get subscription information at run time by querying an external data source that provides values used to specify a recipient, report parameters, or application format. You might use data-driven subscriptions if you have a very large recipient list or if you want to vary report output for each recipient. To use data-driven subscriptions, you must have expertise in building queries and an understanding of how parameters are used. Report server administrators typically create and manage these subscriptions.

Delivery Extensions

Subscriptions use delivery extensions to determine how to distribute a report and in what format. When a user creates a subscription, he or she can choose one of the available delivery extensions to determine how the report is delivered. Reporting Services includes the following delivery extensions. Developers can create additional delivery extensions to route reports to other locations.

Windows File Share - Delivers a report as a static application file to a shared folder that is accessible on the network.

E-mail - Delivers a notification or a report as an e-mail attachment or URL link.

SharePoint library - Delivers a report as a static application file to a SharePoint library that is accessible from a SharePoint site. The site must be integrated with a report server that

runs in SharePoint integrated mode.

Null - The null delivery provider is a highly specialized delivery extension that is used to preload a cache with ready-to-view parameterized reports. This method is not available to users in individual subscriptions. Null delivery is used by administrators in data-driven subscriptions to improve report server performance by preloading the cache.

Question No : 39

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The solution includes several SQL Server 2008 Integration Services (SSIS) packages. The SSIS packages import data from files located on other servers.

The packages will be deployed to a SQL Server 2008 instance and scheduled to run through the SQL Server Agent service. The SQL Server Agent service runs under a local user account.

The SSIS packages fail to run when the SQL Server Agent jobs are executed.

You need to ensure that the packages run successfully in the production environment.

What should you do?

- A. Configure the SQL Server Agent job step to run as a proxy account.
- B. Configure the SQL Server Agent job to use the sa account as the job owner.
- C. Configure the SQL Server Agent service to use the Local Service account.
- D. Configure the SQL Server Agent service to use a local administrator account.

Answer: A

Explanation:

<http://msdn.microsoft.com/en-us/library/ms189064.aspx>

Creating SQL Server Agent Proxies

A SQL Server Agent proxy defines the security context for a job step. A proxy provides SQL Server Agent with access to the security credentials for a Microsoft Windows user.

Each proxy can be associated with one or more subsystems. A job step that uses the proxy can access the specified subsystems by using the security context of the Windows user.

Before SQL Server Agent runs a job step that uses a proxy, SQL Server Agent impersonates the credentials defined in the proxy, and then runs the job step by using that security context.

Job steps that execute Transact-SQL do not use SQL Server Agent proxies. Transact-SQL job steps run in the security context of the owner of the job. To set the security context for a Transact-SQL job step, use the `database_user_name` parameter in the `sp_add_jobstep` stored procedure.

Question No : 40

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The SQL Server 2008 Integration Services (SSIS) developers use a SQL Server 2008 instance as the primary development environment.

All the SSIS packages contain data connection managers that use SQL Server authentication to extract data. The packages are saved by using the `EncryptAllWithUserKey` package protection level.

You plan a package migration strategy from the development environment to a production environment.

Migration will be performed by using an automated script.

You need to ensure that the packages execute without error in the production environment.

What should you do?

- A. Create a package configuration for every package that uses a SQL Server table.
- B. Create a package configuration for every package that uses an XML configuration file.
- C. Export each package and change the package protection level to `DontSaveSensitive`.
- D. Export each package and change the package protection level to `EncryptSensitiveWithPassword`.

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms141747.aspx>

Setting the Protection Level of Packages

To protect the data in an Integration Services package, you can set a protection level that helps protect just sensitive data or all the data in the package. Furthermore, you can encrypt this data with a password or a user key, or rely on the database to encrypt the data. Also, the protection level that you use for a package is not necessarily static, but changes throughout the life cycle of the package. You often set one protection level during development and another as soon as you deploy the package.

Do not save sensitive (DontSaveSensitive) - Suppresses the values of sensitive properties in the package when the package is saved. This protection level does not encrypt, but instead it prevents properties that are marked sensitive from being saved with the package and therefore makes the sensitive data unavailable to other users. If a different user opens the package, the sensitive information is replaced with blanks and the user must provide the sensitive information.

Encrypt all with password (EncryptAllWithPassword) - Uses a password to encrypt the whole package. The package is encrypted by using a password that the user supplies when the package is created or exported. To open the package in SSIS Designer or run the package by using the dtexec command prompt utility, the user must provide the package password. Without the password the user cannot access or run the package.

Encrypt all with user key (EncryptAllWithUserKey) - Uses a key that is based on the current user profile to encrypt the whole package. Only the user who created or exported the package can open the package in SSIS Designer or run the package by using the dtexec command prompt utility.

Encrypt sensitive with password (EncryptSensitiveWithPassword) - Uses a password to encrypt only the values of sensitive properties in the package. DPAPI is used for this encryption. Sensitive data is saved as a part of the package, but that data is encrypted by using a password that the current user supplies when the package is created or exported. To open the package in SSIS Designer, the user must provide the package password. If the password is not provided, the package opens without the sensitive data and the current user must provide new values for sensitive data. If the user tries to execute the package without providing the password, package execution fails.

Encrypt sensitive with user key (EncryptSensitiveWithUserKey) - Uses a key that is based on the current user profile to encrypt only the values of sensitive properties in the package. Only the same user who uses the same profile can load the package. If a different user opens the package, the sensitive information is replaced with blanks and the current user must provide new values for the sensitive data. If the user attempts to execute the

package, package execution fails. DPAPI is used for this encryption.

Rely on server storage for encryption (ServerStorage) - Protects the whole package using SQL Server database roles. This option is supported only when a package is saved to the SQL Server msdb database. It is not supported when a package is saved to the file system from Business Intelligence Development Studio.

Question No : 41

You design a Business Intelligence (BI) solution by using SQL Server 2008.

The SQL Server 2008 instance hosts a database. The database is currently scheduled for a full backup on a monthly basis. The 4-terabyte database contains 3.5 terabyte of data in a read-only filegroup. The database uses the bulklogged recovery model.

You need to back up the database changes to a tape drive every night by using minimum storage space and time.

Which backup strategy should you use?

- A. File backup
- B. Partial backup
- C. Differential backup
- D. Differential Partial backup

Answer: D

Explanation:

<http://msdn.microsoft.com/en-us/library/ms191539.aspx>

Partial Backups

Partial backups were introduced in SQL Server 2005. Partial backups are designed for use under the simple recovery model to improve flexibility for backing up very large databases that contain one or more read-only filegroups. However, partial backups work on all databases, regardless of recovery model.

Note

For an overview of the various types of backups, see either Backup Under the Simple Recovery Model or Backup Under the Full Recovery Model.

A partial backup resembles a full database backup, but a partial backup does not contain all the filegroups. Instead, a partial backup contains all the data in the primary filegroup,

every read/write filegroup, and any optionally-specified read-only files. Partial backups are useful whenever you want to exclude read-only filegroups. A partial backup of a read-only database contains only the primary filegroup.

Question No : 42

You administer a server that runs SQL Server 2008 Integration Services (SSIS) packages. The packages are stored on the file system.

You need to integrate the backup and restore processes for the SSIS packages along with the regular SQL Server backup process.

What should you do?

- A. Deploy the packages to SQL Server 2008 and back up the msdb system database.
- B. Deploy the packages to SQL Server 2008 and back up the master system database.
- C. In the ProtectionLevel package property, select the ServerStorage option.
- D. In the ProtectionLevel package property, select the DontSaveSensitive option.

Answer: A

Explanation:

Question No : 43

You administer a SQL Server 2008 Reporting Services (SSRS) environment.

You restore the ReportServer and ReportServerTempDB databases to a new server. When you browse to the Report Manager Web page, you receive an error message. You are unable to view the folder structure and the reports.

You need to view the folder structure and the reports.

What should you do?

- A. Restore the symmetric key.
- B. Restore the msdb database.
- C. Restore the master database.
- D. Configure the IIS virtual directory.

Answer: A

Explanation:

<http://technet.microsoft.com/en-us/library/ms156016.aspx>

Report Server Database

A report server is a stateless server that uses the SQL Server Database Engine to store metadata and object definitions. A Reporting Services installation uses two databases to separate persistent data storage from temporary storage requirements. The databases are created together and bound by name. By default, the database names are reportserver and reportservertempdb, respectively.

The databases can run on a local or remote Database Engine instance. Choosing a local instance is useful if you have sufficient system resources or want to conserve software licenses, but running the databases on a remote computer can improve performance.

The report server database is a SQL Server database that stores the following content:

- Items managed by a report server (reports and linked reports, shared data sources, report models, folders, resources) and all of the properties and security settings that are associated with those items.
- Subscription and schedule definitions.
- Report snapshots (which include query results) and report history.
- System properties and system-level security settings.
- Report execution log data.
- Symmetric keys and encrypted connection and credentials for report data sources.

Administering a Report Server Database

A Reporting Services deployment uses two SQL Server relational databases for internal storage. By default, the databases are named ReportServer and ReportServerTempdb. ReportServerTempdb is created with the primary report server database and is used to store temporary data, session information, and cached reports.

In Reporting Services, database administration tasks include backing up and restoring the report server databases and managing the encryption keys that are used to encrypt and decrypt sensitive data.

Moving the Report Server Databases to Another Computer

Both the reportserver and reportservertempdb databases must be moved or copied together. A Reporting

Services installation requires both databases; the reportservertempdb database must be

related by name to the primary reportserver database you are moving.

Moving a database does not effect scheduled operations that are currently defined for report server items.

- Schedules will be recreated the first time that you restart the Report Server service.
- SQL Server Agent jobs that are used to trigger a schedule will be recreated on the new database instance.

You do not have to move the jobs to the new computer, but you might want to delete jobs on the computer that will no longer be used.

- Subscriptions, cached reports, and snapshots are preserved in the moved database. If a snapshot is not picking up refreshed data after the database is moved, clear the snapshot options in Report Manager, click

Apply to save your changes, re-create the schedule, and click Apply again to save your changes.

- Temporary report and user session data that is stored in reportservertempdb are persisted when you move that database.

Use the following steps to move the databases:

1. Backup the encryption keys for the report server database you want to move. You can use the Reporting Services Configuration tool backup the keys.
2. Stop the Report Server service. You can use the Reporting Services Configuration tool to stop the service.
3. Start SQL Server Management Studio and open a connection to the SQL Server instance that hosts the report server databases.
- 4 Right-click the report server database, point to Tasks, and click Detach. Repeat this step for the report server temporary database.
5. Copy or move the .mdf and .ldf files to the Data folder of the SQL Server instance you want to use. Because you are moving two databases, make sure that you move or copy all four files.
6. In Management Studio, open a connection to the new SQL Server instance that will host the report server databases.
- 7.. Right-click the Databases node, and then click Attach.
8. Click Add to select the report server database .mdf and .ldf files that you want to attach. Repeat this step for the report server temporary database.
9. After the databases are attached, verify that the RSExecRole is a database role in the report server database and temporary database. RSExecRole must have select, insert, update, delete, and reference permissions on the report server database tables, and execute permissions on the stored procedures. For more information, see How to: Create the RSExecRole.
10. Start the Reporting Services Configuration tool and open a connection to the report server.
11. On the Database page, select the new SQL Server instance, and then click Connect.

12. Select the report server database that you just moved, and then click Apply.
13. On the Encryption Keys page, click Restore. Specify the file that contains the backup copy of the keys and the password to unlock the file.
14. Restart the Report Server service.

Question No : 44

You administer a SQL Server 2005 Reporting Services server. Your company publishes reports to a public Web site. Customers can view the reports without providing user credentials.

You plan to upgrade the server to use SQL Server 2008 Reporting Services.

You need to ensure that customers can continue to view the reports without providing user credentials.

What should you do?

- A. Enable Basic authentication.
- B. Enable Anonymous access on the IIS virtual directory.
- C. Use a custom authentication extension.
- D. Select Windows Authentication and add the Guest user account.

Answer: C

Explanation:

Authentication

All users or automated processes that request access to Report Server must be authenticated before access is allowed. Reporting Services provides default authentication based on Windows integrated security and assumes trusted relationships where client and network resources are in the same domain or a trusted domain.

You can change the authentication settings to narrow the range of accepted requests to specific security packages for Windows integrated security, use Basic authentication, or use a custom forms-based authentication extension that you provide. To change the authentication type to a method other than the default, you must deploy a custom authentication extension. Previous versions of SSRS relied on IIS to perform all types of authentication. Because SSRS 2008 no longer depends on IIS, there is a new

authentication subsystem that supports this. The Windows authentication extension supports multiple authentication types so that you can precisely control which HTTP requests a report server will accept

Question No : 45

You administer a Microsoft SQL Server 2005 Reporting Services (SSRS) instance.

The instance has the following features:

You plan to upgrade the instance to SQL Server 2008 Reporting Services. You need to upgrade the instance without loss of functionality.

What should you do?

- A. Uninstall Internet Information Services (IIS).
- B. Upgrade the data extension to .NET Framework 3.5.
- C. Edit the RSWebapplication.config file to refer to the upgraded SSRS endpoint location.
- D. Install a new instance of SSRS. Migrate the existing configuration files and database to the new instance.

Answer: D

Explanation:

Question No : 46

You administer a SQL Server 2000 server.

The SQL Server 2000 server hosts a SQL Server 2000 relational data warehouse and a SQL Server 2000 Analysis Services database (OLAP database).

You plan to migrate to a new SQL Server 2008 server in a new untrusted domain.

You need to ensure that both the relational data warehouse and the OLAP database are migrated in the minimum possible time.

What should you do?

A.

- Use the Copy Database Wizard to migrate the relational data warehouse.
- Use the Migration Wizard to migrate the OLAP database and process the OLAP database.

B.

- Use the Copy Database Wizard to migrate the relational data warehouse.
- Use the Migration Wizard to migrate the OLAP database and do not process the OLAP database.

C.

- Perform a detach and attach of the relational data warehouse files from SQL Server 2000 to SQL Server2008.
- Use the Migration Wizard to migrate the OLAP database and process the OLAP database.

D.

- Perform a detach and attach of the relational data warehouse files from SQL Server 2000 to SQL Server2008.
- Use the Migration Wizard to migrate the OLAP database and do not process the OLAP database.

Answer: C

Explanation:

Explanation/Reference:

<http://msdn.microsoft.com/en-us/library/ms190794.aspx>

Detaching and Attaching Databases

The data and transaction log files of a database can be detached and then reattached to the same or another instance of SQL Server. Detaching and attaching a database is useful if you want to change the database to a different instance of SQL Server on the same computer *or to move the database*.

Note: The SQL Server on-disk storage format is the same in the 64-bit and 32-bit environments. Therefore, attach works across 32-bit and 64-bit environments. A database detached from a server instance running in one environment can be attached on a server instance that runs in another environment.

In SQL Server 2008, you can use detach and attach operations to upgrade a user database from SQL Server 2000 or SQL Server 2005. However, the following restrictions apply:

Copies of the master, model or msdb database created using SQL Server 2000 or SQL Server 2005 cannot be attached.

<http://technet.microsoft.com/en-us/library/ms143409.aspx>

<http://technet.microsoft.com/en-us/library/ms174860.aspx>

Migrating Existing Analysis Services Databases

You can use the Analysis Services Migration Wizard to upgrade Microsoft SQL Server 2000 Analysis Services databases to Microsoft SQL Server 2008 Analysis Services. During migration, the wizard copies SQL Server 2000 Analysis Services database objects and then *re-creates* them on an instance of SQL Server 2008 Analysis Services. The source databases are left intact and are not modified. After you verify that the new databases are fully operational, you can manually delete the old databases.

As a best practice, you should migrate your databases one at a time, or in small batches. This will allow you to verify that each database object appears as expected on the destination server, before you migrate additional objects. When you use the Migration Wizard, the MSSQLServerOLAPService service must be running on both the source and the destination server.

Using the Migration Wizard

You can start the Migration Wizard from an Analysis Services server node in the Object Browser in SQL Server Management Studio. You can also start the wizard at the command prompt, by running the program MigrationWizard.exe After you migrate a database, you must process the database from the original data source before you can query the database. As an administrator, you keep the Microsoft SQL Server Analysis Services objects in the production databases current by processing them. Processing is the step, or series of steps, that populate Analysis Services objects with data from relational data sources. Processing is different depending on the type of object and the selection of processing options. While the processing job is working, the affected Analysis Services objects can be accessed for querying. The processing job works inside a transaction and the transaction can be committed or rolled back. If the processing job fails, the transaction is rolled back. If the processing job succeeds, an exclusive lock is put on the object when changes are being committed, which means the object is temporarily unavailable for query or processing. During the commit phase of the transaction, queries can still be sent to the object, but they will be queued until the commit is completed. For more information about locking and unlocking during processing, see Locking and Unlocking Databases (XMLA).

Question No : 47

You design a Business Intelligence (BI) solution by using SQL Server 2008.

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Trying our product !



- ★ **100%** Guaranteed Success
- ★ **100%** Money Back Guarantee
- ★ **365 Days** Free Update
- ★ **Instant Download** After Purchase
- ★ **24x7** Customer Support
- ★ Average **99.9%** Success Rate
- ★ More than **69,000** Satisfied Customers Worldwide
- ★ Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 One Year Free Update <p>Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 Money Back Guarantee <p>To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 Security & Privacy <p>We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Guarantee & Policy | Privacy & Policy | Terms & Conditions

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.