

Exam Questions 70-463

Implementing a Data Warehouse with Microsoft SQL Server 2012

<https://www.2passeasy.com/dumps/70-463/>



NEW QUESTION 1

You are developing a project that contains multiple SQL Server Integration Services (SSIS) packages. The packages will be deployed to the SSIS catalog. One of the steps in each package accesses an FTP site to download sales transaction data.

You create project parameters to store the username and password that are used to access the FTP site. You need to ensure that the username and password values are encrypted when they are deployed. What should you do?

- A. Set the Sensitive property of the parameters to True.
- B. Set the ProtectionLevel property of the package to EncryptSensitiveWithUserKey.
- C. Change the parameters to package parameters.
- D. Change the project to the Legacy Deployment mode

Answer: A

NEW QUESTION 2

You maintain a SQL Server Integration Services (SSIS) package. The package was developed by using SQL Server 2008 Business Intelligence Development Studio (BIDS).

The package includes custom scripts that must be upgraded. You need to upgrade the package to SQL Server 2012. Which tool should you use?

- A. SSIS Upgrade Wizard in SQL Server 2008 BIDS
- B. SSIS Upgrade Wizard in SQL Server Data Tools
- C. SQL Server DTEXECUI utility (dtexecui.exe)
- D. SQL Server dtexec utility (dtexec.exe)

Answer: B

NEW QUESTION 3

DRAG DROP

You are editing a SQL Server Integration Services (SSIS) package that contains a task with a sensitive property.

You need to create a project parameter and configure it so that its value is encrypted when it is deployed to the SSIS catalog.

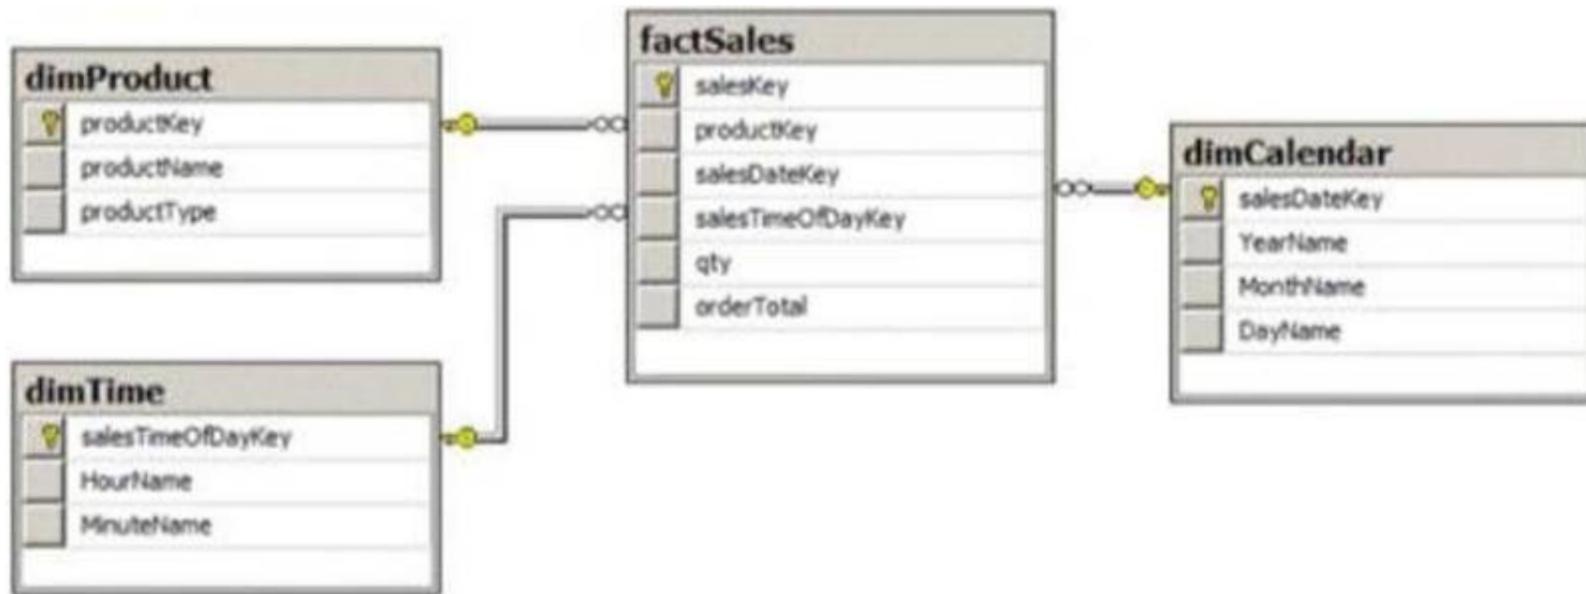
Which three steps should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Answer:

Explanation:

NEW QUESTION 4

You are reviewing the design of an existing fact table named factSales, which is loaded from a SQL Azure database by a SQL Server Integration Services (SSIS) package each day. The fact table has approximately 1 billion rows and is dimensioned by product, sales date, and sales time of day.



The database administrator is concerned about the growth of the database. Users report poor reporting performance against this database. Reporting requirements have recently changed and the only remaining report that uses this fact table reports sales by product name, sale month, and sale year. No other reports will be created against this table.

You need to reduce the report processing time and minimize the growth of the database. What should you do?

- A. Partition the table by product type.
- B. Create a view over the fact table to aggregate sales by month.
- C. Change the granularity of the fact table to month.
- D. Create an indexed view over the fact table to aggregate sales by month.

Answer: C

NEW QUESTION 5

You are designing a data warehouse with two fact tables. The first table contains sales per month and the second table contains orders per day. Referential integrity must be enforced declaratively.

You need to design a solution that can join a single time dimension to both fact tables. What should you do?

- A. Create a time mapping table.
- B. Change the level of granularity in both fact tables to be the same.
- C. Merge the fact tables.
- D. Create a view on the sales table.

Answer: B

NEW QUESTION 6

You are designing a data warehouse hosted on SQL Azure. The data warehouse currently includes the dimUser and dimDistrict dimension tables and the factSales fact table. The dimUser table contains records for each user permitted to run reports against the warehouse; and the dimDistrict table contains information about sales districts.

The system is accessed by users from certain districts, as well as by area supervisors and users from the corporate headquarters.

You need to design a table structure to ensure that certain users can see sales data for only certain districts. Some users must be permitted to see sales data from multiple districts.

What should you do?

- A. Add a district column to the dimUser table.
- B. Partition the factSales table on the district column.
- C. Create a userDistrict table that contains primary key columns from the dimUser and dimDistrict tables.
- D. For each district, create a view of the factSales table that includes a WHERE clause for the district.

Answer: C

NEW QUESTION 7

You are reviewing the design of a customer dimension table in an existing data warehouse hosted on SQL Azure.

The current dimension design does not allow the retention of historical changes to customer attributes such as Postcode.

You need to redesign the dimension to enable the full historical reporting of changes to multiple customer attributes including Postcode.

What should you do?

- A. Add StartDate and EndDate columns to the customer dimension.
- B. Add an IsCurrent column to the customer dimension.
- C. Enable Snapshot Isolation on the data warehouse.
- D. Add CurrentValue and PreviousValue columns to the customer dimension.

Answer: A

NEW QUESTION 8

You are designing an enterprise star schema that will consolidate data from three independent data marts. One of the data marts is hosted on SQL Azure.

Most of the dimensions have the same structure and content. However, the geography dimension is slightly different in each data mart.

You need to design a consolidated dimensional structure that will be easy to maintain while ensuring that all dimensional data from the three original solutions is represented.

What should you do?

- A. Create a junk dimension for the geography dimension.
- B. Implement change data capture.
- C. Create a conformed dimension for the geography dimension.
- D. Create three geography dimension

Answer: C

NEW QUESTION 9

You are developing a SQL Server Integration Services (SSIS) package. To process complex scientific data originating from a SQL Azure database, a custom task component is added to the project. You need to ensure that the custom component is deployed on a test environment correctly. What should you do?

- A. Add an OnError event handler to the SSIS project.
- B. Open a command prompt and run the gacutil command.
- C. Configure the SSIS solution to use the Project Deployment Model.
- D. Open a command prompt and run the dtexec /dumperror /conn command.
- E. Configure the output of a component in the package data flow to use a data tap.
- F. Open a command prompt and execute the package by using the SQL Log provider and running the dtexecui.exe utility.
- G. Open a command prompt and run the dtexec /rep /conn command.
- H. Run the dtutil command to deploy the package to the SSIS catalog and store the configuration in SQL Server.
- I. Use an msi file to deploy the package on the server.
- J. Open a command prompt and run the dtutil /copy command.
- K. Create a reusable custom logging component and use it in the SSIS project.

Answer: B

Explanation: Reference:

<http://msdn.microsoft.com/en-us/library/ms403356.aspx>

NEW QUESTION 10

DRAG DROP

A SQL Server Integration Services (SSIS) project has been deployed to the SSIS catalog. The project includes a project Connection Manager to connect to the data warehouse.

The SSIS catalog includes two Environments:

? Development

? QA

Each Environment defines a single Environment Variable named `ConnectionString` of type string. The value of each variable consists of the connection string to the development or QA data warehouses. You need to be able to execute deployed packages by using either of the defined Environments.

Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

In the deployed project, set the **ConnectionString** property of the Connection Manager to use the **ConnectionString** Environment Variable.

In the deployed package, set the **ConnectionString** property of the Connection Manager to use the **ConnectionString** Environment Variable.

Add the Environment reference to the deployed project.

Deploy the project to the **Development** and **QA** Environments.

Add permissions to the **Development** and **QA** Environments.

Execute the deployed package by using the required Environment.

➤

➤

Answer:

Explanation:

NEW QUESTION 10

You are developing a SQL Server Integration Services (SSIS) package. You need to design a package to change a variable value during package execution by using the least amount of development effort. What should you use?

- A. Expression task
- B. Script task
- C. Execute SQL task
- D. Execute Process task
- E. Term Extraction transformation

Answer: A

Explanation: Reference:
<http://msdn.microsoft.com/en-us/library/hh213137.aspx>

NEW QUESTION 14

DRAG DROP

A Data Flow task in a SQL Server Integration Services (SSIS) package produces run-time errors. You need to edit the package to log specific error messages. Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Answer:

Explanation:

NEW QUESTION 19

You are using a SQL Server Integration Services (SSIS) project that is stored in the SSIS catalog. An Environment has been defined in the SSIS catalog. You need to add the Environment to the project. Which stored procedure should you use?

- A. catalog.set_environment_reference_type
- B. catalog.set_environment_property
- C. catalog.create_environment_reference
- D. catalog.create_environment

Answer: C

NEW QUESTION 24

You are developing a SQL Server Integration Services (SSIS) package that imports data into a data warehouse hosted on SQL Azure. The package uses a Foreach container to process text files found in a folder. The package must be deployed to a single server by using the Project Deployment model.

Multiple SQL Server Agent jobs call the package. Each job is executed on a different schedule. Each job passes a different folder path to the package. You need to configure the package to accept the folder path from each job. Which package configuration should you use?

- A. .dtsConfig file
- B. Registry Entry
- C. Environment Variable
- D. Parent Package Variable
- E. XML Configuration File

Answer: C

Explanation: Reference:

<http://msdn.microsoft.com/en-us/library/ms345179.aspx> <http://msdn.microsoft.com/en-us/library/ms141708.aspx> <http://msdn.microsoft.com/en-us/library/hh213214.aspx> <http://msdn.microsoft.com/en-us/library/hh213296.aspx> <http://msdn.microsoft.com/en-us/library/hh213293.aspx>

NEW QUESTION 27

DRAG DROP

You are building a SQL Server Integration Services (SSIS) package to load data from all files that are automatically copied to a directory each night through an external FTP process.

You need to load data from all copied files to a destination table in SQL Server.

Which three steps should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Open the Control Flow page of the package.

On the Collection page, select the **Foreach File Enumerator** option and configure the directory and files to process.

On the Collection page, select the **Foreach ADO Enumerator** option and configure the ADO object source variable.

Create a File System task.

Create a For Loop container.

Create a Foreach Loop container.

Open the Data Flow page of the package.

Answer:

Explanation:

References:

<http://msdn.microsoft.com/en-us/library/ms141724.aspx> <http://msdn.microsoft.com/en-us/library/ms137728.aspx>

NEW QUESTION 31

You are developing a SQL Server Integration Services (SSIS) package that imports data from a relational database to a data warehouse.

You are importing data from a relational table named Projects. The table has change data capture enabled on all columns.

You need to process only the most recent values from rows that have been inserted or updated since the previous execution of the package.

Which query should you use as the data source?

- A. `SELECT * FROM cdc.fn_cdc_get_all_changes_Projects (@from_Isn, @to_Isn, N 'all ')`
- B. `SELECT * FROM cdc.fn_cdc_get_all_changes_Projects (@from_Isn, @to_Isn, N 'all update old ')`
- C. `SELECT * FROM cdc.Projects_CT WHERE @from_Isn >= _$start_Isn AND @to_Isn < _$start_Isn`
- D. `SELECT * FROM cdc.fn_cdc_get_net_changes_Projects (@from_Isn, @to_Isn)`

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

Answer: D

Explanation:

Reference: <http://msdn.microsoft.com/en-us/library/bb522511.aspx> Reference: <http://msdn.microsoft.com/en-us/library/bb510627.aspx> Reference: <http://msdn.microsoft.com/en-us/library/cc645937.aspx>

NEW QUESTION 32

You are developing a SQL Server Integration Services (SSIS) package to load data into a data warehouse.



You need to establish the correct order for loading each table to maximize parallel processing. Which order should you use?

- A.
 1. DimCurrency, DimScenario, DimAccount in parallel
 2. DimOrganization
 3. FactFinance
- B.

- A. DimCurrency, DimOrganization in parallel2. DimScenario, DimAccount in parallel3. FactFinance C.1. DimCurrency, FactFinance in parallel2. DimOrganization, DimScenario, DimAccount in parallel D.1. FactFinance2. DimOrganization, DimScenario, DimAccount in parallel3. DimCurrency E.1. DimCurrency2. DimOrganization3. DimScenario, DimAccount in parallel4. FactFinance

Answer: A

Explanation: References:

<http://msdn.microsoft.com/en-us/library/ms139892.aspx> <http://msdn.microsoft.com/en-us/library/ms141261.aspx>

NEW QUESTION 35

DRAG DROP

You are designing an extract, transform, load (ETL) process with SQL Server Integration Services (SSIS). Two packages, Package A and Package B, will be designed. Package A will execute Package B. Both packages must reference a file path corresponding to an input folder where files will be located for further processing.

You need to design a solution so that the file path can be easily configured with the least administrative and development effort.

Which four actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Create one SSIS project containing only Package A and another SSIS project containing only Package B .	
Add a Parent Package Variable configuration in Package B to set the FilePath variable.	
Add a package parameter named FilePath to Package A and Package B .	
Add a configuration file in Package A to set the FilePath value.	
Configure the project to use an Environment that sets the FilePath parameter.	
Deploy the project to the SSIS catalog.	
Create an SSIS project containing Package A and Package B .	
Add a project parameter named FilePath .	

Answer:

Explanation: Reference:

<http://msdn.microsoft.com/en-us/library/hh479588.aspx> <http://msdn.microsoft.com/en-us/library/hh213290.aspx> <http://msdn.microsoft.com/en-us/library/hh213373.aspx>

NEW QUESTION 39

You are creating a SQL Server Integration Services (SSIS) package that implements a Type 3 Slowly Changing Dimension (SCD). You need to add a task or component to the package that allows you to implement the SCD logic. What should you use?

- A. a Script component
- B. an SCD component
- C. an Aggregate component
- D. a Merge component

Answer: D

NEW QUESTION 43

DRAG DROP

You develop a SQL Server Integration Services (SSIS) project by using the Project Deployment model.

The project contains many packages. It is deployed on a server named Development1. The project will be deployed to several servers that run SQL Server 2012.

The project accepts one required parameter. The data type of the parameter is a string.

A SQL Agent job is created that will call the master.dtsx package in the project. A job step is created for the SSIS package.

The job must pass the value of an SSIS Environment Variable to the project parameter. The value of the Environment Variable must be configured differently on each server that runs SQL Server. The value of the Environment Variable must provide the server name to the project parameter.

You need to configure SSIS on the Development1 server to pass the Environment Variable to the package.

Which four actions should you perform in sequence by using SQL Server Management Studio? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

In the SSIS catalog, create an Environment named DevelopmentEnv .	
In the DevelopmentEnv Environment, create an Environment Variable. Name the Environment Variable PackageEnvironment . Set the Environment Variable value to Development1 .	
In the SQL Agent job, choose the DevelopmentEnv Environment in the Configuration tab of the SQL Server Integration Services job step.	
In the project configuration, create an Environment reference to DevelopmentEnv . Assign the value of the PackageEnvironment Environment Variable to the project parameter.	
In the project configuration, edit the value of the project parameter to DevelopmentEnv .	
In the SSIS catalog, create a folder named DevelopmentEnv .	

Answer:

Explanation: References:

<http://msdn.microsoft.com/en-us/library/hh479588.aspx> <http://msdn.microsoft.com/en-us/library/hh213230.aspx> <http://msdn.microsoft.com/en-us/library/hh213214.aspx>
http://sqlblog.com/blogs/jamie_thomson/archive/2010/11/13/ssis-server-catalogs-environmentsenvironment-variables-in-ssis-in-denali.aspx

NEW QUESTION 47

You are developing a SQL Server Integration Services (SSIS) package that imports data into a data warehouse. You add an Execute SQL task to the control flow. The task must execute a simple INSERT statement. The task has the following requirements:

- The INSERT statement must use the value of a string package variable. The variable name is StringVar.
- The Execute SQL task must use an OLE DB Connection Manager.

In the Parameter Mapping tab of the Execute SQL task, StringVar has been added as the only parameter.

You must configure the SQLStatement property of the Execute SQL task. Which SQL statement should you use?

- A. INSERT INTO dbo.Table (variablevalue) VALUES (\$StringVar)
- B. INSERT INTO dbo.Table (variablevalue) VALUES (0)
- C. INSERT INTO dbo.Table (variablevalue) VALUES (@0)
- D. INSERT INTO dbo.Table (variablevalue) VALUES (?)

Answer: D

Explanation: References:

<http://msdn.microsoft.com/en-us/library/ms141003.aspx> <http://msdn.microsoft.com/en-us/library/ms140355.aspx> <http://msdn.microsoft.com/en-us/library/cc280502.aspx>

NEW QUESTION 52

You are creating a SQL Server Master Data Services (MDS) model. This model is used to store a master list of products.

An attribute must be added to the Product entity to define the sales manager responsible for each product.

You need to create an attribute in the Product entity that prevents users from entering invalid sales manager values.

Which type of attribute should you create?

- A. Recursive
- B. Explicit
- C. Domain-based
- D. User-defined
- E. Derived
- F. Parent

Answer: C

Explanation: References:

<http://msdn.microsoft.com/en-us/library/bb190163.aspx> <http://msdn.microsoft.com/en-us/library/ee633737.aspx> <http://msdn.microsoft.com/en-us/library/ee633759.aspx> <http://msdn.microsoft.com/en-us/library/ee633745.aspx> <http://msdn.microsoft.com/en-us/library/ee633724.aspx>
<http://msdn.microsoft.com/en-us/library/ee633733.aspx> <http://msdn.microsoft.com/en-us/library/ff487058.aspx>

NEW QUESTION 53

You are creating a SQL Server Master Data Services (MDS) model for a company.

The source data for the company is stored in a single table that contains the manager-to-subordinate relationships.

You need to create a hierarchy representing the organizational structure of the company. Which hierarchy type should you use?

- A. Organizational
- B. Recursive
- C. Non-Mandatory Explicit
- D. Many-to-Many

Answer: B

NEW QUESTION 58

You are the data steward for a Business Intelligence project.

You must identify duplicate rows stored in a SQL Server table and output discoveries to a CSV file. A Data Quality Services (DQS) knowledge base has been created to support this project.

You need to produce the CSV file with the least amount of development effort. What should you do?

- A. Create an Integration Services package and use a Data Profiling transform.
- B. Create a custom .NET application based on the Knowledgebase class.
- C. Create a data quality project.
- D. Create a CLR stored procedure based on the Knowledgebase class.
- E. Create a Master Data Services (MDS) business rule.

Answer: C

Explanation: Reference:

<http://msdn.microsoft.com/en-us/library/hh213052.aspx>

NEW QUESTION 63

You manage a SQL Server Master Data Services (MDS) environment.
A new application requires access to the product data that is available in the MDS repository.
You need to design a solution that gives the application access to the product data with the least amount of development effort.
What should you do?

- A. Create a Subscription View in MDS.
- B. Access the product entity tables in the MDS database directly.
- C. Use SQL Server Integration Services (SSIS) to extract the data and put it in a staging database.
- D. Use change data capture on the product entity table

Answer: A

NEW QUESTION 68

Occasionally a job that executes an existing SQL Server Integration Services (SSIS) package does not complete and nothing is processed.
You need to ensure that package logging occurs. Your solution must minimize deployment and development efforts.
What should you do?

- A. Create a reusable custom logging component.
- B. Use the gacutil command.
- C. Use the Project Deployment Wizard.
- D. Run the package by using the dtexec /rep /conn command.
- E. Add a data tap on the output of a component in the package data flow.
- F. Create an OnError event handler.
- G. Use the dtutil /copy command.
- H. Deploy the package by using an msi file.
- I. Run the package by using the dtexec /dumperror /conn command.
- J. Run the package by using the dtexecui.exe utility and the SQL Log provider.
- K. Deploy the package to the Integration Services catalog by using dtutil and use SQL Server to store the configuration.

Answer: J

Explanation: References:

<http://msdn.microsoft.com/en-us/library/ms140246.aspx> <http://msdn.microsoft.com/en-us/library/hh231187.aspx>

NEW QUESTION 70

You are writing a SQL Server Integration Services (SSIS) package that transfers data from a legacy system.
Data integrity in the legacy system is very poor. Invalid rows are discarded by the package but must be logged to a CSV file for auditing purposes.
You need to establish the best technique to log these invalid rows while minimizing the amount of development effort.
What should you do?

- A. Add a data tap on the output of a component in the package data flow.
- B. Deploy the package by using an msi file.
- C. Run the package by using the dtexecui.exe utility and the SQL Log provider.
- D. uses the dtutil /copy command.
- E. Deploy the package to the Integration Services catalog by using dtutil and use SQL Server to store the configuration.
- F. Create an OnError event handler.
- G. uses the Project Deployment Wizard.
- H. Use the gacutil command.
- I. Create a reusable custom logging component.
- J. Run the package by using the dtexec /rep /conn command.
- K. Run the package by using the dtexec /dumperror /conn command.

Answer: A

Explanation: Reference:

<http://www.rafael-salas.com/2012/01/ssis-2012-quick-peek-to-data-taps.html> <http://msdn.microsoft.com/en-us/library/hh230989.aspx> <http://msdn.microsoft.com/en-us/library/jj655339.aspx>

NEW QUESTION 73

You are deploying a new SQL Server Integration Services (SSIS) project to the test environment. A package in the project uses a custom task component.
You need to ensure that the custom object is deployed on the test environment correctly. What should you do?

- A. Run the package by using the dtexec /rep /conn command.
- B. Create a reusable custom logging component.
- C. Create an OnError event handler.
- D. Use the gacutil command.
- E. Use the dtutil /copy command.
- F. Deploy the package to the Integration Services catalog by using dtutil and use SQL Server to store the configuration.
- G. Run the package by using the dtexec /dumperror /conn command.
- H. Use the Project Deployment Wizard.
- I. Deploy the package by using an msi file.
- J. Add a data tap on the output of a component in the package data flow.
- K. Run the package by using the dtexecui.exe utility and the SQL Log provider.

Answer: D

Explanation: Reference:

<http://msdn.microsoft.com/en-us/library/ms403356.aspx>

NEW QUESTION 78

You are deploying a project to the SQL Server Integration Services (SSIS) catalog. You need to ensure that the project is deployed to the production environment with the least amount of administrative effort. What should you do?

- A. Create an On Error event handler.
- B. Add a data tap on the output of a component in the package data flow.
- C. Use the gacutil command.
- D. Create a reusable custom logging component.
- E. Run the package by using the dtexecui.exe utility and the SQL Log provider.
- F. Run the package by using the dtexec /rep /conn command.
- G. Use the Integration Services Deployment Wizard.
- H. Run the package by using the dtexec /dumperror /conn command.
- I. Use the dtutil /copy command.
- J. Deploy the package to the Integration Services catalog by using dtutil and use SQL Server to store the configuration.
- K. Deploy the package by using an msi file

Answer: G

Explanation: Reference:

<http://msdn.microsoft.com/en-us/library/hh479588.aspx> <http://msdn.microsoft.com/en-us/library/hh213290.aspx> <http://msdn.microsoft.com/en-us/library/hh213373.aspx>

NEW QUESTION 80

You are creating a SQL Server Master Data Services (MDS) model for a company. The source data for the company is stored in a single table that contains the manager-to-subordinate relationships. You need to create a hierarchy representing the organizational structure of the company. Which hierarchy type should you use?

- A. Natural
- B. Explicit
- C. Parent
- D. Recursive

Answer: D

Explanation: Reference: <http://technet.microsoft.com/en-us/library/ff487006.aspx> Reference: <http://msdn.microsoft.com/en-us/library/ee633747.aspx> Reference: <http://technet.microsoft.com/en-us/library/ee633759.aspx>

NEW QUESTION 85

You are using the Knowledge Discovery feature of the Data Quality Services (DQS) client application to modify an existing knowledge base. In the mapping configuration, two of the three columns are mapped to existing domains in the knowledge base. The third column, named Group, does not yet have a domain. You need to complete the mapping of the Group column. What should you do?

- A. Map a composite domain to the source column.
- B. Create a composite domain that includes the Group column.
- C. Add a domain for the Group column.
- D. Add a column mapping for the Group column.

Answer: C

NEW QUESTION 86

You are installing SQL Server Data Quality Services (DQS). You need to give specific users access to the Data Quality Server. Which SQL Server application should you use?

- A. SQL Server Configuration Manager
- B. SQL Server Data Tools
- C. SQL Server Management Studio
- D. Data Quality Client

Answer: C

Explanation: Ref: <http://msdn.microsoft.com/en-us/library/hh213045.aspx>

NEW QUESTION 87

You manage a SQL Server Master Data Services (MDS) environment. A new application requires access to the product data that is available in the MDS repository. You need to design a solution that gives the application access to the product data with the least amount of development effort. What should you do?

- A. Use sp_addlinkedserver to add a linked server to access the MDS database tables directly.
- B. Create an OLE DB connection string that sets the Provider property to MDS.

- C. Use transactional replication for data synchronization.
- D. Create a Subscription View in MD

Answer: D

NEW QUESTION 89

You are completing the installation of the Data Quality Server component of SQL Server Data Quality Services (DQS). You need to complete the post-installation configuration. What should you do?

- A. Install the Analysis Services OLE DB Provider.
- B. Make the data available for DQS operations.
- C. Run the Data Quality Server Installer.
- D. Install ADOMD.NE

Answer: C

Explanation: References:
<http://msdn.microsoft.com/en-us/library/ff877917.aspx> <http://msdn.microsoft.com/en-us/library/gg492277.aspx>

NEW QUESTION 90

DRAG DROP

You are developing a SQL Server Integration Services (SSIS) project by using the Project Deployment model. The project contains many packages. It is deployed on a server named SQLTest1. The project will be deployed to several servers that run SQL Server 2012. The project accepts one required parameter. The data type of the parameter is a string. A SQL Agent job is created that will call the Loading.dtsx package in the project. A job step is created for the SSIS package. The job must pass the value of an SSIS Environment Variable to the project parameter. The value of the Environment Variable must be configured differently on each server that runs SQL Server. The value of the Environment Variable must provide the server name to the project parameter. You need to configure SSIS on the SQLTest1 server to pass the Environment Variable to the package. Which four actions should you perform in sequence by using SQL Server Management Studio?
 (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Answer:

Explanation:

NEW QUESTION 95

You are creating a SQL Server Integration Services (SSIS) package that implements a Type 3 Slowly Changing Dimension (SCD). You need to add a task or component to the package that allows you to implement the SCD logic. What should you use?

- A. a Data Conversion component
- B. an Execute SQL task that executes a MERGE statement on the database
- C. a Merge component
- D. an Expression task

Answer: C

Explanation: Note: Type of Slowly Changing Dimensions

Slowly Changing Dimensions are categorized into three types named: Type 1, Type 2, and Type3. The Type 1 SCD does not maintain the history of changing attributes, it overwrites values of the attributes. Type 2 maintains historical values for changing attributes. Type 3 that we do not use much maintains separate columns for changed attributes. SSIS SCD wizard supports both Type 1 and Type 2.

Using MERGE instead of SCD wizard

Replacement of SCD wizard with MERGE is not a straightforward technique. If the SCD has both Type 1 and Type 2 types attributes, they need to be handled separately.

NEW QUESTION 99

You are developing a SQL Server Integration Services (SSIS) package to load data into a SQL Server table on Server

- A. The package includes a data flow and is executed on Server
- B. The destination table has its own identity column.The destination data load has the following requirements: The identity values from the source table must be use
- C. Default constraints on the destination table must be ignore
- D. Batch size must be 100,000 rows.You need to add a destination and configure it to meet the requirement
- E. Which destination should you use?
- F. OLE DB Destination with Fast Load
- G. SQL Server Destination
- H. ADO NET Destination without Bulk Insert
- I. ADO NET Destination with Bulk Insert
- J. OLE DB Destination without Fast Load

Answer: A

Explanation: Reference: <http://msdn.microsoft.com/en-us/library/ms141237.aspx> Reference: <http://msdn.microsoft.com/en-us/library/ms139821.aspx> Reference: <http://msdn.microsoft.com/en-us/library/ms141095.aspx>

NEW QUESTION 104

DRAG DROP

You are editing a SQL Server Integration Services (SSIS) package that uses checkpoints. The package performs the following steps:

1. Download a sales transaction file by using FTP.
2. Truncate a staging table.
3. Load the contents of the file to the staging table.
4. Merge the data with another data source for loading to a data warehouse.



The checkpoints are currently working such that if any of the four steps fail, the package will restart from the failed step the next time it executes.

You need to modify the package to ensure that if either the Truncate Staging Table or the Load Sales to Staging task fails, the package will always restart from the Truncate Staging Table task the next time the package runs.

Which three steps should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

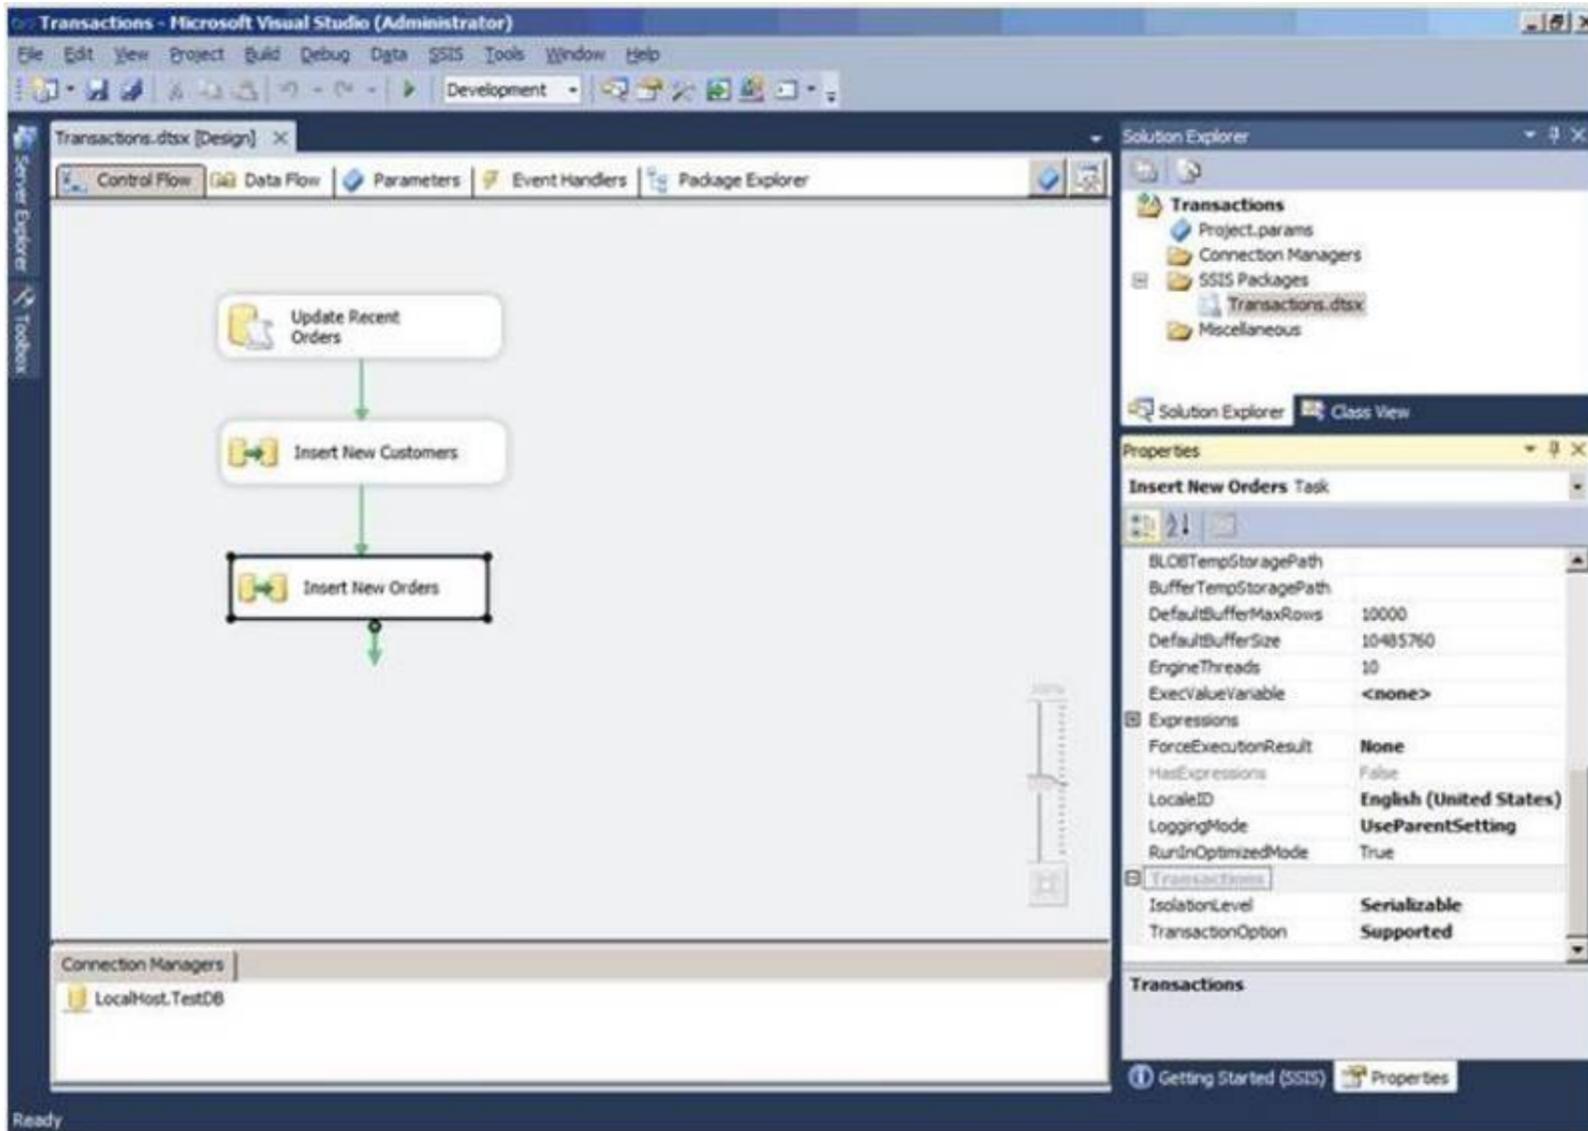
Answer:

Explanation:

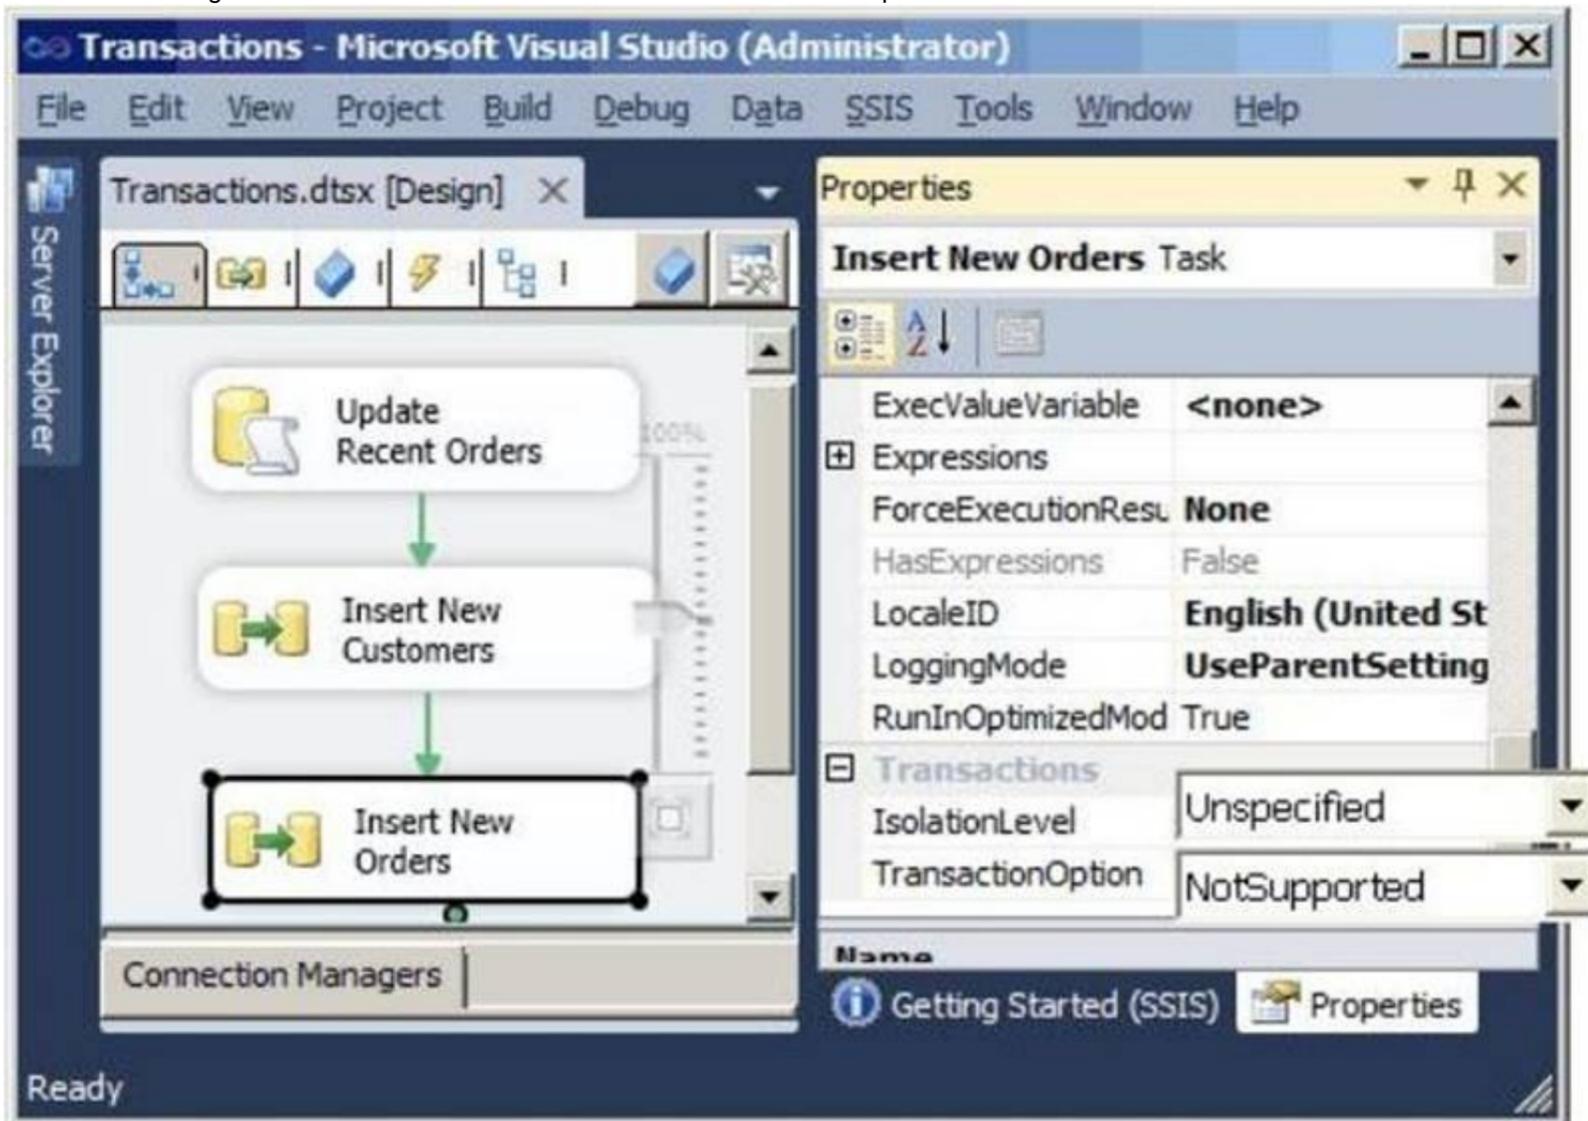
NEW QUESTION 109

HOTSPOT

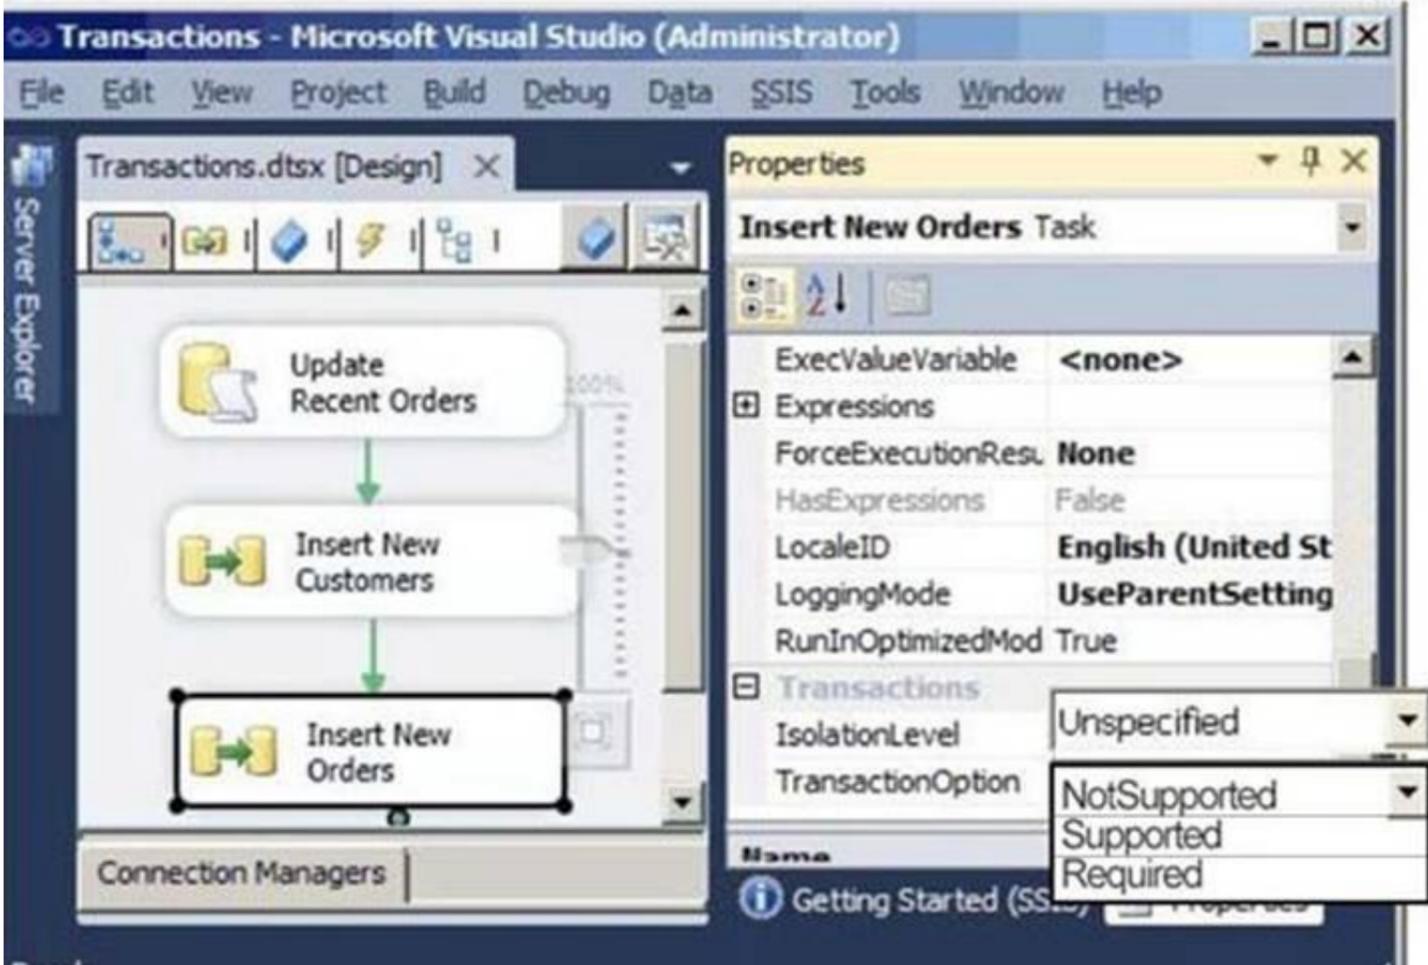
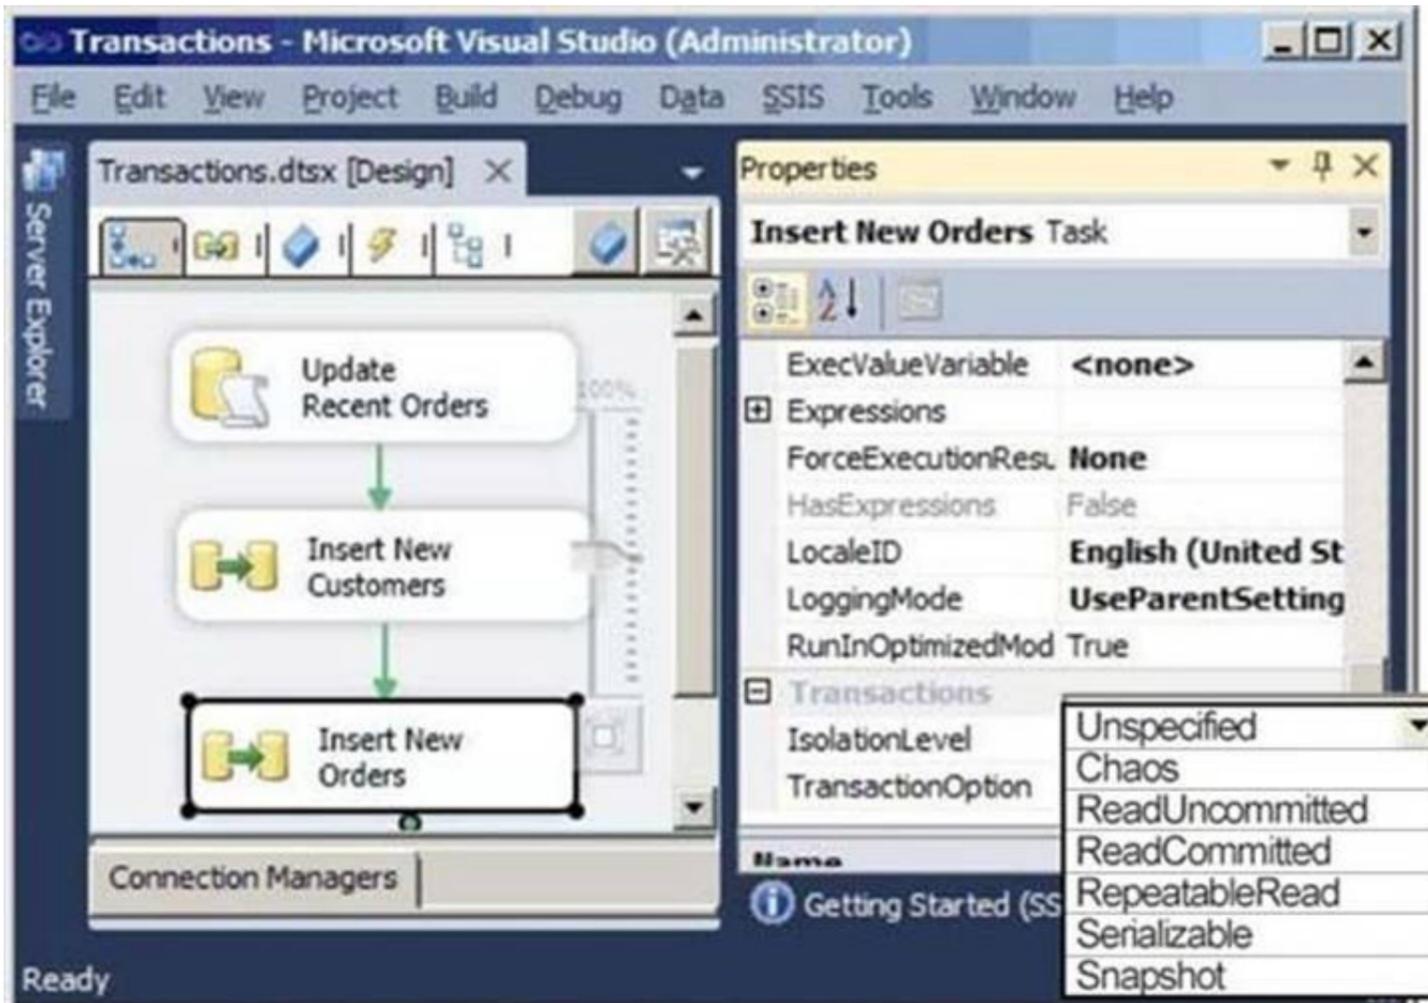
You are designing a package control flow. The package moves sales order data from a SQL Azure transactional database to an on-premise reporting database. The package will run several times a day, while new sales orders are being added to the transactional database. The current design of the package control flow is shown in the answer area. (Click the Exhibit button.)



The Insert New Orders Data Flow task must meet the following requirements: Usage of the tempdb database should not be impacted. Concurrency should be maximized, while only reading committed transactions. If the task fails, only that task needs to be rolled back. You need to configure the Insert New Orders Data Flow task to meet the requirements.

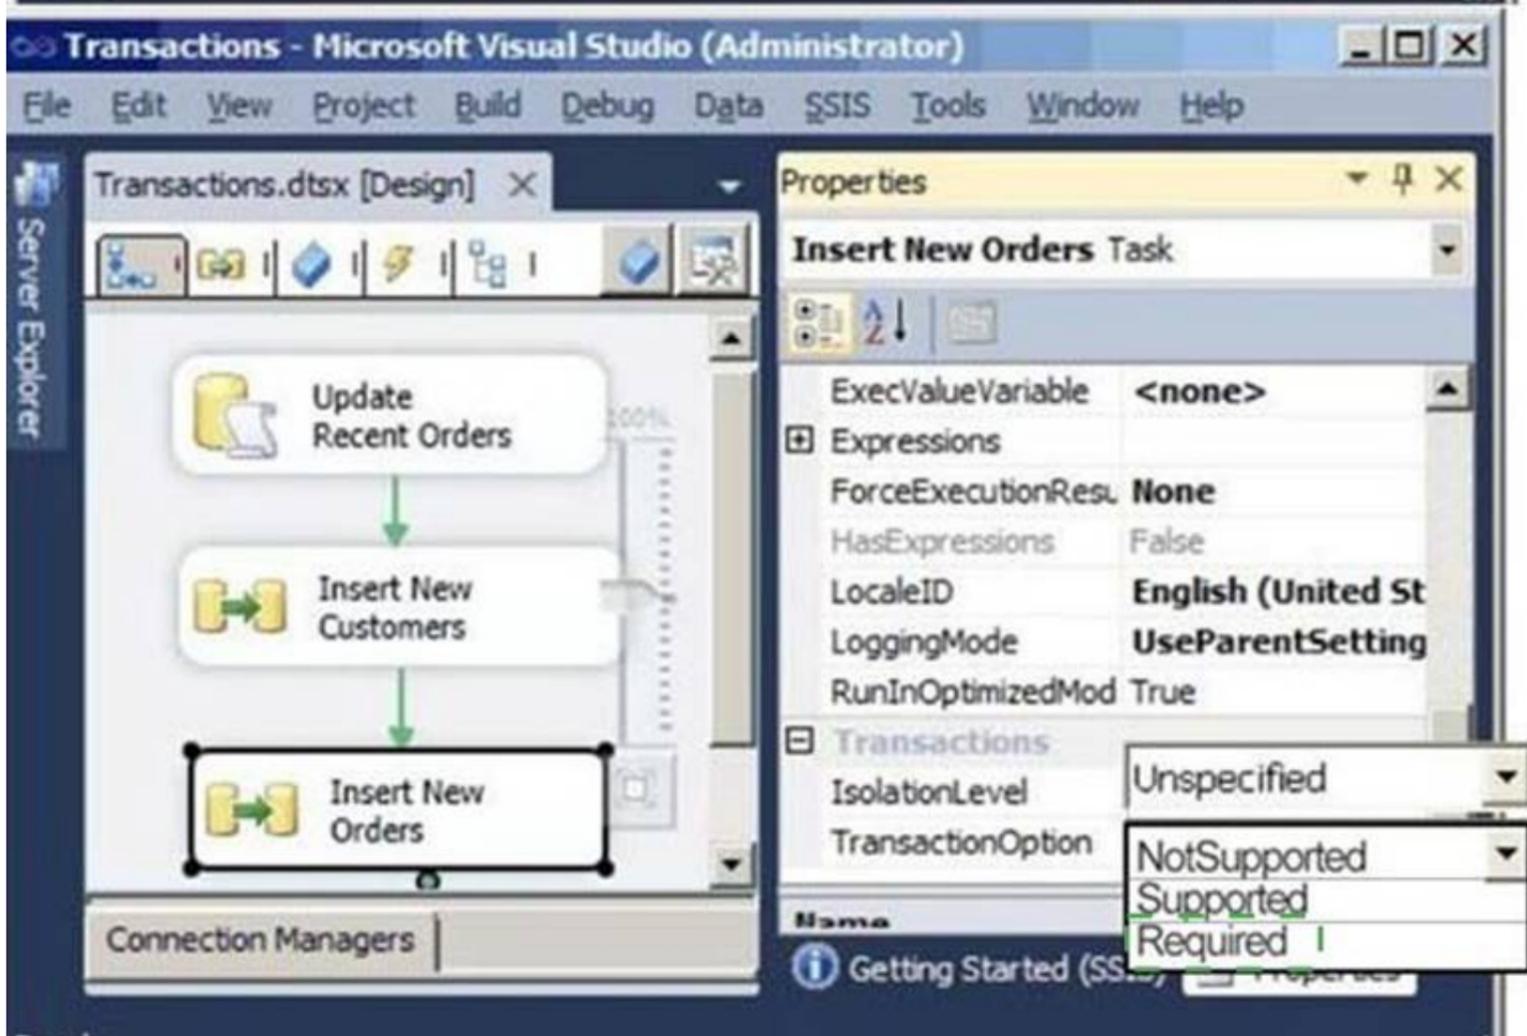
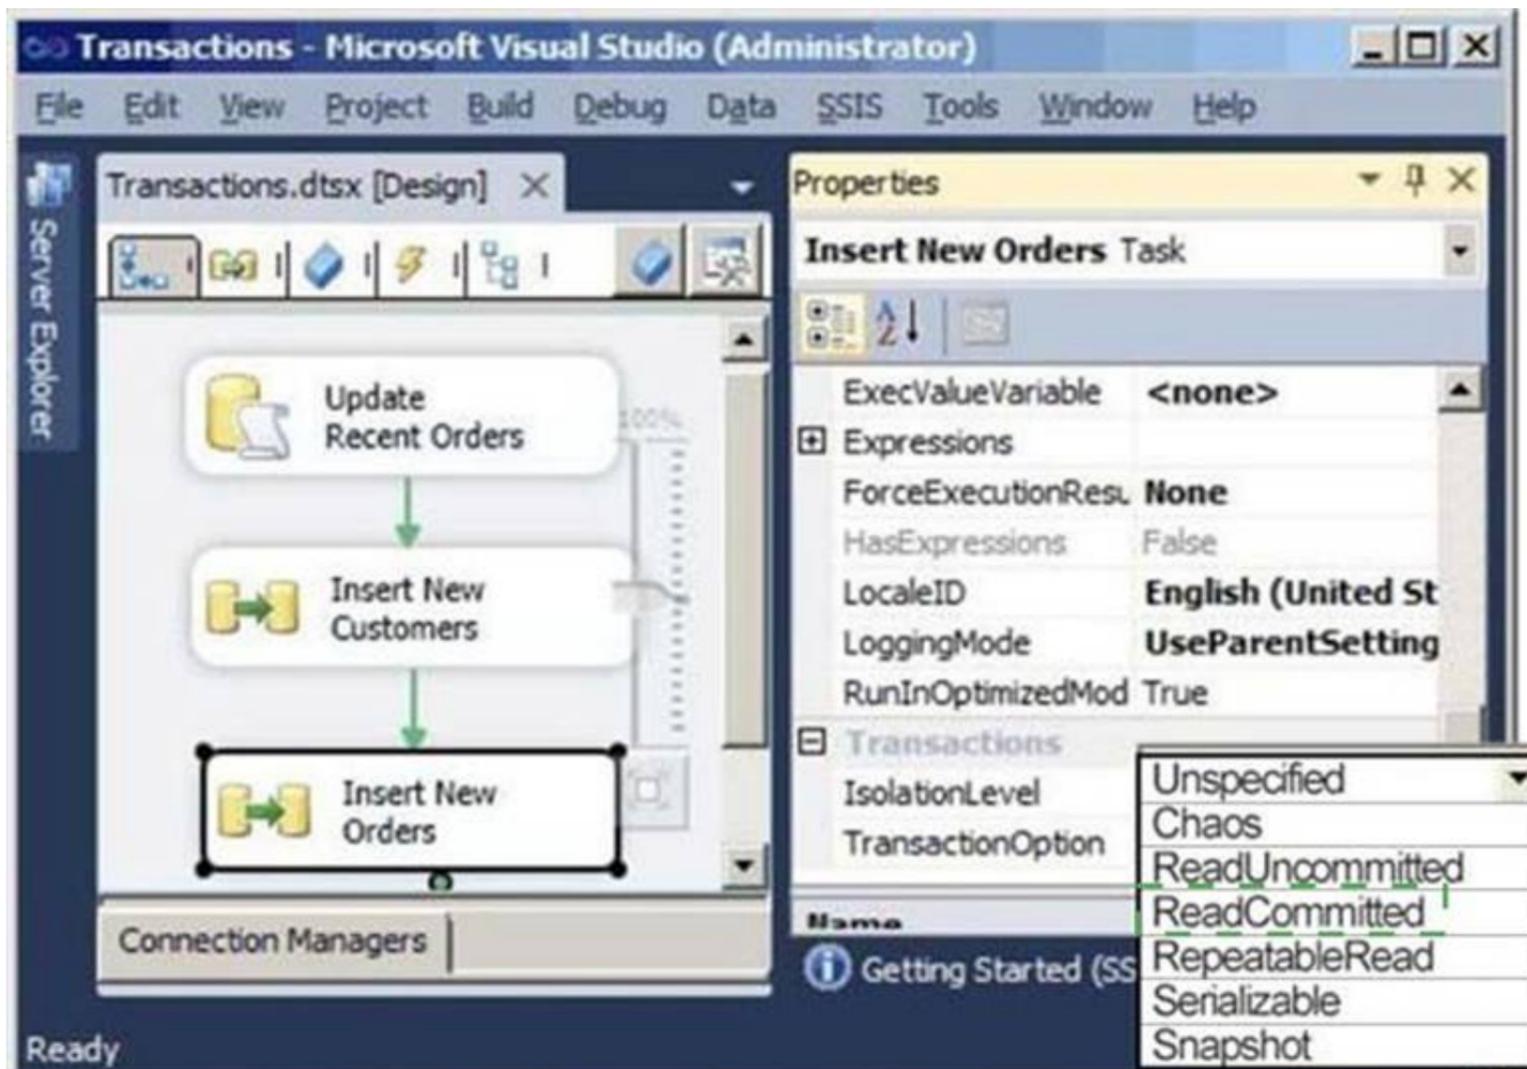


How should you configure the transaction properties? To answer, select the appropriate setting or settings in the answer area.



Answer:

Explanation:



NEW QUESTION 112

You are designing a SQL Server Integration Services (SSIS) package that uploads a file to a table named Orders in a SQL Azure database. The company's auditing policies have the following requirements:

- An entry must be written to a dedicated SQL Server log table named OrderLog.
- The entry must be written as soon as the file upload task completes.

You need to meet the company's policy requirements. Which event handler should you use?

- A. OnProgress
- B. OnInformation
- C. OnPostExecute
- D. OnComplete

Answer: C

Explanation: Reference: <http://msdn.microsoft.com/en-us/library/ms140223.aspx>

NEW QUESTION 115

You are designing a data warehouse for a fresh food distribution business that stores sales by individual product. It stores sales targets by product category. Products are classified into subcategories and categories.

Each product is included in only a single product subcategory, and each subcategory is included in only a single category.

The data warehouse will be a data source for an Analysis Services cube. The data warehouse contains two fact tables:

- factSales, used to record daily sales by product
- factProductTarget, used to record the monthly sales targets by product category

Reports must be developed against the warehouse that reports product sales by product, category and subcategory, and product sales targets.

You need to design the product dimension. The solution should use as few tables as possible while supporting all the requirements.

What should you do?

- A. Create two product tables, dimProduct and dimProductCategory
- B. Connect factSales to dimProduct and factProductTarget to dimProductCategory with foreign key constraint
- C. Direct the cube developer to use key granularity attributes.
- D. Create one product table, dimProduct, which contains product detail, category, and subcategory column
- E. Connect factSales to dimProduct with a foreign key constrain
- F. Direct the cube developer to use a non-key granularity attribute for factProductTarget.
- G. Create three product tables, dimProduct, dimProductCategory, and dimProductSubcategory, and a fourth bridge table that joins products to their appropriate category and subcategory table records with foreign key constraint
- H. Direct the cube developer to use key granularity attributes.
- I. Create three product tables, dimProduct, dimProductCategory, and dimProductSubcategory
- J. Connect factSales to all three product tables and connect factProductTarget to dimProductCategory with foreign key constraint
- K. Direct the cube developer to use key granularity attributes.

Answer: B

NEW QUESTION 120

You are reviewing the design of a student dimension table in an existing data warehouse hosted on SQL Azure.

The current dimension design does not allow the retention of historical changes to student attributes such as ParentOccupation.

You need to redesign the dimension to enable the full historical reporting of changes to multiple student attributes including ParentOccupation.

What should you do?

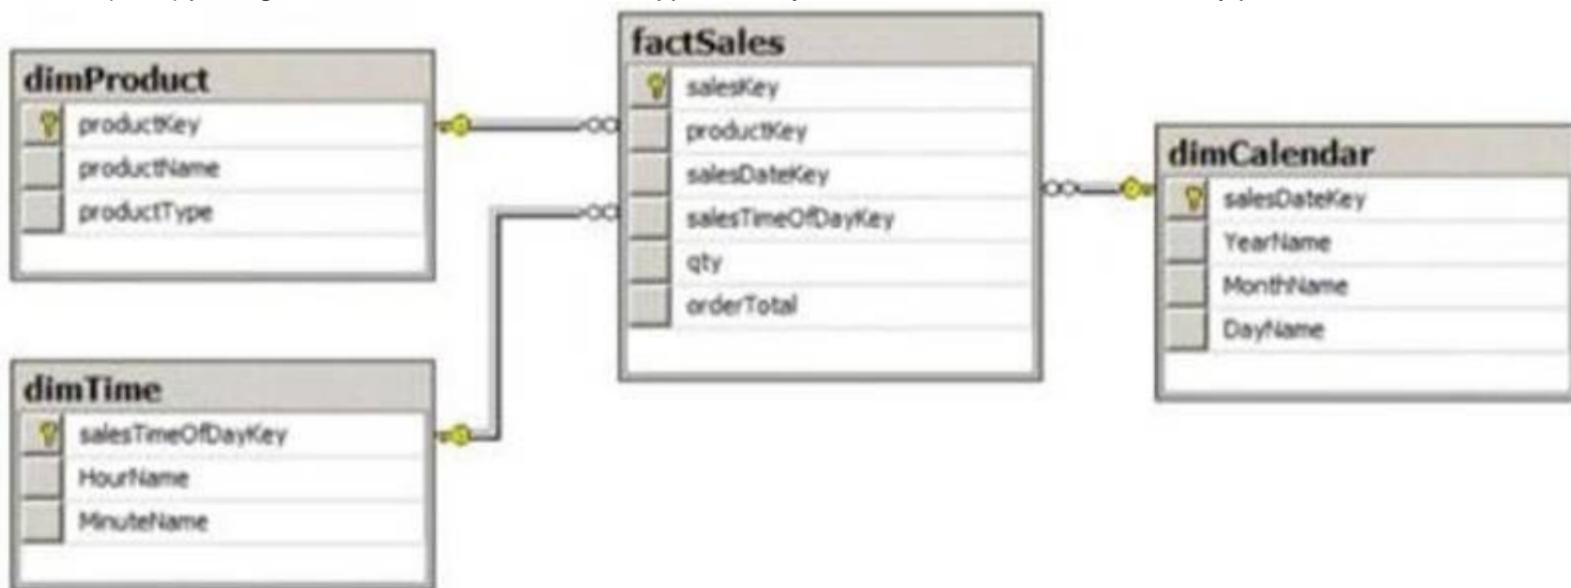
- A. Add CurrentValue and PreviousValue columns to the student dimension.
- B. Enable Snapshot Isolation on the data warehouse.
- C. Add an IsCurrent column to the student dimension.
- D. Add StartDate and EndDate columns to the student dimension.

Answer: D

Explanation: Adding a start and end date will give you this ability as when a record is inserted and given a start and end date, you'll have the ability to determine when they were active therefore giving you a retention of historical changes

NEW QUESTION 124

You are reviewing the design of an existing fact table named factSales, which is loaded incrementally from a SQL Azure database by a SQL Server Integration Services (SSIS) package each hour. The fact table has approximately 4 billion rows and is dimensioned by product, sales date, and sales time of day.



The database administrator is concerned about the rapid growth of the database and users experience poor reporting performance against this database. Reporting requirements have recently changed and the only remaining report that uses this fact table reports sales by product name, sale month, and sale year. No other reports will be created against this table.

You need to reduce the report processing time and minimize the growth of the database. What should you do?

- A. Create an indexed view over the fact table to sum orderTotal by month.
- B. Create a view over the fact table to sum orderTotal by month.
- C. Change the granularity of the fact table to month.
- D. Partition the fact table by productKe

Answer: C

NEW QUESTION 126

You are adding a new capability to several dozen SQL Server Integration Services (SSIS) packages. The new capability is not available as an SSIS task. Each package must be extended with the same new capability. You need to add the new capability to all the packages without copying the code between packages. What should you do?

- A. Use the Expression task.
- B. Use the Script task.
- C. Develop a custom task.
- D. Use the Script component.
- E. Develop a custom component.

Answer: C

Explanation: References:

<http://msdn.microsoft.com/en-us/library/ms135965.aspx> <http://msdn.microsoft.com/en-us/library/ms345161.aspx>

NEW QUESTION 128

You are creating a SQL Server Integration Services (SSIS) package to retrieve product data from two different sources. One source is hosted in a SQL Azure database. Each source contains products for different distributors. Products for each distributor source must be combined for insertion into a single product table destination. You need to select the appropriate data flow transformation to meet this requirement. Which transformation types should you use? (Each answer represents a complete solution. Choose all that apply.)

- A. Slowly Changing Dimension
- B. pivot
- C. Lookup
- D. Union All
- E. Merge

Answer: DE

NEW QUESTION 130

You are designing a SQL Server Integration Services (SSIS) data flow to load sales transactions from a source system into a data warehouse hosted on SQL Azure. One of the columns in the data source is named ProductCode. Some of the data to be loaded will reference products that need special processing logic in the data flow. You need to enable separate processing streams for a subset of rows based on the source product code. Which data flow transformation should you use?

- A. Audit
- B. Source Assistant
- C. Script Task
- D. Conditional Split

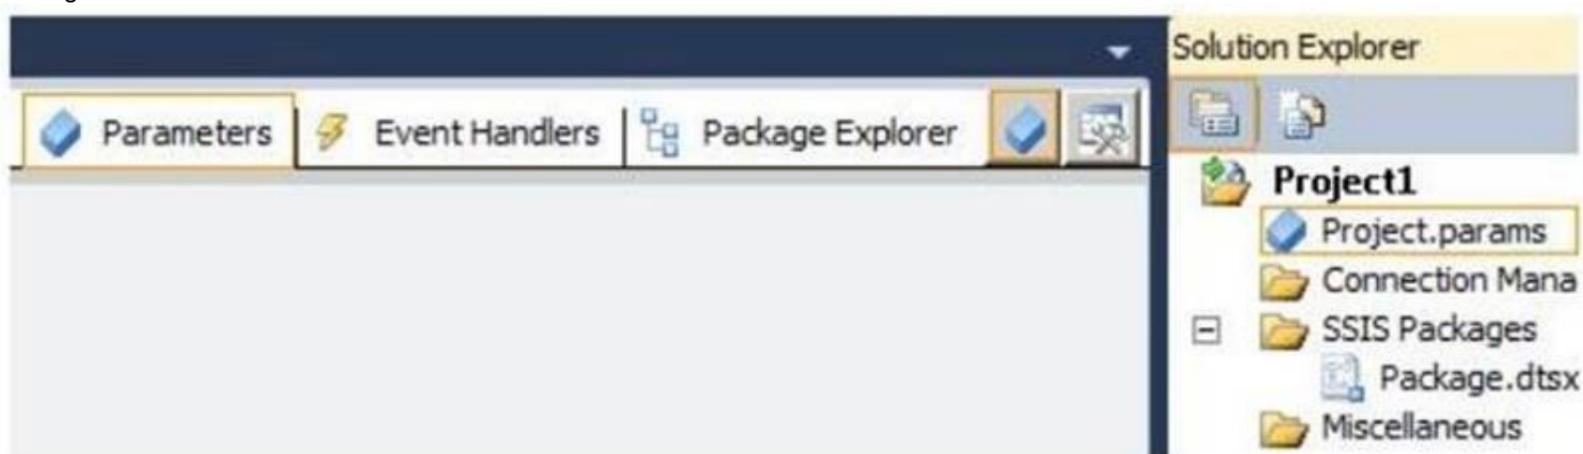
Answer: D

Explanation: <http://msdn.microsoft.com/en-us/library/ms137640.aspx> <http://msdn.microsoft.com/en-us/library/ms141150.aspx> <http://msdn.microsoft.com/en-us/library/ff929138.aspx> <http://msdn.microsoft.com/en-us/library/ff929116.aspx>

NEW QUESTION 134

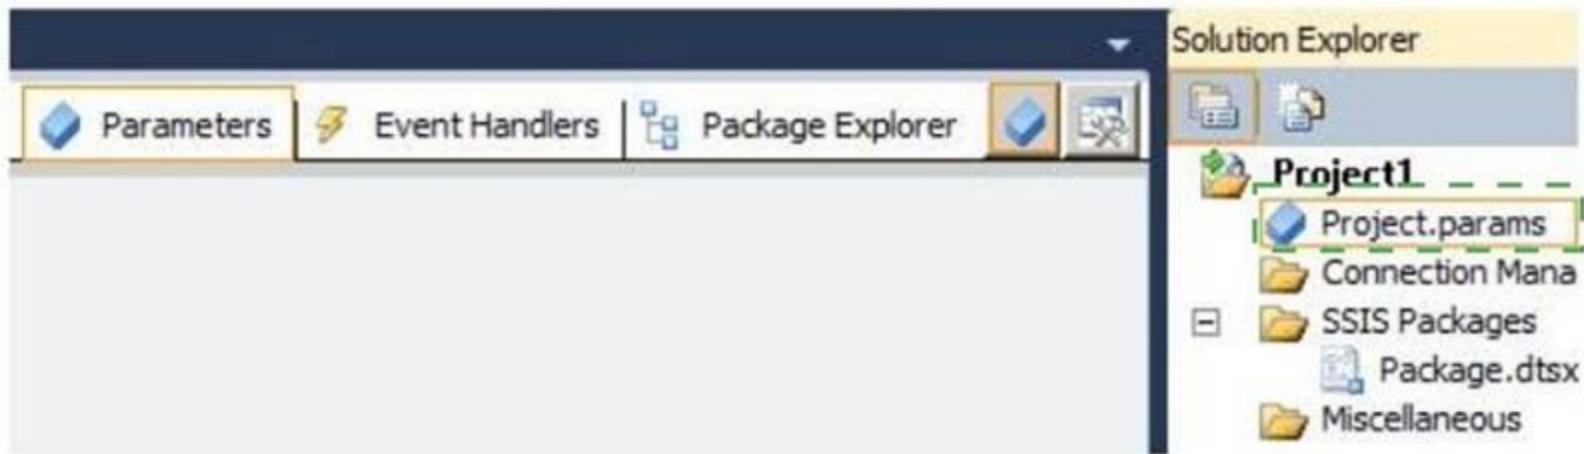
HOTSPOT

You are editing a SQL Server Integration Services (SSIS) project named Project1 in SQL Server Data Tools. A package Connection Manager has been parameterized with project scope. You need to display the parameters that have been generated for the Connection Manager. What should you use? To answer, select the appropriate setting or settings in the answer area.



Answer:

Explanation:



NEW QUESTION 136

You are troubleshooting an existing SQL Server Integration Services (SSIS) package. On several occasions, the package execution does not finish and no data seems to have been transferred. You need to ensure that package logging occurs. Your solution must minimize deployment and development efforts. What should you do?

- A. Run the dtutil command to deploy the package to the SSIS catalog and store the configuration in SQL Server.
- B. Use an msi file to deploy the package on the server.
- C. Open a command prompt and execute the package by using the SQL Log provider and running the dtexecui.exe utility.
- D. Create a reusable custom logging component and use it in the SSIS project.
- E. Open a command prompt and run the dtutil /copy command.
- F. Configure the SSIS solution to use the Project Deployment Model.
- G. Add an OnError event handler to the SSIS project.
- H. Open a command prompt and run the gacutil command.
- I. Open a command prompt and run the dtexec /dumperror /conn command.
- J. Open a command prompt and run the dtexec /rep /conn command.
- K. Configure the output of a component in the package data flow to use a data ta

Answer: C

Explanation: References:

- <http://msdn.microsoft.com/en-us/library/ms141212.aspx>
- <http://www.mssqltips.com/sqlservertip/2450/ssis-package-deployment-model-in-sql-server-2012-part-1-of-2/>
- <http://www.mssqltips.com/sqlservertip/2450/ssis-package-deployment-model-in-sql-server-2012-part-2-of-2/>

NEW QUESTION 140

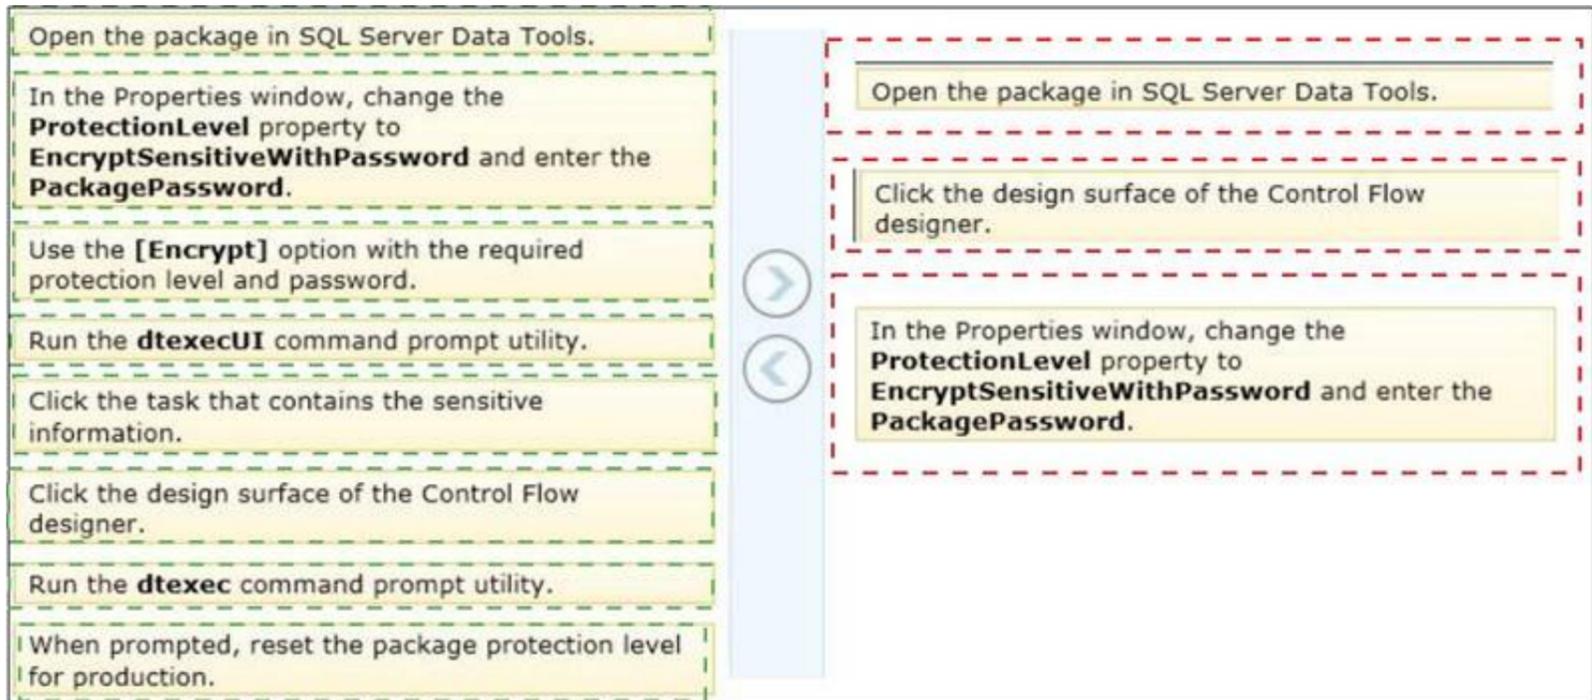
DRAG DROP

You are developing a SQL Server Integration Services (SSIS) package that is ready for deployment to a production server. The package contains sensitive information secured by using the EncryptSensitiveWithUserKey package protection level. You are preparing the package for deployment by the production operations team. You need to ensure that the production operations team can open and execute the package without re-entering the sensitive information. Which three steps should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Open the package in SQL Server Data Tools.	
In the Properties window, change the ProtectionLevel property to EncryptSensitiveWithPassword and enter the PackagePassword .	
Use the [Encrypt] option with the required protection level and password.	
Run the dtexecUI command prompt utility.	
Click the task that contains the sensitive information.	
Click the design surface of the Control Flow designer.	
Run the dtexec command prompt utility.	
When prompted, reset the package protection level for production.	

Answer:

Explanation:



NEW QUESTION 143

A SQL Server Integration Services (SSIS) package was deployed two weeks ago with the Project Deployment Model. Sometimes the package is started as part of a multistep SQL job. At other times, the package is started manually by a database administrator by using the Object Explorer in SQL Server Management Studio. You need to identify the authenticated user responsible for starting the package each time it executes. How can you find this information?

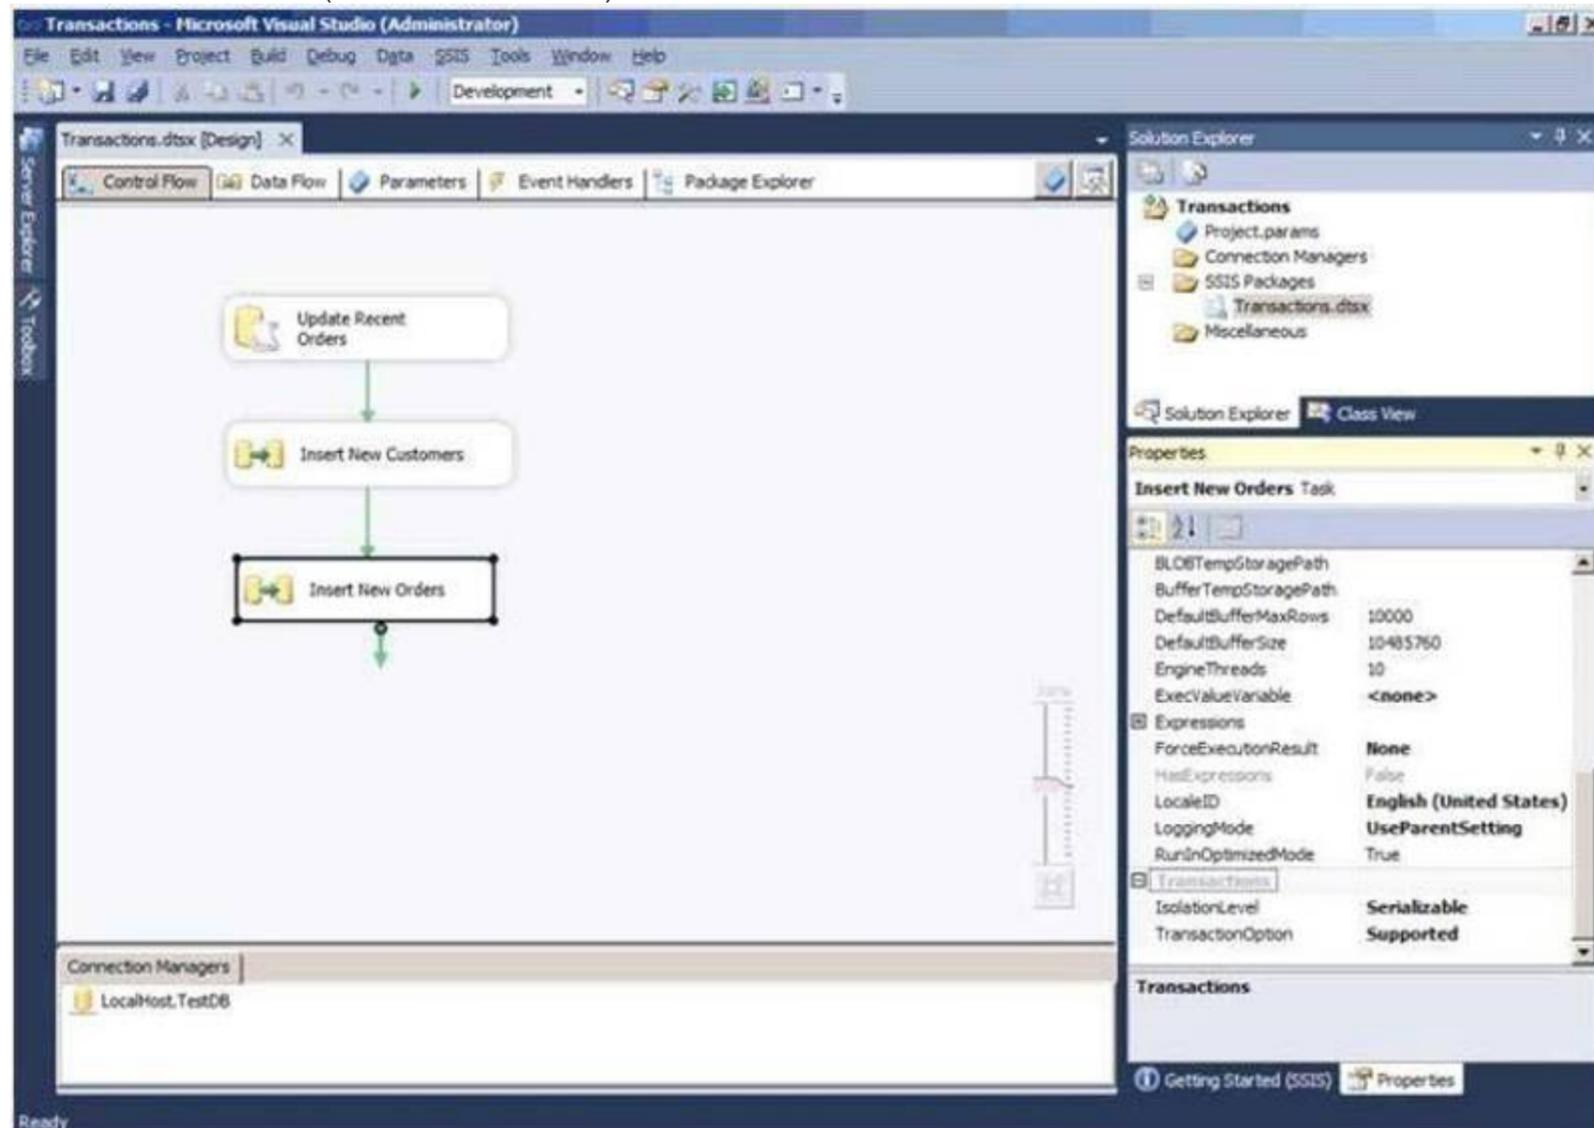
- A. In the SSISDB.[catalog], query the .[executions] view.
- B. In the SSISDB.[catalog] , query the [event_messages] view.
- C. In SQL Server Management Studio, view the SQL Agent Job History.
- D. In SQL Server Management Studio, view the SQL Agent Error Log.
- E. In SQL Server Management Studio, view the SQL Server Lo

Answer: A

NEW QUESTION 147

CORRECT TEXT

You are designing a package control flow. The package moves sales order data from a SQL Azure transactional database to an on-premise reporting database. The package will run several times a day, while new sales orders are being added to the transactional database. The current design of the package control flow is shown in the answer area. (Click the Exhibit button.)



The Insert New Orders Data Flow task must meet the following requirements:

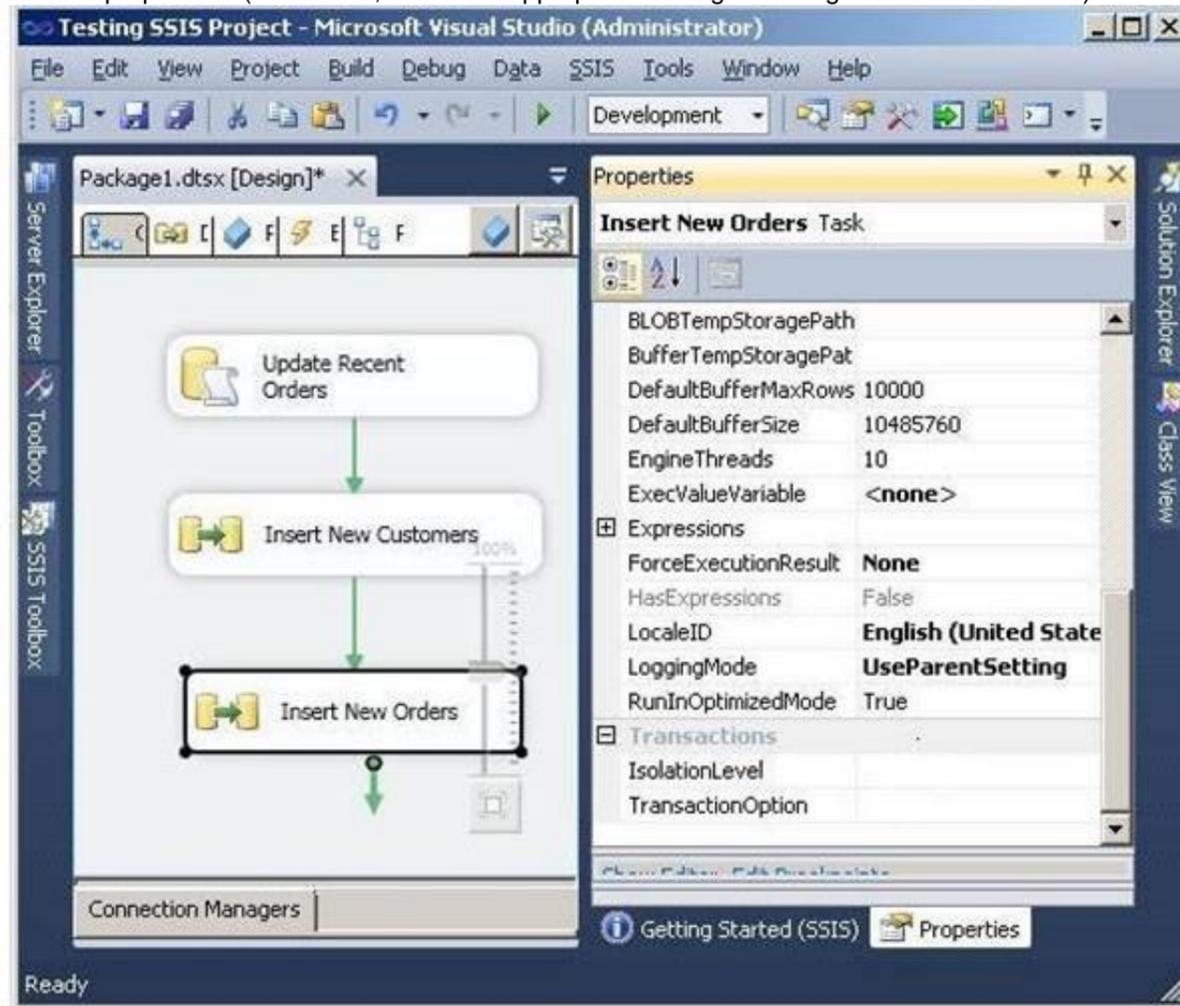
- Usage of the tempdb database should not be impacted.

- Concurrency should be maximized, while only reading committed transactions.
- If the task fails, only that task needs to be rolled back.

You need to configure the Insert New Orders Data Flow task to meet the requirements.

How should you configure the transaction properties? (To answer, select the appropriate setting or settings in the answer area.)

- Chaos
- NotSupported
- ReadCommitted
- ReadUncommitted
- RepeatableRead
- Required
- Serializable
- Snapshot
- Supported
- Unspecified



Answer:

Explanation: IsolationLevel = ReadCommitted.

TransactionOption = Required

References:

<http://msdn.microsoft.com/en-us/library/ms137690.aspx> <http://msdn.microsoft.com/en-us/library/ms137749.aspx> <http://msdn.microsoft.com/en-us/library/microsoft.sqlserver.dts.runtime.dtscontainer.isolationlevel.aspx> <http://msdn.microsoft.com/en-us/library/ms173763.aspx>

NEW QUESTION 148

You develop a SQL Server Integration Services (SSIS) package that imports SQL Azure data into a data warehouse every night.

The SQL Azure data contains many misspellings and variations of abbreviations. To import the data, a developer used the Fuzzy Lookup transformation to choose the closest-matching string from a reference table of allowed values. The number of rows in the reference table is very large.

If no acceptable match is found, the Fuzzy Lookup transformation passes a null value. The current setting for the Fuzzy Lookup similarity threshold is 0.50.

Many values are incorrectly matched.

You need to ensure that more accurate matches are made by the Fuzzy Lookup transformation without degrading performance.

What should you do?

- A. Decrease the maximum number of matches per lookup.
- B. Change the similarity threshold to 0.55.
- C. Change the Exhaustive property to True.
- D. Increase the maximum number of matches per lookup.

Answer: B

Explanation: Reference: <http://msdn.microsoft.com/en-us/library/ms137786.aspx>

NEW QUESTION 153

You are designing a data warehouse with two fact tables. The first table contains sales per month and the second table contains orders per day.

Referential integrity must be enforced declaratively.

You need to design a solution that can join a single time dimension to both fact tables. What should you do?

- A. Join the two fact tables.
- B. Merge the fact tables.
- C. Create a time dimension that can join to both fact tables at their respective granularity.
- D. Create a surrogate key for the time dimension.

Answer: C

Explanation: References:

<http://msdn.microsoft.com/en-us/library/ms174537.aspx> <http://technet.microsoft.com/en-us/library/ms174832.aspx> <http://msdn.microsoft.com/en-us/library/ms174832.aspx>

us/library/ms174884.aspx
<http://decipherinfosys.wordpress.com/2007/02/01/surrogate-keys-vs-natural-keys-for-primary-key/> <http://www.agiledata.org/essays/keys.html>
<http://www.databasejournal.com/features/mssql/article.php/3922066/SQL-Server-Natural-Key-Verses-Surrogate-Key.htm> <http://www.jamesserra.com/archive/2012/01/surrogate-keys/>

NEW QUESTION 154

You administer a Microsoft SQL Server 2012 database. The database contains a table that has the following definition:

```
CREATE TABLE [Sales].[Customer] (
    [CustomerID] int NOT NULL,
    [CustomerName] nvarchar(50) NOT NULL,
    [TerritoryID] int NULL,
    [LastContactDate] datetimeoffset NULL,
    [CustomerType] nchar(1) NOT NULL,
    [Notes] varchar(250) NULL
)
```

You want to export data from the table to a flat file by using the SQL Server Import and Export Wizard.

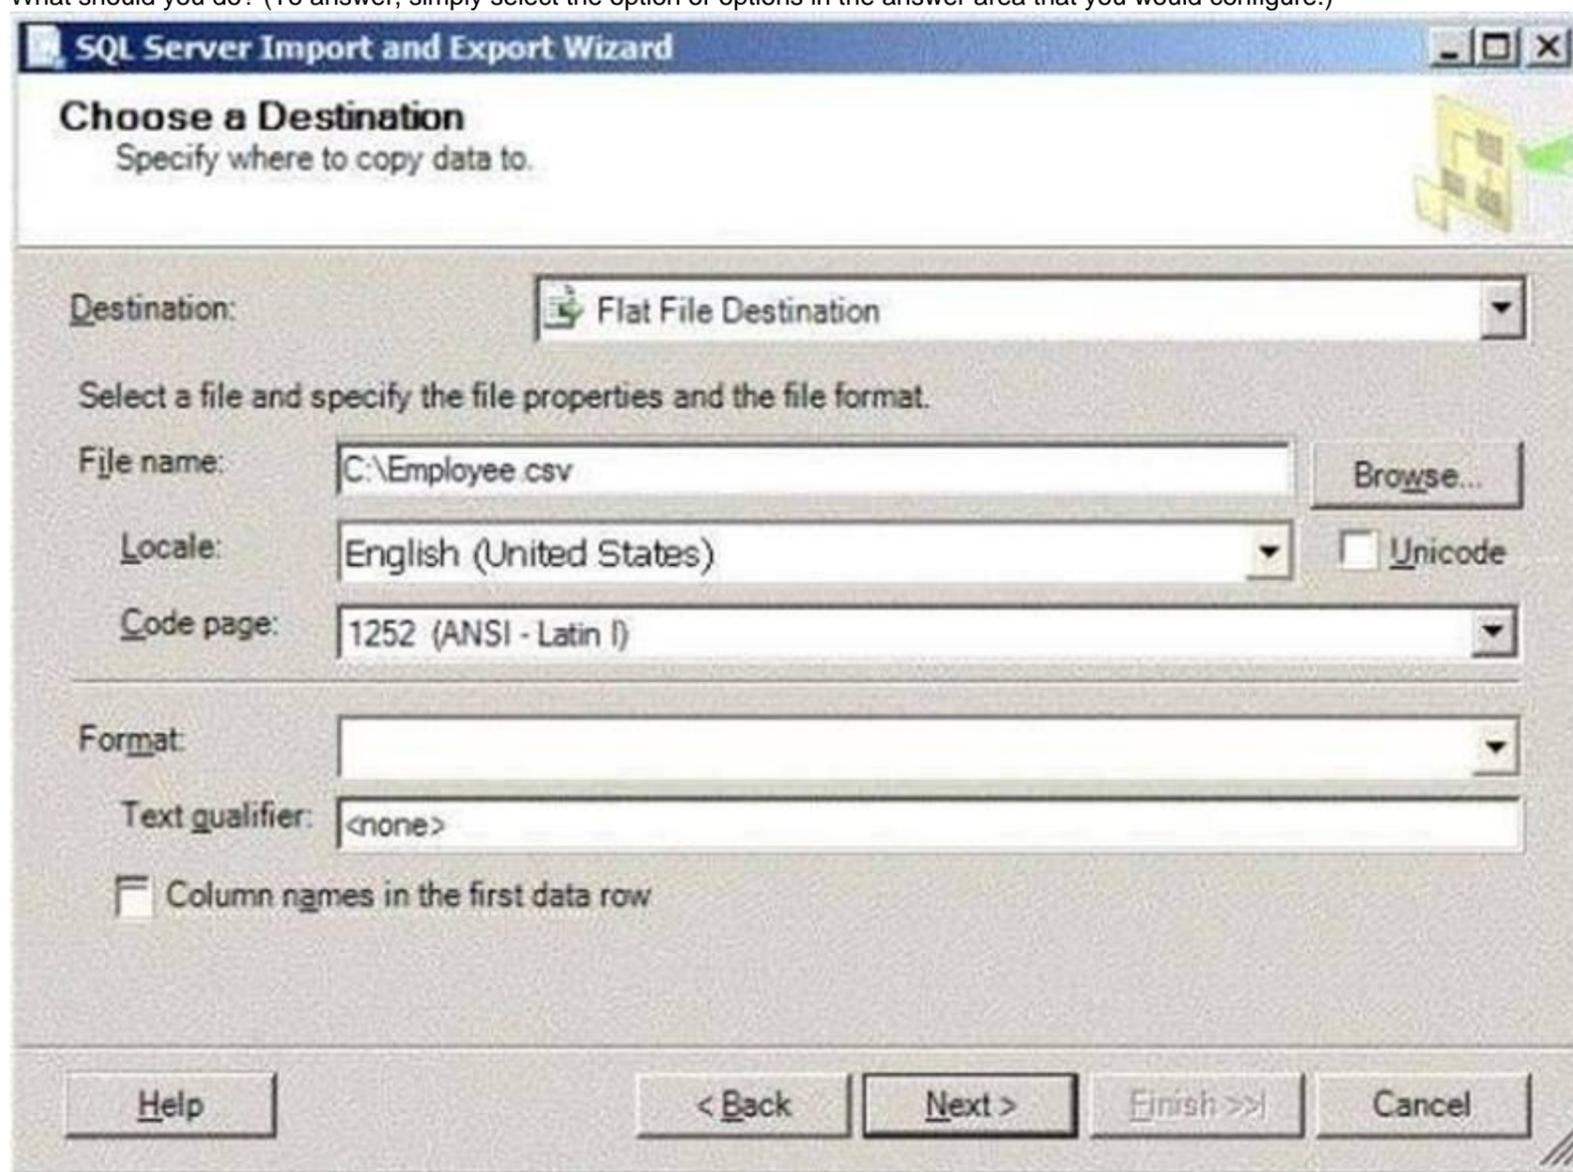
You need to ensure that the following requirements are met:

The first row of the file contains the first row of data.

Each record is of the same length. The date follows the U.S. date format.

The file supports international characters.

What should you do? (To answer, simply select the option or options in the answer area that you would configure.)

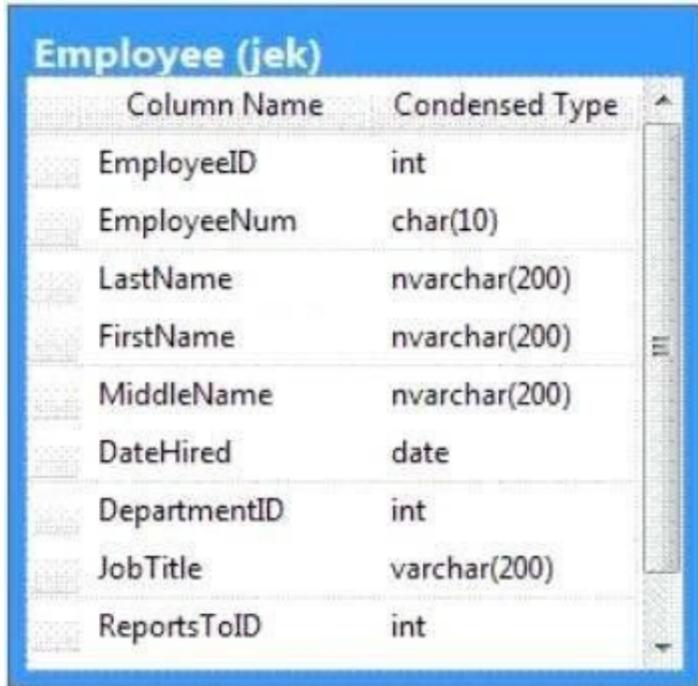


Answer:

Explanation: Each record is of the same length = Fixed Width. The file supports international characters = Unicode. The date follows the U.S. date format = Locale (already configured). The first row of the file contains the first row of data = Uncheck "Column names in the first data row"

NEW QUESTION 158

You administer a Microsoft SQL Server 2012 database. The database contains a table named Employee. Part of the Employee table is shown in the exhibit.



Column Name	Condensed Type
EmployeeID	int
EmployeeNum	char(10)
LastName	nvarchar(200)
FirstName	nvarchar(200)
MiddleName	nvarchar(200)
DateHired	date
DepartmentID	int
JobTitle	varchar(200)
ReportsToID	int

Column name	Description
EmployeeID(pk)	Uniquely identifies the employee record in the table Used throughout the database by all the other tables that reference the Employee table
EmployeeNum	An alphanumeric value calculated according to company requirements Has to be unique within the Employee table Exists only within the Employee table
DepartmentID	References another table named Department that contains data for each department in the company
ReportsToID	Contains the EmployeeID of the manager to whom an employee reports
ReportsToID	Contains the EmployeeID of the manager to whom an employee reports

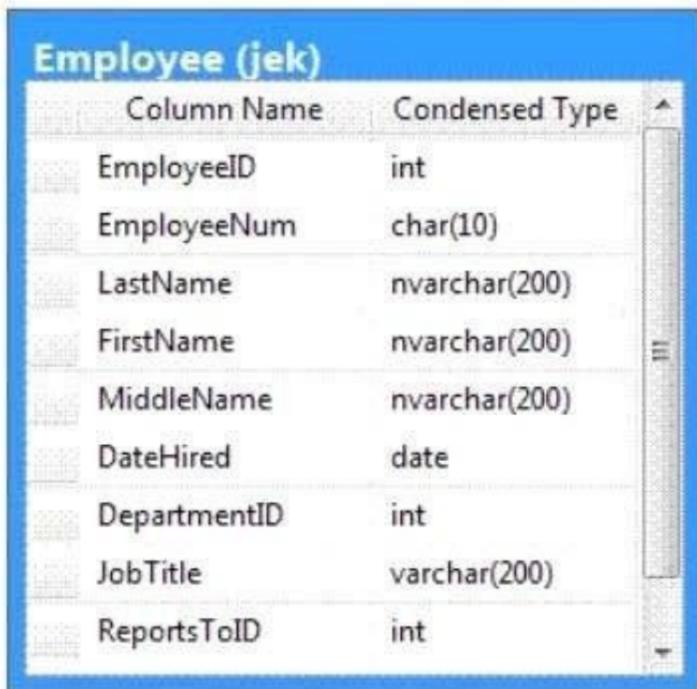
Confidential information about the employees is stored in a separate table named EmployeeDat

- A. One record exists within EmployeeData for each record in the Employee tabl
- B. You need to assign the appropriate constraints and table properties to ensure data integrity and visibilit
- C. On which column in the Employee table should you create a unique constraint?
- D. DateHired
- E. DepartmentID
- F. EmployeeID
- G. EmployeeNum
- H. FirstName
- I. JobTitle
- J. LastName
- K. MiddleName
- L. ReportsToID

Answer: D

NEW QUESTION 161

You administer a Microsoft SQL Server 2012 database. The database contains a table named Employee. Part of the Employee table is shown in the exhibit. (Click the Exhibit button.)



Column Name	Condensed Type
EmployeeID	int
EmployeeNum	char(10)
LastName	nvarchar(200)
FirstName	nvarchar(200)
MiddleName	nvarchar(200)
DateHired	date
DepartmentID	int
JobTitle	varchar(200)
ReportsToID	int

Column name	Description
EmployeeID	Uniquely identifies the employee record in the table Used throughout the database by all the other tables that reference the Employee table
EmployeeNum	An alphanumeric value calculated according to company requirements Has to be unique within the Employee table Exists only within the Employee table
DepartmentID	References another table named Department that contains data for each department in the company
ReportsToID	Contains the EmployeeID of the manager to whom an employee reports

Confidential information about the employees is stored in a separate table named EmployeeData. One record exists within EmployeeData for each record in the Employee table. You need to assign the appropriate constraints and table properties to ensure data integrity and visibility. On which column in the Employee table should you use an identity specification to include a seed of 1,000 and an increment of 1?

- A. DateHired
- B. DepartmentID
- C. EmployeeID
- D. EmployeeNum
- E. FirstName
- F. JobTitle
- G. LastName
- H. MiddleName
- I. ReportsToID

Answer: C

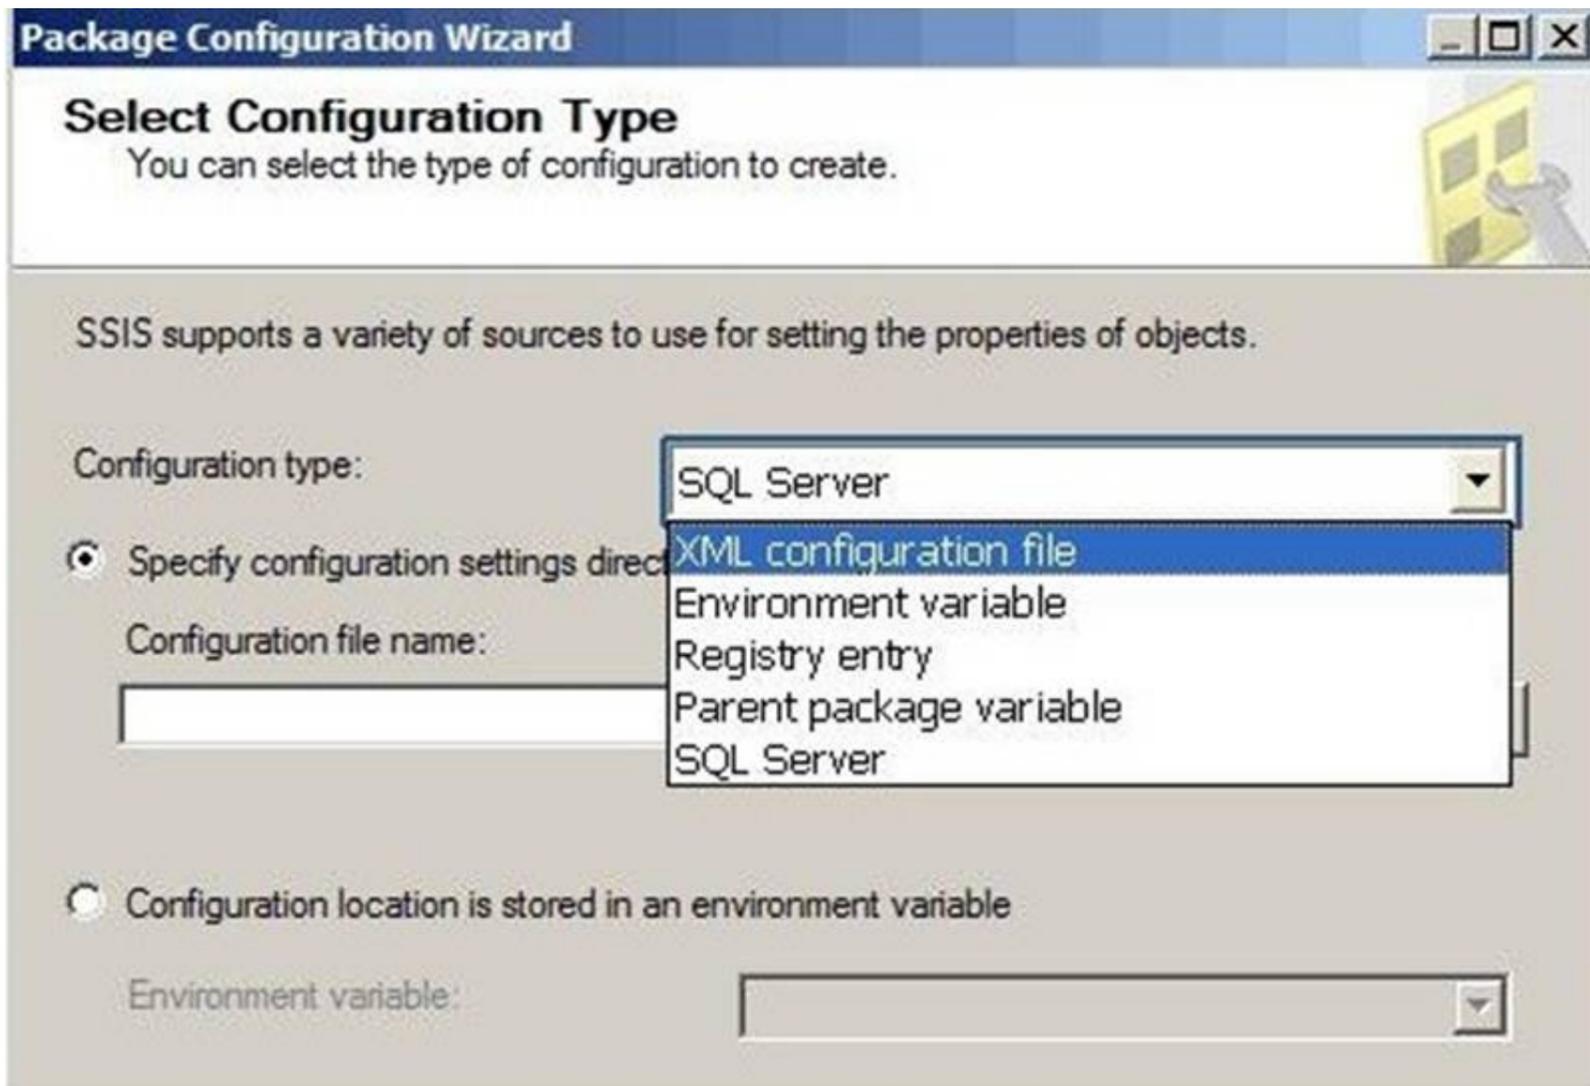
NEW QUESTION 164

HOTSPOT

You are designing a SQL Server Integration Services (SSIS) package configuration strategy. The package configuration must meet the following requirements:

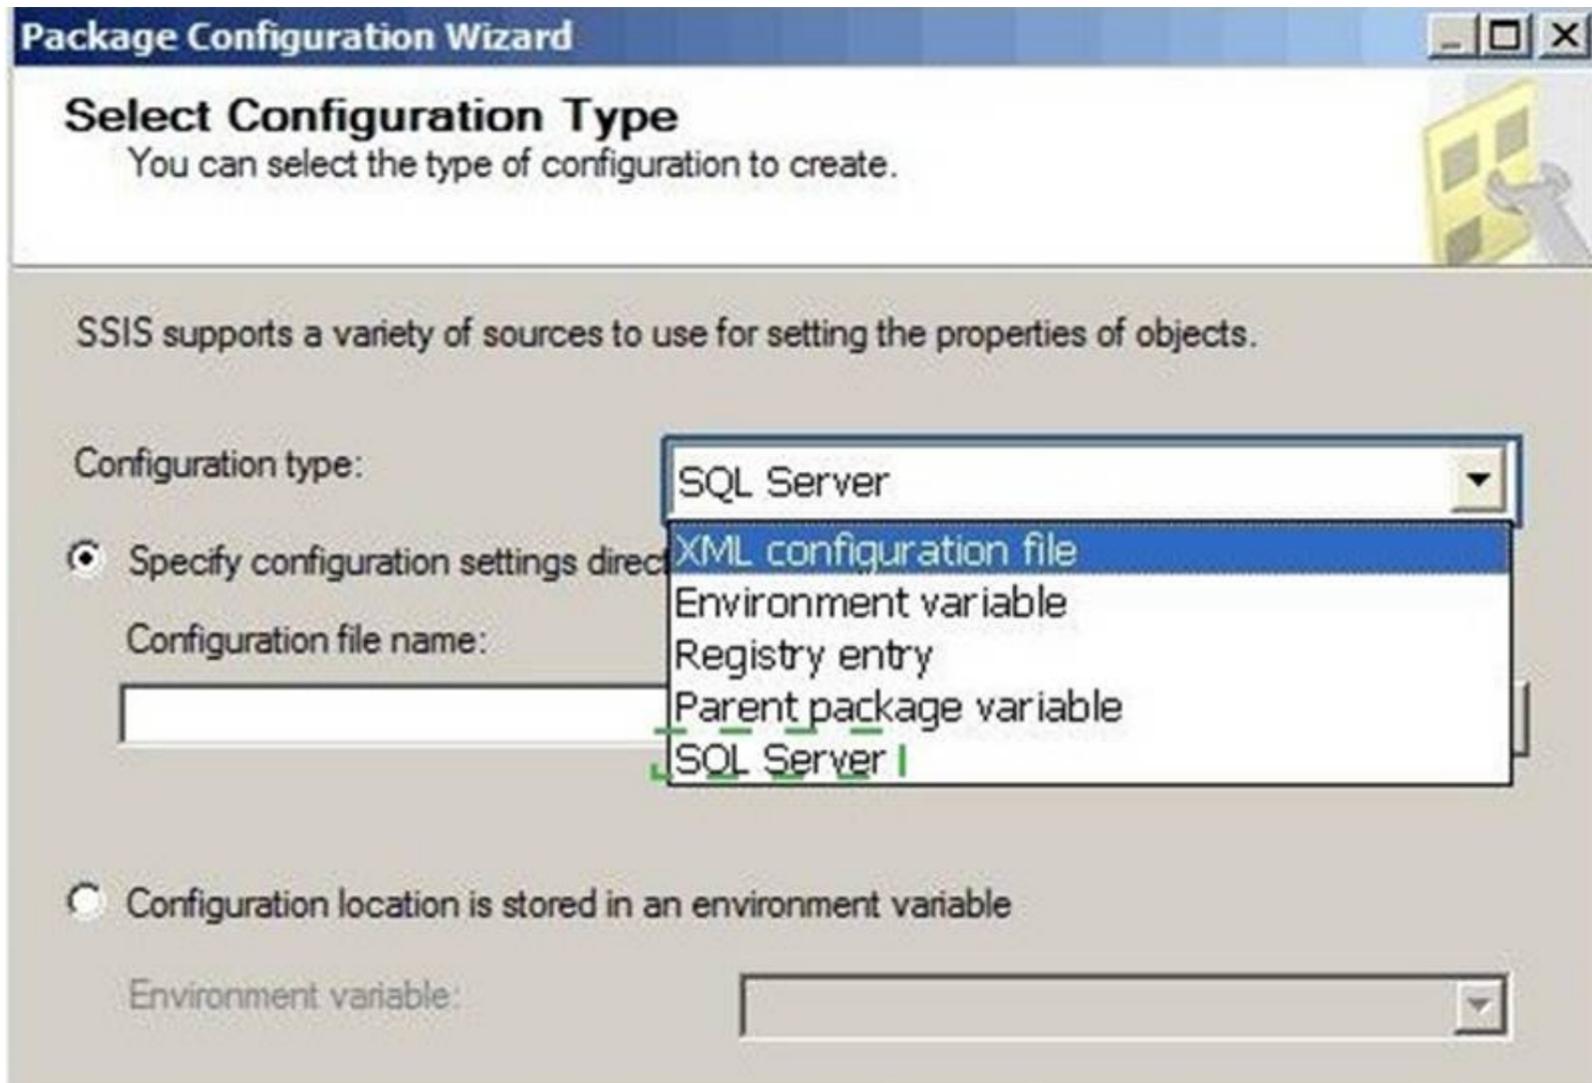
- Include multiple properties in a configuration.
- Support several packages with different configuration settings.

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



Answer:

Explanation:



NEW QUESTION 168

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

You have the following requirements:

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

- Minimize development effort.

You need to design a solution that meets the requirements.

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

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

Answer: A

NEW QUESTION 170

DRAG DROP

You plan to deploy a SQL Server Integration Services (SSIS) project by using the project deployment model.

You need to monitor control flow tasks to determine whether any of them are running longer than usual. Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Answer:

Explanation: 1. Connect to SSISDB.

2. Query against execution_component_phases and executions.

3. Execute query Ref:

<http://msdn.microsoft.com/en-us/library/hh230981.aspx> <http://msdn.microsoft.com/en-us/library/ff878089.aspx>

NEW QUESTION 172

DRAG DROP

You are validating whether a SQL Server Integration Services (SSIS) package named Master.dtsx in the SSIS catalog is executing correctly.

You need to display the number of rows in each buffer passed between each data flow component of the package.

Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Answer:

Explanation: Ref: <http://msdn.microsoft.com/en-us/library/hh230986.aspx>

NEW QUESTION 173

You are designing a fact table in a SQL Server database. The fact table must meet the following requirements:

- Include a columnstore index.
- Allow users to choose up to 10 dimension tables and up to five facts at one time.
- Maximize performance of queries that aggregate measures by using any of the 10 dimensions.
- Support billions of rows.
- Use the most efficient design strategy.

You need to design the fact table to meet the requirements. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. Design a fact table with 5 dimensional key columns and 10 measure column
- B. Place the columnstore index on the dimensional key columns.
- C. Design a fact table with 5 dimensional key columns and 10 measure column
- D. Place the columnstore index on the measure columns.
- E. Design a fact table with 10 dimensional key columns and 5 measure column
- F. Place the columnstore index on the dimensional key columns and the measure columns.
- G. Design a fact table with 10 dimensional key columns and 5 measure column
- H. Place the columnstore index on only the measure columns.

Answer: C

NEW QUESTION 174

You are designing a SQL Server Integration Services (SSIS) solution that will load multiple Online Transactional Processing (OLTP) data sources into a SQL Server data mart.

You have the following requirements:

- Ensure that the process supports the creation of an exception report that details possible duplicate key values, null ratios within columns, and column-length distributions of values.
- Ensure that users can generate the exception report in an XML format.
- Use the minimum development effort.

You need to design the SSIS solution to meet the requirements. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. Use a Data Profiling task
- B. Use a Data Flow task to extract the XML output of the Data Profiling task into a SQL Server table
- C. Query the table to view the exceptions.
- D. Use Data Flow tasks to process the clean data.
- E. Use a Data Profiling task
- F. Read the exceptions in Data Profile Viewer.
- G. Design a stored procedure that examines data for common dirty data pattern
- H. Use an Execute SQL task.

Answer: C

NEW QUESTION 177

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

You have the following requirements:

- The deployment model must support changing the content of connection strings by using parameters at execution time.
- The deployment model must automatically log events to the SSISDB database.
- Maximize performance at execution time.

You need to design a solution that meets the requirements. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

- A. Use a project deployment mode
- B. Modify connection manager properties to use project parameters.
- C. Use a package deployment mode
- D. Save each SSIS package to a file share that can be accessed from all environments.
- E. Use a package deployment mode
- F. Configure connections in an XML configuration file referenced by an environment variable that corresponds to the SQL Server environment of each SSIS package.
- G. Use a project deployment mode
- H. Modify connection manager properties to use package parameters.

Answer: A

NEW QUESTION 179

DRAG DROP

A new dedicated server is used to execute resource-intensive SQL Server Integration Services (SSIS) 2012 packages. The environment that you are deploying the packages to has the following constraints:

The operating system is Windows Server 2008 R2. The SSIS packages are stored in the SSIS catalog.

Some of these SSIS packages use 32-bit custom components.

You need to install only the components that are required to deploy and run the packages on the new server.

Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Answer:

Explanation: Note: Step 1:

Step 2:
 If you require 32-bit versions of the Integration Services runtime and tools to run certain packages in 32-bit mode, you must also install SQL Server Data Tools (SSDT).

Note:
 For a complete installation of Integration Services, together with the tools and documentation for developing and managing packages, select both Integration Services and the following Shared Features:

- * SQL Server Data Tools to install the tools for designing packages.
- * Management Tools - Complete to install SQL Server Management Studio for managing packages.
- * Client Tools SDK to install managed assemblies for Integration Services programming.

Step 3: How to Modify the Service Configuration File to Connect to a Named Instance of SQL Server To connect to a named instance of SQL Server, you must modify the service configuration file, MsDtsSrvr.ini.xml.

NEW QUESTION 181

HOTSPOT

You are developing a SQL Server Integration Service (SSIS) package. The package loads a customer dimension table by using a data flow task.

Changes to the customer attributes must be tracked over time.

You need to produce a checksum value to identify the rows that have changed since the last Extract, Transform and Load (ETL) process execution. You need to use the least amount of development effort to achieve this goal.

Which transformation should you use? (To answer, select the appropriate transformation in the answer area.)

▲ Favorites

-  Cache Transform
-  CDC Splitter
-  Character Map
-  Data Conversion
-  Import Column
-  Script Component

Answer:

Explanation:

▲ Favorites

-  Cache Transform
-  CDC Splitter
-  Character Map
-  Data Conversion
-  Import Column
-  Script Component

NEW QUESTION 185

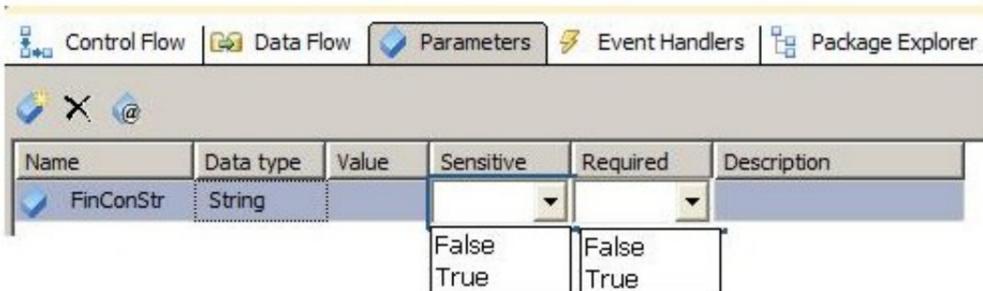
HOTSPOT

A SQL Server Integration Services (SSIS) package is designed to download data from a financial database hosted in SQL Azure.

The connection string to the financial database is defined as a project parameter named FinConStr. The parameter value must be stored securely and must be set explicitly every time the package is executed.

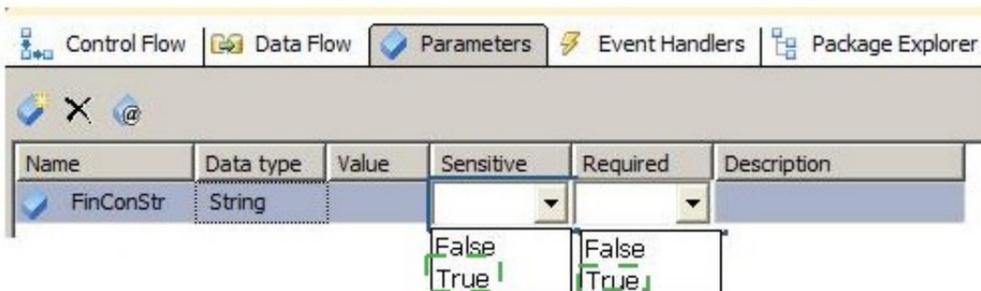
You need to configure the parameter to meet the requirements.

What should you do? (To answer, configure the appropriate option or options in the dialog box in the answer area.)



Answer:

Explanation:



NEW QUESTION 190

HOTSPOT

You are the Master Data Services (MDS) administrator at your company.

An existing user must be denied access to a certain hierarchy node for an existing model. You need to configure the user's permissions.

Which user management menu item should you select? (To answer, configure the appropriate option or options in the dialog box in the answer area.)



Answer:

Explanation:



NEW QUESTION 195

DRAG DROP

You are loading a dataset into SQL Server. The dataset contains numerous duplicates for the Artist and Song columns.

The values in the Artist column in the dataset must exactly match the values in the Artist domain in the knowledge base. The values in the Song column in the dataset can be a close match with the values in the Song domain.

You need to use SQL Server Data Quality Services (DQS) to define a matching policy rule to identify duplicates.

How should you configure the Rule Editor? (To answer, drag the appropriate answers to the answer area.)

0
50
100

Prerequisite Choices

No
Yes

Rule Editor Show overlapping clusters ▾

Domain	Similarity	Weight	Prerequisite
Artist ▾	Similar ▾		
Song ▾	Similar ▾		

Answer:

Explanation:

0
50
100

Prerequisite Choices

No
Yes

Rule Editor Show overlapping clusters ▾

Domain	Similarity	Weight	Prerequisite
Artist ▾	Similar ▾	0	Yes
Song ▾	Similar ▾	100	No

NEW QUESTION 200

HOTSPOT

You are developing a SQL Server Integration Services (SSIS) package. An Execute SQL task in the package checks product stock levels and sets a package variable named InStock to TRUE or FALSE depending on the stock level found.

After the successful execution of the Execute SQL task, one of two data flow tasks must run, depending on the value of the InStock variable.

You need to set the precedence constraints.

Which value for the evaluation operation should you use? (To answer, select the appropriate option for the evaluation operation in the answer area.)

Precedence Constraint Editor ✕

A precedence constraint defines the workflow between two executables. The precedence constraint can be based on a combination of the execution results and the evaluation of expressions.

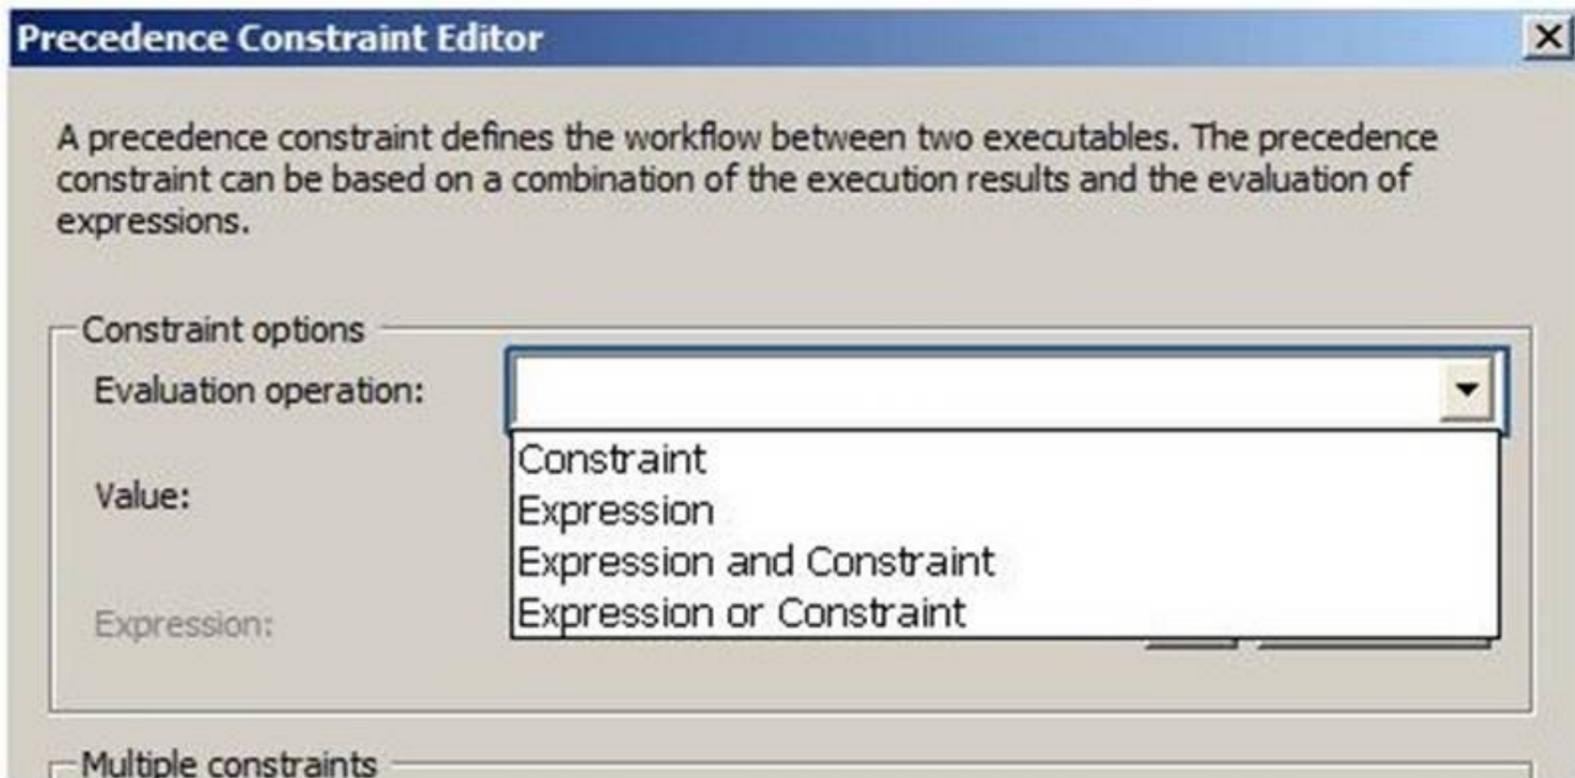
Constraint options

Evaluation operation:

Value:

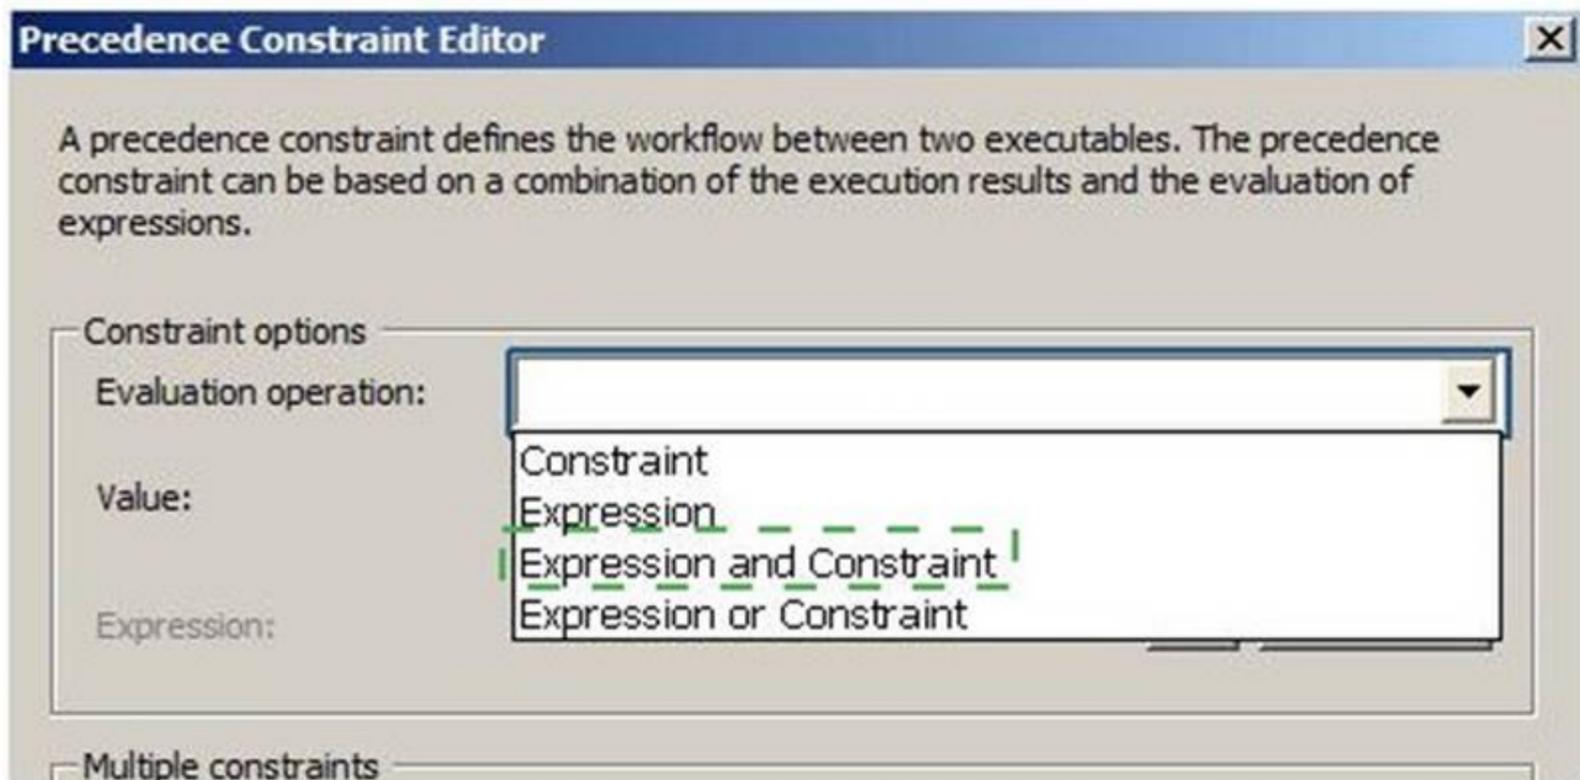
Expression:

Multiple constraints



Answer:

Explanation:



NEW QUESTION 203

You are deploying a new SQL Server Integration Services (SSIS) package to several servers. The package must meet the following requirements:

- .NET Common Language Runtime (CLR) integration in SQL Server must not be enabled.
- The Connection Managers used in the package must be configurable without editing the package.
- The deployment procedure must be automated as much as possible. You need to set up a deployment strategy that meets the requirements. What should you do?

- A. Use the gacutil command.
- B. Use the dtutil /copy command.
- C. Use the Project Deployment Wizard.
- D. Create an OnError event handler.
- E. Create a reusable custom logging component.
- F. Run the package by using the dtexec /rep /conn command.
- G. Run the package by using the dtexec /dumperror /conn command.
- H. Run the package by using the dtexecui.exe utility and the SQL Log provider.
- I. Add a data tap on the output of a component in the package data flow.
- J. Deploy the package by using an msi file.
- K. Deploy the package to the Integration Services catalog by using dtutil and use SQL Server to store the configuration.

Answer: B

NEW QUESTION 206

You maintain a SQL Server Integration Services (SSIS) package. The package was developed by using SQL Server 2008 Business Intelligence Development Studio (BIDS).

The package includes custom scripts that must be upgraded. You need to upgrade the package to SQL Server 2012. Which tool should you use?

- A. SQL Server dtexec utility (dtexec.exe)
- B. SQL Server DTEXecUI utility (dtexecui.exe)
- C. SSIS Upgrade Wizard in SQL Server Data Tools
- D. SQL Server Integration Services Deployment Wizard

Answer: C

Explanation: Use the SSIS Package Upgrade Wizard to upgrade SQL Server 2005 Integration Services (SSIS) packages and SQL Server 2008 Integration Services (SSIS) packages to the package format for the current (2012) release of SQL Server Integration Services.

NEW QUESTION 208

You are developing a SQL Server Integration Services (SSIS) package to implement an incremental data load strategy. The package reads data from a source system that uses the SQL Server change data capture (CDC) feature.

You have added a CDC Source component to the data flow to read changed data from the source system.

You need to add a data flow transformation to redirect rows for separate processing of insert, update, and delete operations.

Which data flow transformation should you use?

- A. Audit
- B. Merge Join
- C. Merge
- D. CDC Splitter

Answer: D

Explanation: The CDC splitter splits a single flow of change rows from a CDC source data flow into different data flows for Insert, Update and Delete operations
Ref: <http://msdn.microsoft.com/en-us/library/hh758656.aspx>

NEW QUESTION 212

A SQL Server Integration Services (SSIS) package imports daily transactions from several files into a SQL Server table named Transaction. Each file corresponds to a different store and is imported in parallel with the other files. The data flow tasks use OLE DB destinations in fast load data access mode.

The number of daily transactions per store can be very large and is growing. The Transaction table does not have any indexes.

You need to minimize the package execution time. What should you do?

- A. Partition the table by day and store.
- B. Create a clustered index on the Transaction table.
- C. Run the package in Performance mode.
- D. Increase the value of the Row per Batch property.

Answer: D

Explanation: * Data Access Mode – This setting provides the 'fast load' option which internally uses a BULK INSERT statement for uploading data into the destination table instead of a simple INSERT statement (for each single row) as in the case for other options.

* BULK INSERT parameters include: ROWS_PER_BATCH =rows_per_batch

Indicates the approximate number of rows of data in the data file.

By default, all the data in the data file is sent to the server as a single transaction, and the number of rows in the batch is unknown to the query optimizer. If you specify ROWS_PER_BATCH (with a value

> 0) the server uses this value to optimize the bulk-import operation. The value specified for ROWS_PER_BATCH should approximately the same as the actual number of rows.

NEW QUESTION 213

You develop a SQL Server Integration Services (SSIS) package in a project by using the Project Deployment Model. It is regularly executed within a multi-step SQL Server Agent job.

You make changes to the package that should improve performance.

You need to establish if there is a trend in the durations of the next 10 successful executions of the package. You need to use the least amount of administrative effort to achieve this goal.

What should you do?

- A. After 10 executions, in SQL Server Management Studio, view the Execution Performance subsection of the All Executions report for the package.
- B. Configure the package to send you an email upon completion that includes information about the duration of the package.
- C. After 10 executions, view the emails.
- D. Enable logging to the Application Event Log in the package control flow for the OnInformation even
- E. After 10 executions, view the Application Event Log.
- F. Enable logging to the Application Event Log in the package control flow for the OnPostExecute even
- G. After 10 executions, view the Application Event Log.

Answer: A

Explanation: The All Executions Report displays a summary of all Integration Services executions that have been performed on the server. There can be multiple executions of the sample package. Unlike the Integration Services Dashboard report, you can configure the All Executions report to show executions that have started during a range of dates. The dates can span multiple days, months, or years.

The report displays the following sections of information.

* Filter

Shows the current filter applied to the report, such as the Start time range.

* Execution Information

Shows the start time, end time, and duration for each package execution. You can view a list of the parameter values that were used with a package execution,

such as values that were passed to a child package using the Execute Package task.

NEW QUESTION 218

You are developing a SQL Server Integration Services (SSIS) package.

The package sources data from an HTML web page that lists product stock levels.

You need to implement a data flow task that reads the product stock levels from the HTML web page. Which data flow source should you use?

- A. Raw File source
- B. XML source
- C. Custom source component
- D. Flat File source

Answer: C

NEW QUESTION 220

You manage a SQL Server Master Data Services (MDS) environment.

A new application requires access to the product data that is available in the MDS repository. You need to design a solution that gives the application access to the product data with the least amount of development effort.

What should you do?

- A. Use `sp_addlinkedserver` to add a linked server to access the MDS database tables directly.
- B. Create a Subscription View in MDS.
- C. Use transactional replication for data synchronization.
- D. Create an OLE DB connection string that sets the Provider property to MD

Answer: B

Explanation: You create a subscription view to export Master Data Services data to a subscribing system such as SQL Server. You use the subscribing system to view the data in the Master Data Services database. References:

<https://docs.microsoft.com/en-us/sql/master-data-services/overview-exporting-data-master-dataservices?view=sql-server-2017>

NEW QUESTION 223

You are completing the installation of the Data Quality Server component of SQL Server Data Quality Services (DQS). You need to complete the post-installation configuration.

What should you do?

- A. Install the Analysis Services OLE DB Provider.
- B. Run the `DQSInstaller.exe` command.
- C. Run the Configuration component in the Data Quality Client.
- D. Make the data available for DQS operation

Answer: B

NEW QUESTION 226

You are developing a SQL Server Integration Services (SSIS) project by using the Project Deployment Model.

The project will be deployed to an SSIS catalog folder where Environments have already been created. You need to deploy the project. What should you do?

- A. Use an event handler for `OnError` for the package.
- B. Use an event handler for `OnError` for each data flow task.
- C. Use an event handler for `OnTaskFailed` for the package.
- D. View the job history for the SQL Server Agent job.
- E. View the All Messages subsection of the All Executions report for the package.
- F. Store the `System::SourceID` variable in the custom log table.
- G. Store the `System::ServerExecutionID` variable in the custom log table.
- H. Store the `System::ExecutionInstanceGUID` variable in the custom log table.
- I. Enable the SSIS log provider for SQL Server for `OnError` in the package control flow.
- J. Enable the SSIS log provider for SQL Server for `OnTaskFailed` in the package control flow.
- K. Deploy the project by using `dtutil.exe` with the `/COPY DTS` option.
- L. Deploy the project by using `dtutil.exe` with the `/COPY SQL` option.
- M. Deploy the `.ispac` file by using the Integration Services Deployment Wizard.
- N. Create a SQL Server Agent job to execute the `SSISDB.catalog.validate_project` stored procedure.
- O. Create a SQL Server Agent job to execute the `SSISDB.catalog.validate_package` stored procedure.
- P. Create a SQL Server Agent job to execute the `SSISDB.catalog.create_execution` and `SSISDB.catalog.start_execution` stored procedures.
- Q. Create a table to store error informatio
- R. Create an error output on each data flow destination that writes `OnError` event text to the table.
- S. Create a table to store error informatio
- T. Create an error output on each data flow destination that writes `OnTaskFailed` event text to the table.

Answer: M

NEW QUESTION 228

You are developing a SQL Server Integration Services (SSIS) project to read and write data from a Windows Azure SQL Database database to a server that runs SQL Server 2012.

The connection will be used by data flow tasks in multiple SSIS packages. The address of the target Windows Azure SQL Database database will be provided by a project parameter.

You need to create a solution to meet the requirements by using the least amount of administrative effort and maximizing data flow performance.

What should you do?

- A. Use an SSIS Script task that uses the custom assembly to parse the text data when inserting it.
- B. Use an SSIS Script transformation that uses the custom assembly to parse the text data when inserting it.
- C. Create a SQL Common Language Runtime (SQLCLR) function that uses the custom assembly to parse the text data, deploy it in the Windows Azure SQL Database database, and use it when inserting data.
- D. Create a SQL Common Language Runtime (SQLCLR) stored procedure that uses the custom assembly to parse the text data, deploy it in the Windows Azure SQL Database database, and use it when inserting data.

Answer: A

NEW QUESTION 232

You develop a SQL Server Integration Services (SSIS) project by using the Package Deployment Model. A package in the project extracts data from a Windows Azure SQL Database database. The package is deployed to SQL Server.

The package is not producing the desired results.

You need to generate the .mdmp and .tmp debug files in order to troubleshoot the issues. What should you do?

- A. Execute the catalog.add_data_tap stored procedure with the package execution_id.
- B. Execute the catalog.create_execution_dump stored procedure with the package execution_id.
- C. Run the DTEXEC utility with the /DumpOnError option.
- D. Run the DTEXEC utility with the /Reporting V optio

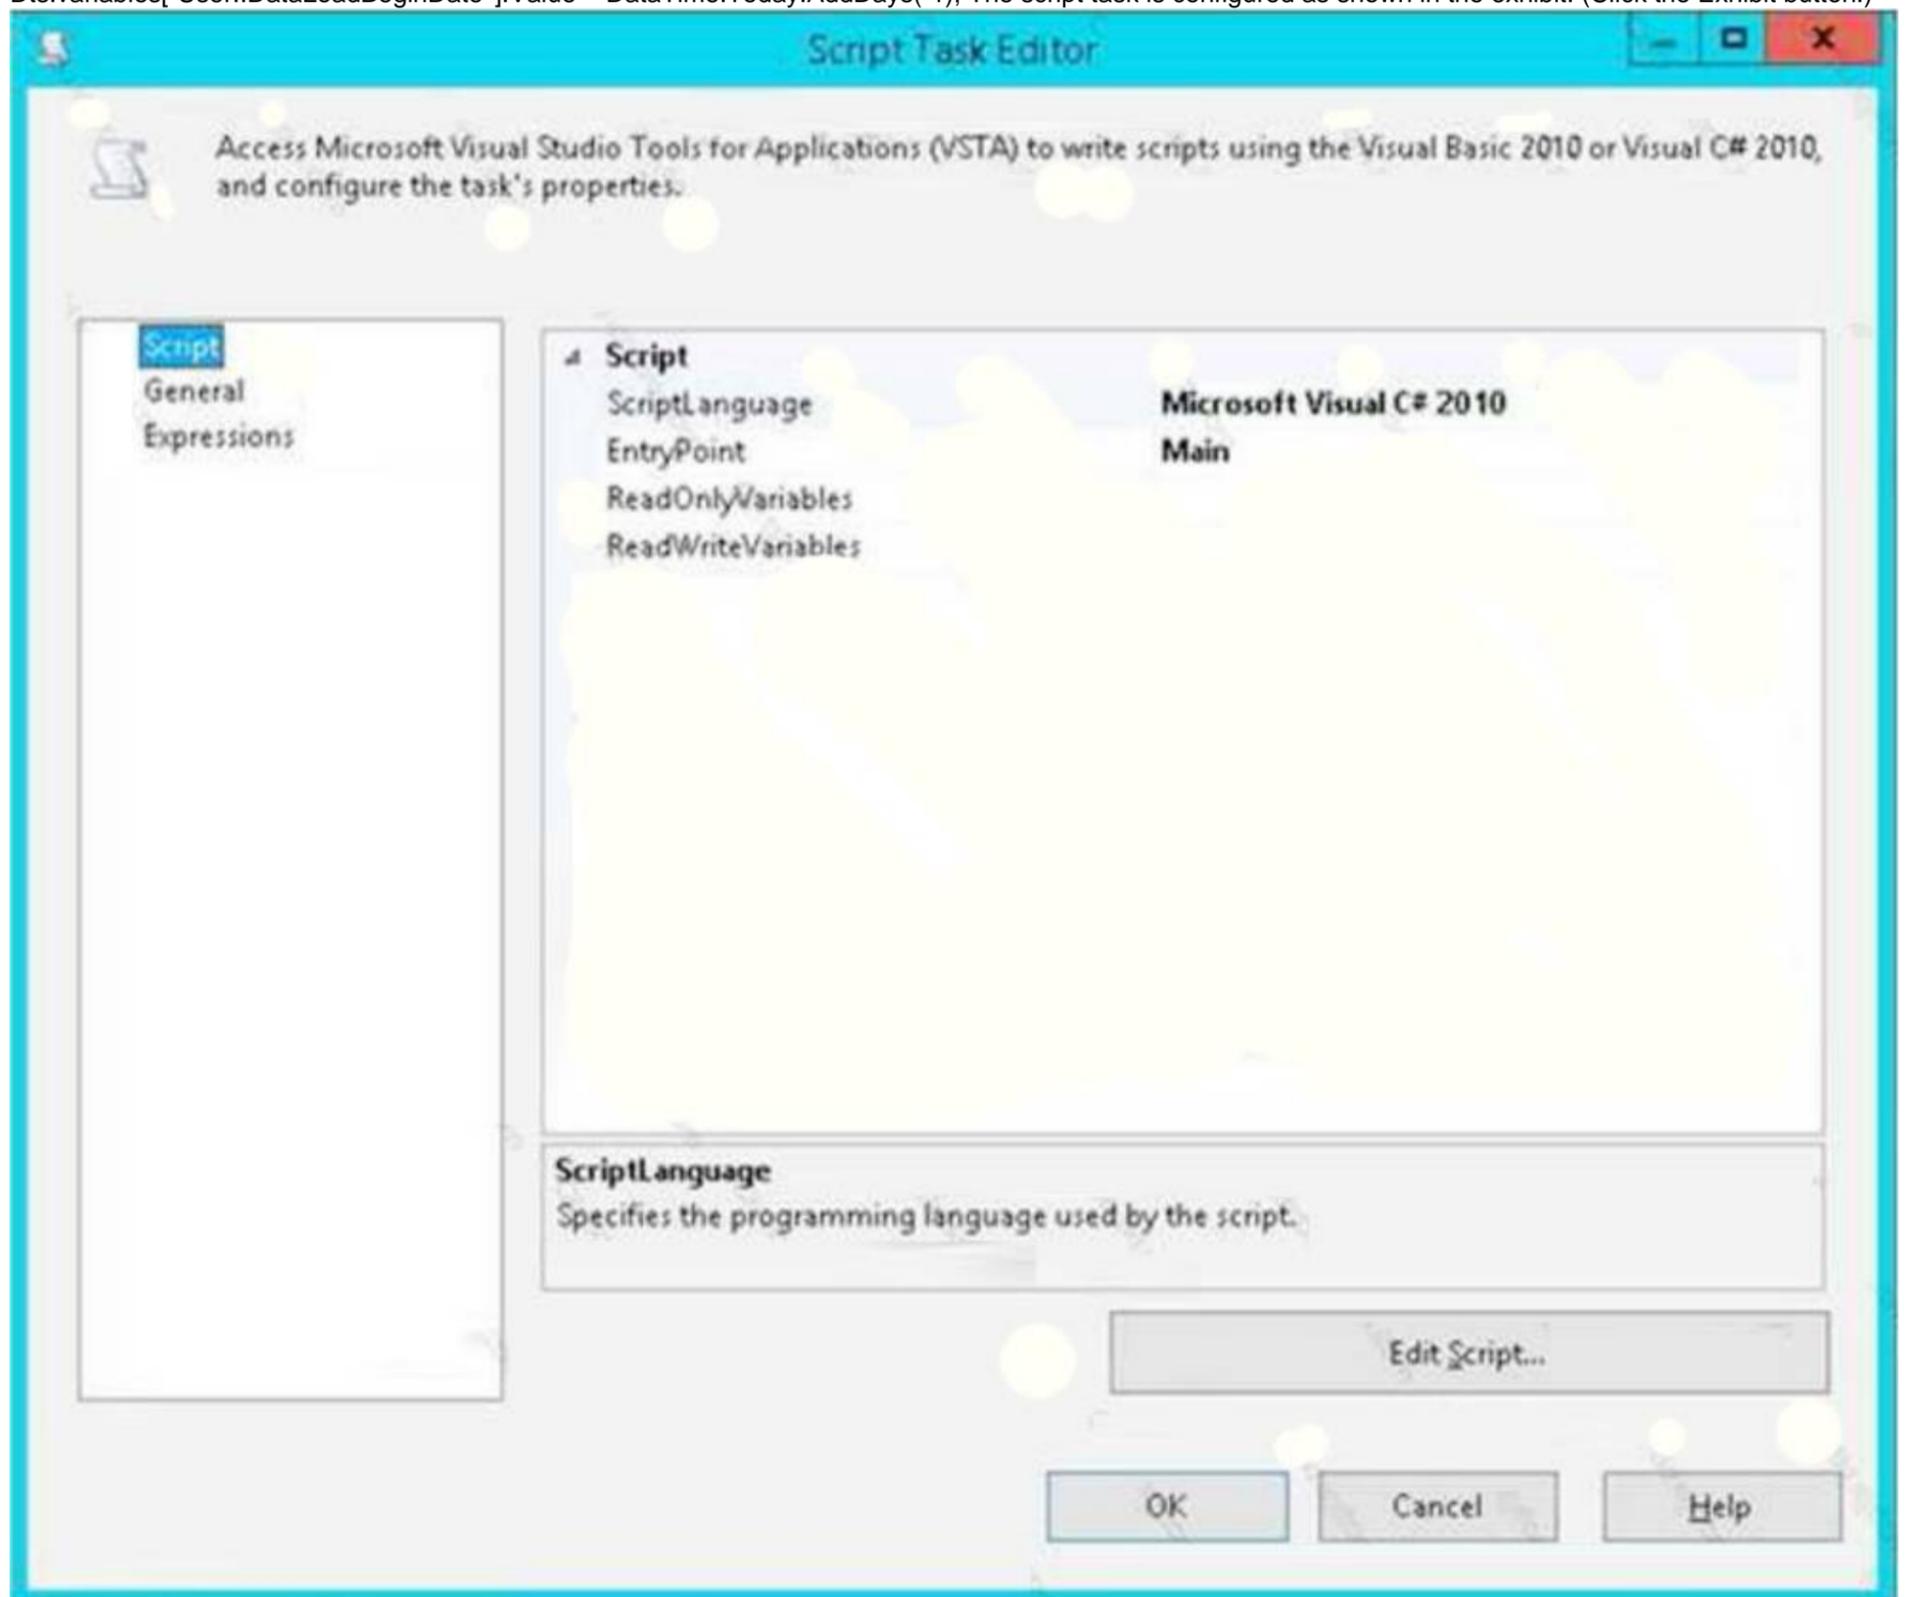
Answer: C

NEW QUESTION 237

You have a SQL Server Integration Services (SSIS) package. The package contains a script task that has the following comment.

// Update DataLoadBeginDate variable to the beginning of yesterday The script has the following code.

Dts.variables["User::DataLoadBeginDate"].Value = DateTime.Today.AddDays(-1); The script task is configured as shown in the exhibit. (Click the Exhibit button.)



When you attempt to execute the package, the package fails and returns the following error message: "Error: Exception has been thrown by the target of an invocation." You need to execute the package successfully. What should you do?

- A. Add the dataLoadBegin Date variable to the ReadOnlyVariables property.
- B. Add the DataLoadBeginDate variable to the ReadWriteVariables property.
- C. Modify the entry point of the script.
- D. Change the scope of the DataLoadBeginDate variable to Packag

Answer: B

Explanation: You add existing variables to the ReadOnlyVariables and ReadWriteVariables lists in the Script Task Editor to make them available to the custom script. Within the script, you access variables of both types through the Variables property of the Dts object.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/extending-packages-scripting/task/usingvariables-in-the-script-task?view=sql-server-2017>

NEW QUESTION 240

You are developing a SQL Server Integration Services (SSIS) package.

The package sources data from an HTML web page that lists product stock levels.

You need to implement a data flow task that reads the product stock levels from the HTML web page. Which data flow source should you use?

- A. Raw File source
- B. XML source
- C. Custom source component
- D. Flat File source

Answer: C

NEW QUESTION 243

You are deploying a new SQL Server Integration Services (SSIS) package to five servers. The package must meet the following requirements:

- .NET Common Language Runtime (CLR) integration in SQL Server must not be enabled.
- The Connection Managers used in the package must be configurable without editing and redeploying the package.
- The deployment procedure must be automated as much as possible.
- Performance must be maximized.

You need to set up a deployment strategy that meets the requirements. What should you do?

- A. Add an OnError event handler to the SSIS project.
- B. Use an msi file to deploy the package on the server.
- C. Open a command prompt and run the gacutil command.
- D. Open a command prompt and run the dtutil /copy command.
- E. Open a command prompt and run the dtexec /rep /conn command.
- F. Open a command prompt and run the dtexec /dumperror /conn command.
- G. Open a command prompt and execute the package by using the SQL Log provider and running the dtexecui.exe utility.
- H. Create a reusable custom logging component and use it in the SSIS project.
- I. Configure the SSIS solution to use the Project Deployment Model.
- J. Configure the output of a component in the package data flow to use a data tap.
- K. Run the dtutil command to deploy the package to the SSIS catalog and store the configuration inSQL Serve

Answer: D

NEW QUESTION 245

HOTSPOT

You are the data steward at your company.

Duplicate customers exist in a Microsoft Excel workbook. You create a Data Quality Services (DQS) knowledge base and matching policy to identify these duplicate customers.

You need to identify the duplicate customers.

Which option should you use? (To answer, select the appropriate option in the answer area.)



Answer:

Explanation:



NEW QUESTION 249

HOTSPOT

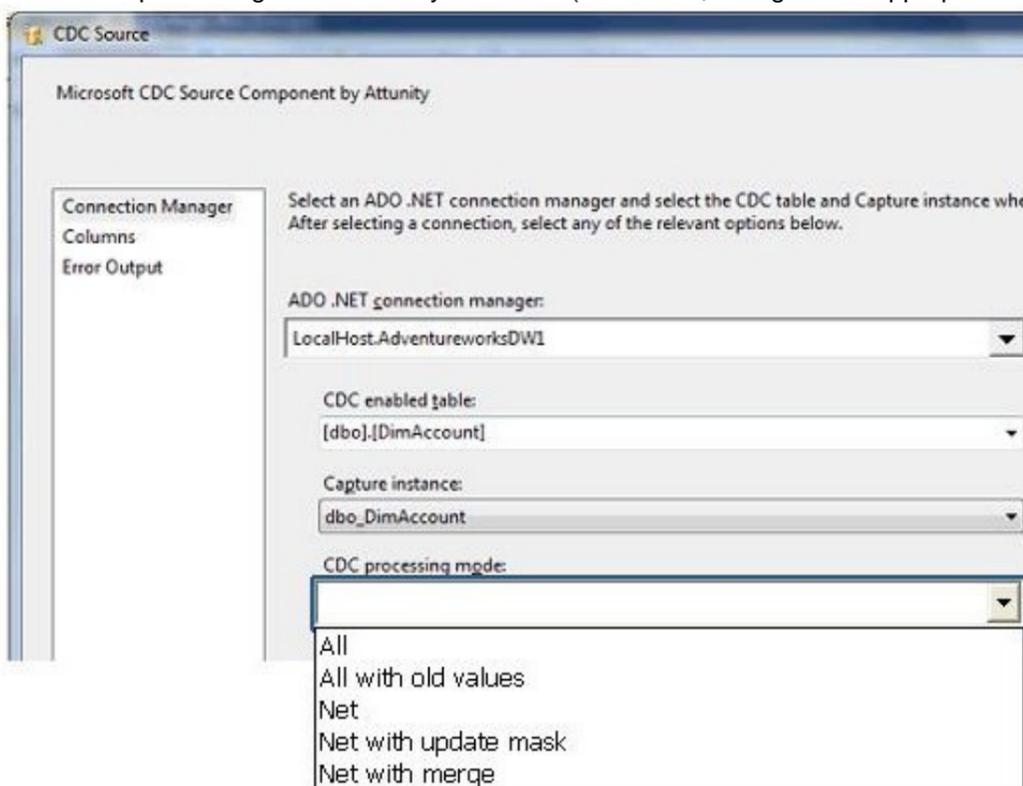
You are developing a SQL Server Integration Services (SSIS) package. The data source for the data flow task is a table that has been configured as a change data capture (CDC) table. You are using a CDC Source component to obtain the CDC data.

The CDC Source component has the following requirements:

The output must include metadata columns that indicate which source columns have changed. The output must return only one change row per source row that is modified in the current CDC processing range.

You need to configure the CDC Source component.

Which CDC processing mode should you select? (To answer, configure the appropriate option or options in the dialog box in the answer area.)



Answer:

Explanation: Ref: <http://msdn.microsoft.com/en-gb/library/hh231004.aspx>

NEW QUESTION 251

DRAG DROP

You are developing a SQL Server Integration Services (SSIS) package.

The package contains an ADO object source variable that holds a result set that was returned by a stored procedure execution.

You need to add and configure a container that will execute several tasks for each row in the ADO object source variable.

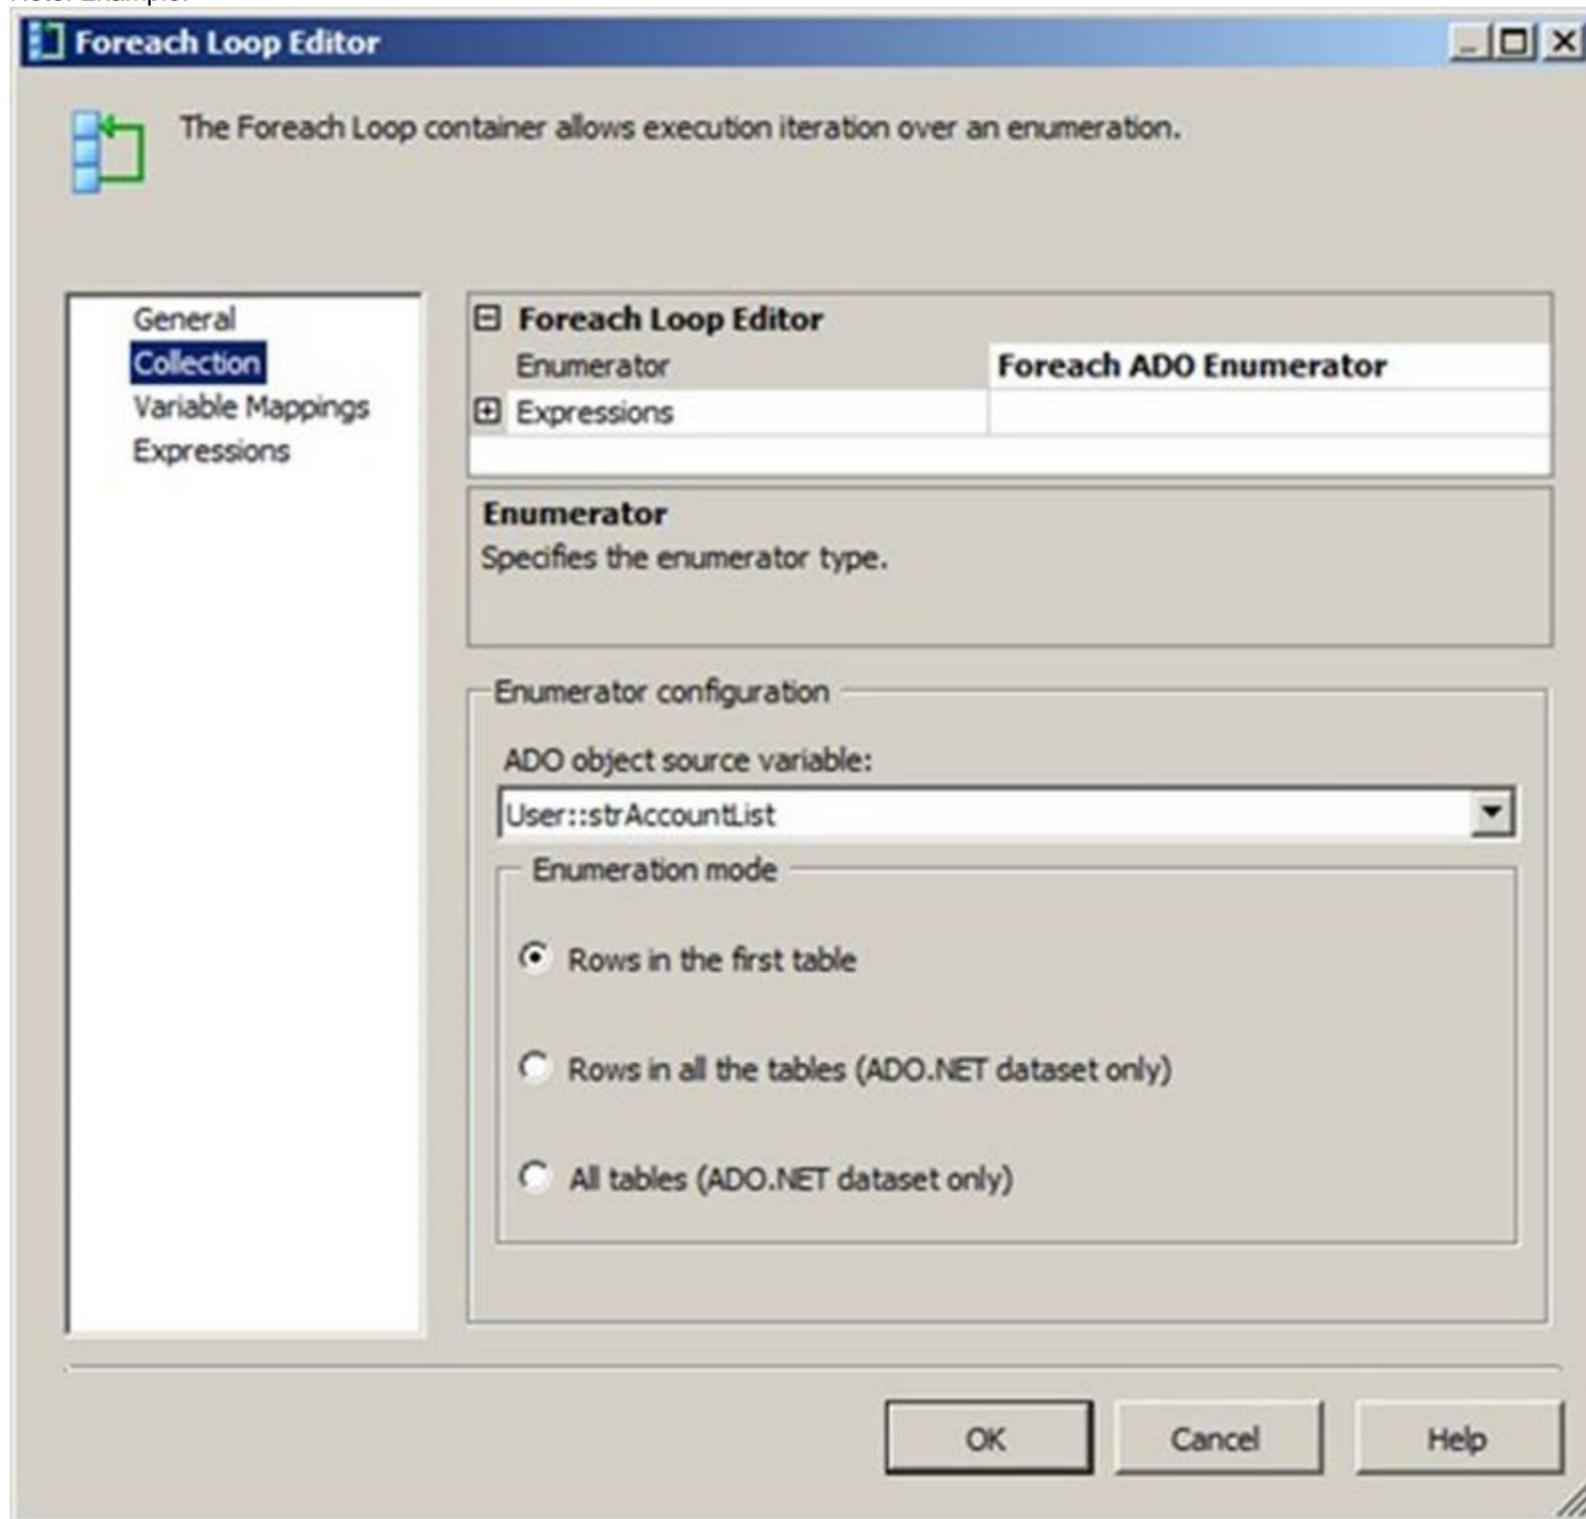
Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

- Add and edit a Foreach Loop container.
- Open the Data Flow designer of the package.
- Set the enumerator property to **Foreach NodeList Enumerator** and select the variable that contains the ADO object source.
- Open the Control Flow designer of the package.
- Add and edit a Sequence container.
- Set the enumerator property to **Foreach ADO Enumerator** and select the variable that contains the ADO object source.



Answer:

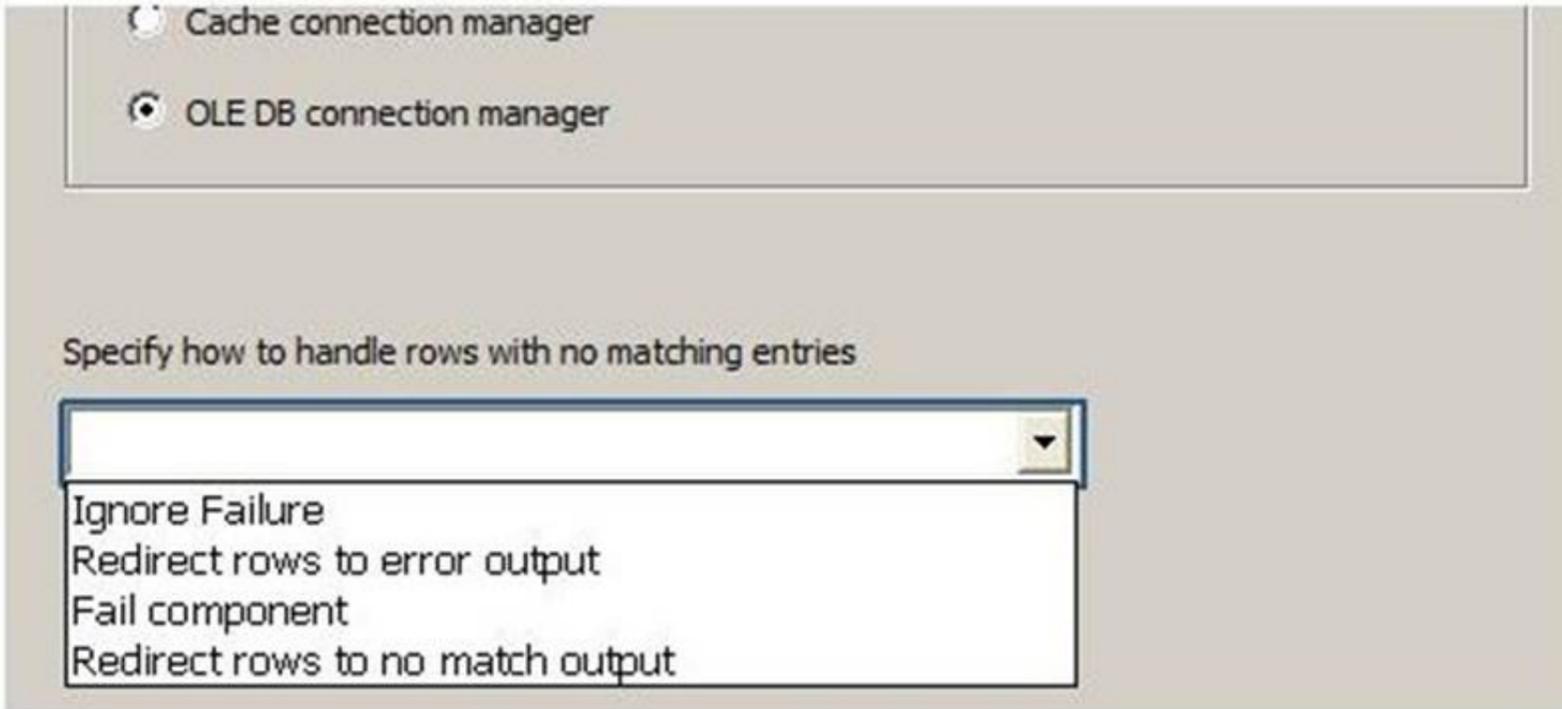
Explanation: Box 1: Open the Control Flow designer of the package. Box 2: Add and edit a Foreach Loop container.
 Box 3: Set the enumerator property to Foreach ADO Enumerator and select the variable that contains the ADO object source.
 Note: Example:



NEW QUESTION 252

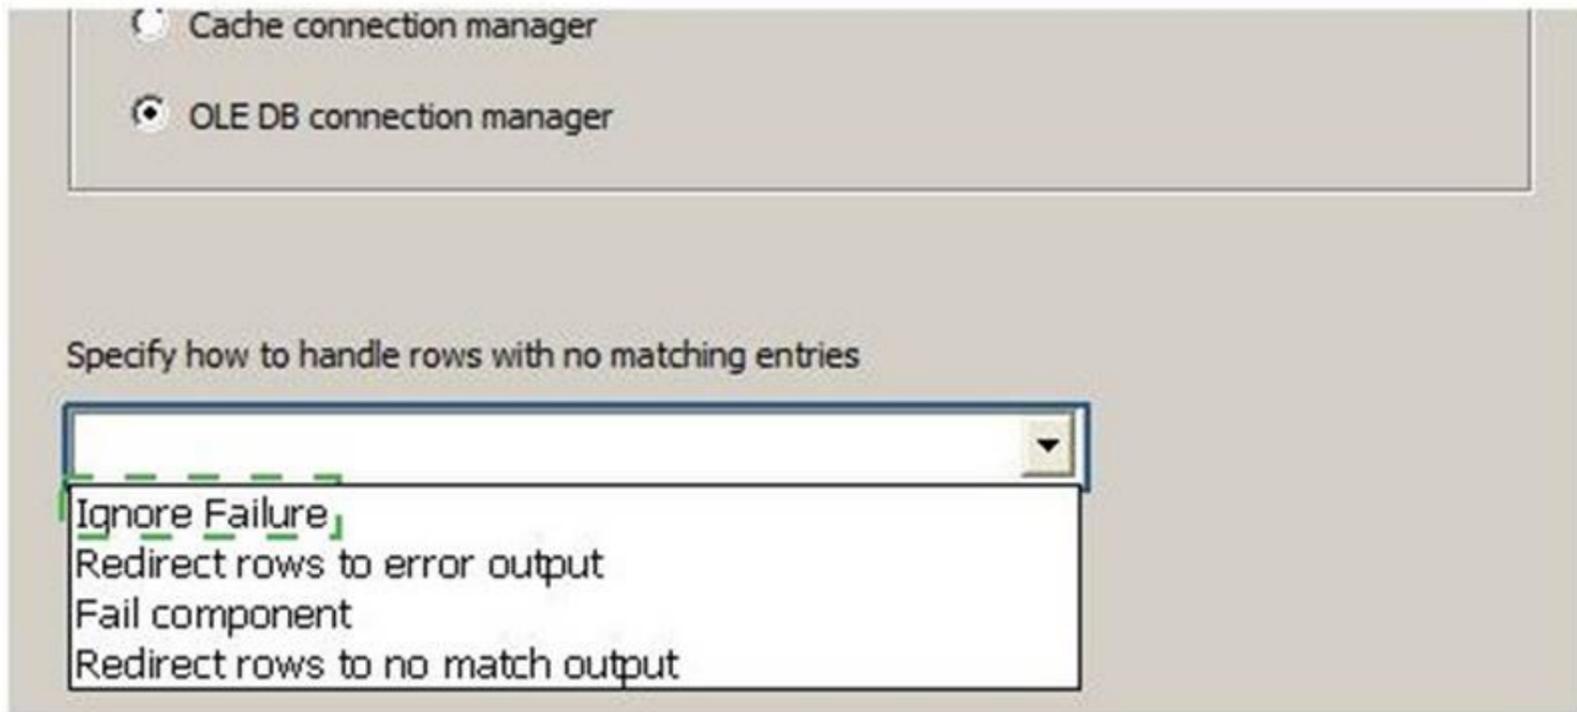
HOTSPOT

You are developing a SQL Server Integration Services (SSIS) package to implement an incremental data load strategy. The package reads rows from a source system and compares them to rows in a destination system. New rows will be inserted and changed rows will be updated. You have used a Lookup transformation and a Conditional Split transformation. The Lookup transformation joins the source and destination table on the business key, and includes all columns from the destination table in the data flow output. The Conditional Split transformation inspects the destination columns and directs data flow to either insert new records or update existing records. You need to configure the Lookup transformation to ensure that all records flow to the Conditional Split transformation, regardless of whether the rows match an existing row in the destination table. Which setting should you select? (To answer, select the appropriate option in the answer area.)



Answer:

Explanation:



NEW QUESTION 256

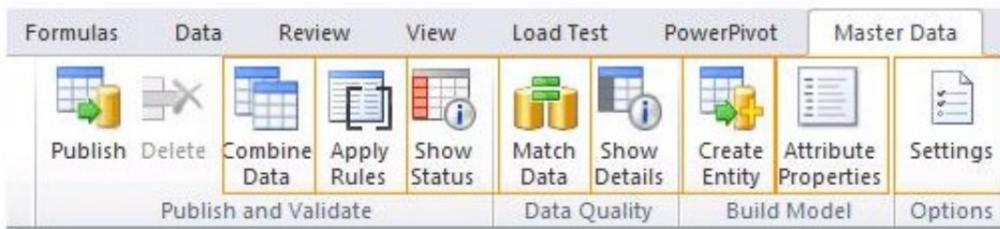
HOTSPOT

You are using the Master Data Services (MDS) Add-in for Excel to configure the entities in a model. The model consists of two entities: one named Customer and one named State. You open the Customer entity.

Currently, data stewards can enter any text value in the Customer entity's State attribute. You must restrict the entry of values in the State attribute to members defined in the State entity.

You need to configure the State attribute of the Customer entity.

Which option should you use? (To answer, select the appropriate area in the answer area.)



Answer:

Explanation:



NEW QUESTION 260

You develop a SQL Server Integration Services (SSIS) package that imports Windows Azure SQL Database data into a data warehouse every night. The Windows Azure SQL Database data contains many misspellings and variations of abbreviations. To import the data, a developer used the Fuzzy Lookup transformation to choose the closest matching string from a reference table of allowed values. The number of rows in the reference table is very large.

If no acceptable match is found, the Fuzzy Lookup transformation passes a null value. The current setting for the Fuzzy Lookup similarity threshold is 0.50. Many values are incorrectly matched.

You need to ensure that more accurate matches are made by the Fuzzy Lookup transformation without degrading performance. What should you do?

- A. Change the similarity threshold to 0.40.
- B. Decrease the maximum number of matches per lookup.
- C. Change the similarity threshold to 0.85.
- D. Increase the maximum number of matches per lookup.

Answer: C

Explanation: * Similarity threshold

Set the similarity threshold at the component level by using the slider. The closer the value is to 1, the closer the resemblance of the lookup value to the source value must be to qualify as a match. Increasing the threshold can improve the speed of matching since fewer candidate records need to be considered.

NEW QUESTION 262

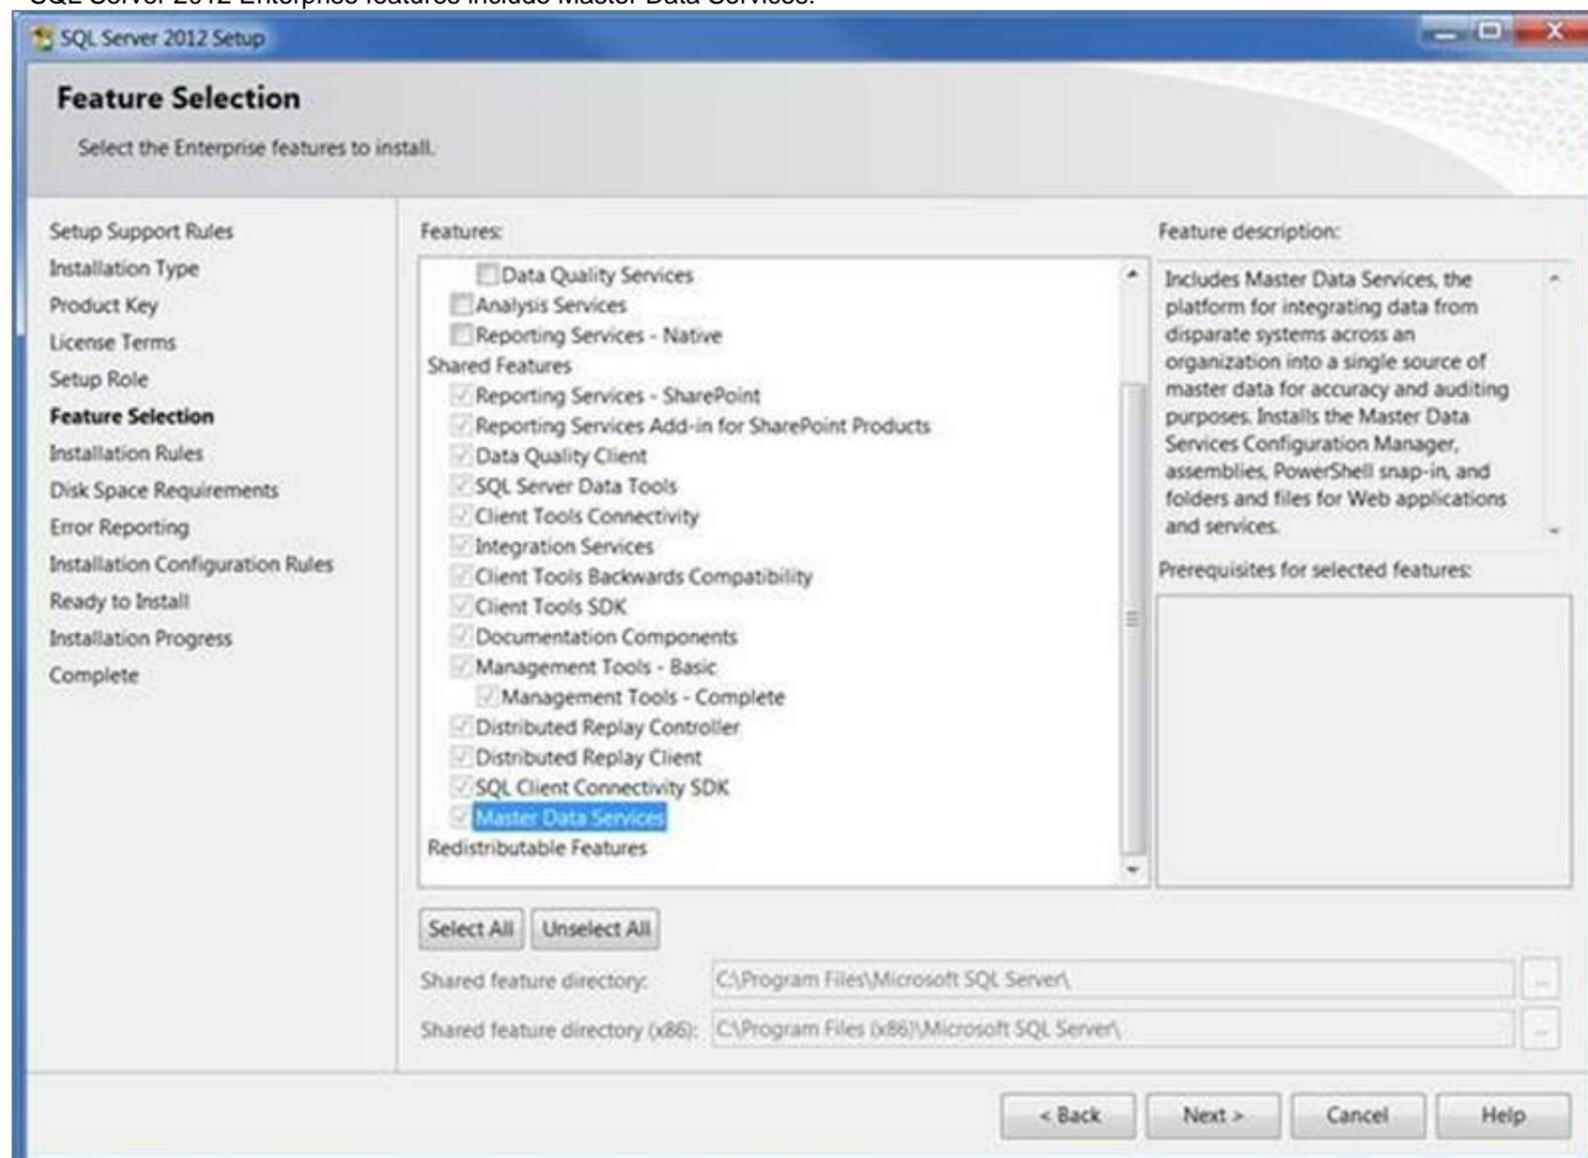
You are preparing to install SQL Server 2012 Master Data Services (MDS). You need to ensure that the database requirements are met. What should you install?

- A. Microsoft SharePoint Foundation 2010 SP1
- B. SQL Server 2012 Enterprise (64-bit) x64 on the database server
- C. SQL Server 2012 Data Center (64-bit) x64 on the database server
- D. SQL Server 2008 Enterprise (64-bit) x64 on the database server

Answer: B

Explanation: * Master Data Services is a new feature introduced in SQL Server 2008 R2 and further enhanced in SQL Server 2012.

* SQL Server 2012 Enterprise features include Master Data Services:



Note:

* Microsoft SQL Server Master Data Services is a Master Data Management (MDM) product from Microsoft, which will ship as a part of the Microsoft SQL Server database. Originally code-named Bulldog, Master Data Services is the rebranding of the Stratature MDM product titled +EDM, which Microsoft acquired in June 2007. Master Data Services is architecturally similar to +EDM, with increased integration with other Microsoft applications as well as some new features. Master Data Services first shipped with Microsoft SQL Server 2008 R2.

NEW QUESTION 263

You are installing the Data Quality Server component of Data Quality Services. You need to provision the hardware and install the software for the server that runs the Data Quality Server. You must ensure that the minimum Data Quality Server prerequisites are met. What should you do?

- A. Install SQL Server 2012 Database Engine.
- B. Install Microsoft SharePoint Server 2010 Enterprise Edition with PowerPivot.
- C. Make sure the server has at least 4 GB of RAM.
- D. Install Microsoft Internet Explorer 6.0 SP1 or later.

Answer: A

Explanation: Data Quality Server Minimum System Requirements

* SQL Server 2012 Database Engine.

* Memory (RAM): Minimum: 2 GB Recommended: 4 GB or more

Note: SQL Server Data Quality Services (DQS) is a new feature in SQL Server 2012 that contains the following two components: Data Quality Server and Data Quality Client.

NEW QUESTION 268

DRAG DROP

Your company is evaluating the data cleansing capabilities of SQL Server Data Quality Services (DQS). A stand-alone server will be used to host DQS, as well as all related services necessary to run this service. You need to install the DQS components on the server by using the minimum amount of steps possible. You also need to grant permissions to a login named DataStewards to enable DataStewards to edit and execute DQS projects. Which four actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Answer:

Explanation: Note:

Step 1-2:

* SQL Server Data Quality Services (DQS) is a new feature in SQL Server 2012 that contains the following two components: Data Quality Server and Data Quality Client.

* You have to use the SQL Server 2012 Setup to install DQS components. When you run the SQL Server Setup, you have to go through a series of installation wizard pages to select appropriate options based on your requirements.

/ Feature Selection

Data Quality Services under Database Engine Services to install the Data Quality Server. Data Quality Client to install Data Quality Client.

Step 3-4:

* The Data Quality Services (DQS) security infrastructure is based upon the SQL Server security infrastructure. A database administrator grants a user a set of permissions by associating the user with a DQS role. Doing so determines the DQS resources that the user has access to and the functional activities that the user is allowed to perform.

* There are four roles for DQS. One is the database administrator (DBA) who deals primarily with product installation, database maintenance, and user management. This role primarily uses the SQL Server Management Studio, rather than within the Data Quality Client application. Their server role is sysadmin.

The three other roles are information workers, data stewards who use the product directly by working in the Data Quality Client application. These roles include the following:

/ The DQS Administrator (dqs_administrator role) can do everything in the scope of the product.

/ The DQS KB Editor (dqs_kb_editor role) can perform all of the DQS activities, except for administration.

/ The DQS KB Operator (dqs_kb_operator role) can edit and execute a project. They cannot perform any kind of knowledge management; they cannot create or change a knowledge base. They can see the activity monitoring data, but cannot terminate an activity or perform administrative duties.

Ref: [http://msdn.microsoft.com/en-gb/library/gg492277\(v=sql.110\).aspx](http://msdn.microsoft.com/en-gb/library/gg492277(v=sql.110).aspx)

NEW QUESTION 269

HOTSPOT

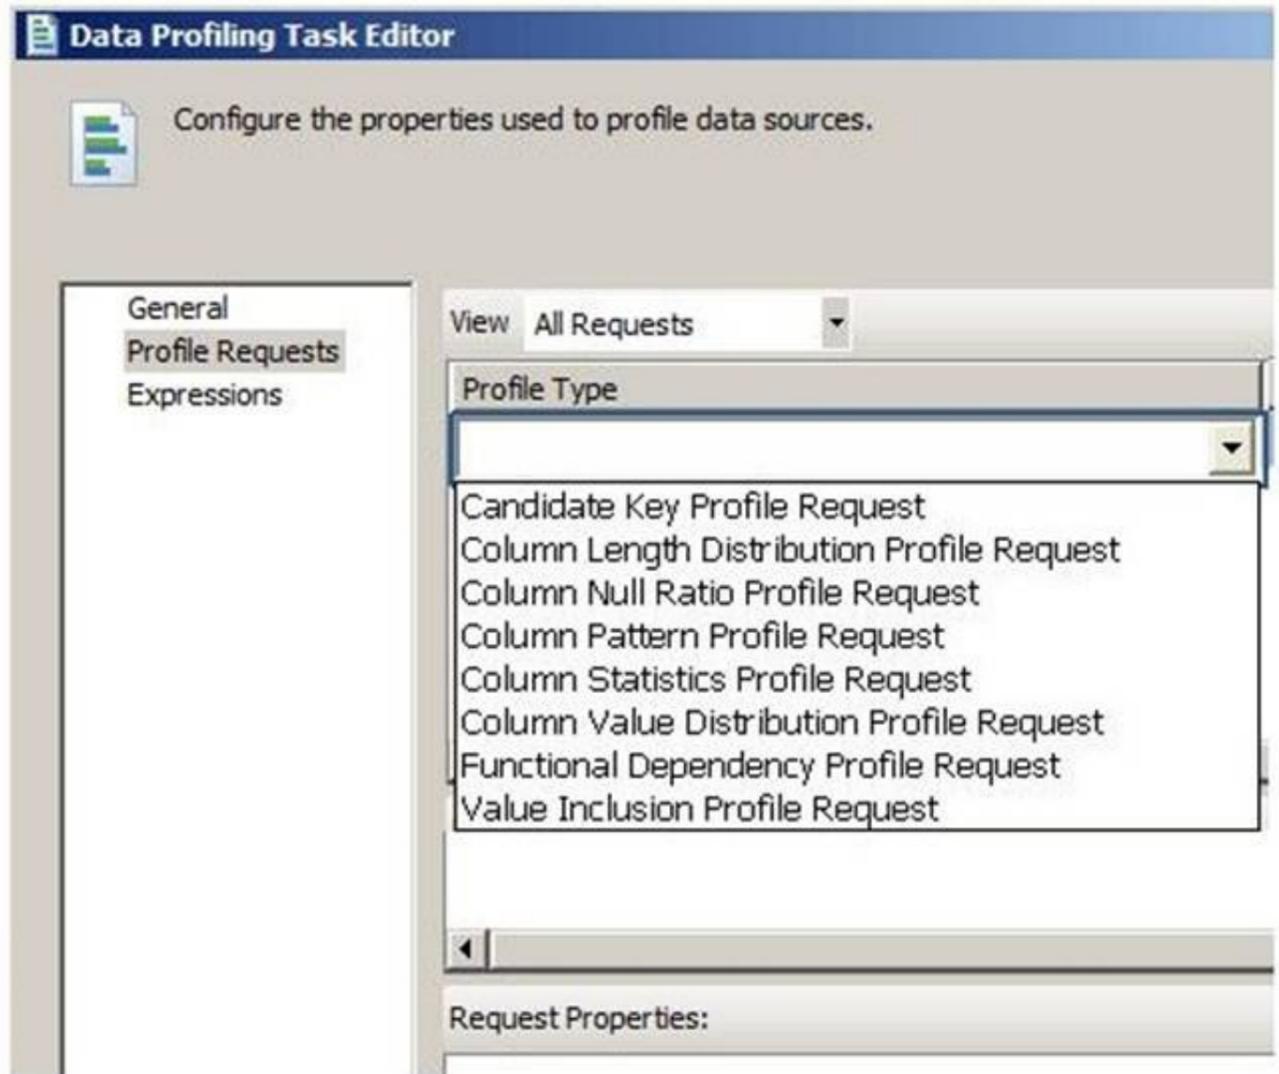
You are developing a SQL Server Integration Services (SSIS) package.

You use a Data Profiling task to examine the data from a source system. You need to establish: The minimum and maximum dates for the datetime columns in the source data

The minimum, maximum, and average values for numeric columns in the source data

You need to use the appropriate profile type in the Data Profiling task.

Which profile type should you use? (To answer, select the appropriate profile type in the answer area.)



Answer:

Explanation:



NEW QUESTION 270

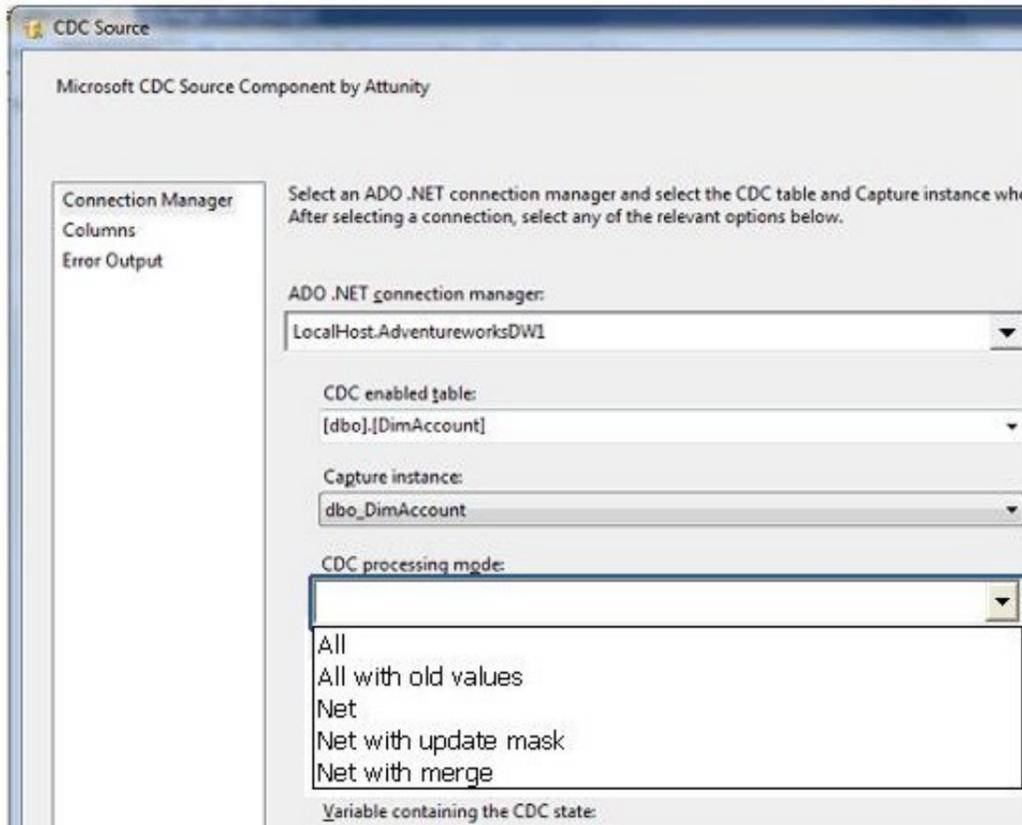
HOTSPOT

You are developing a SQL Server Integration Services (SSIS) package. The data source for the data flow task is a table that has been configured as a change data capture (CDC) table. You are using a CDC Source component to obtain the CDC data.

The data source will be polled once per hour. The data is updated with multiple important status changes per minute.

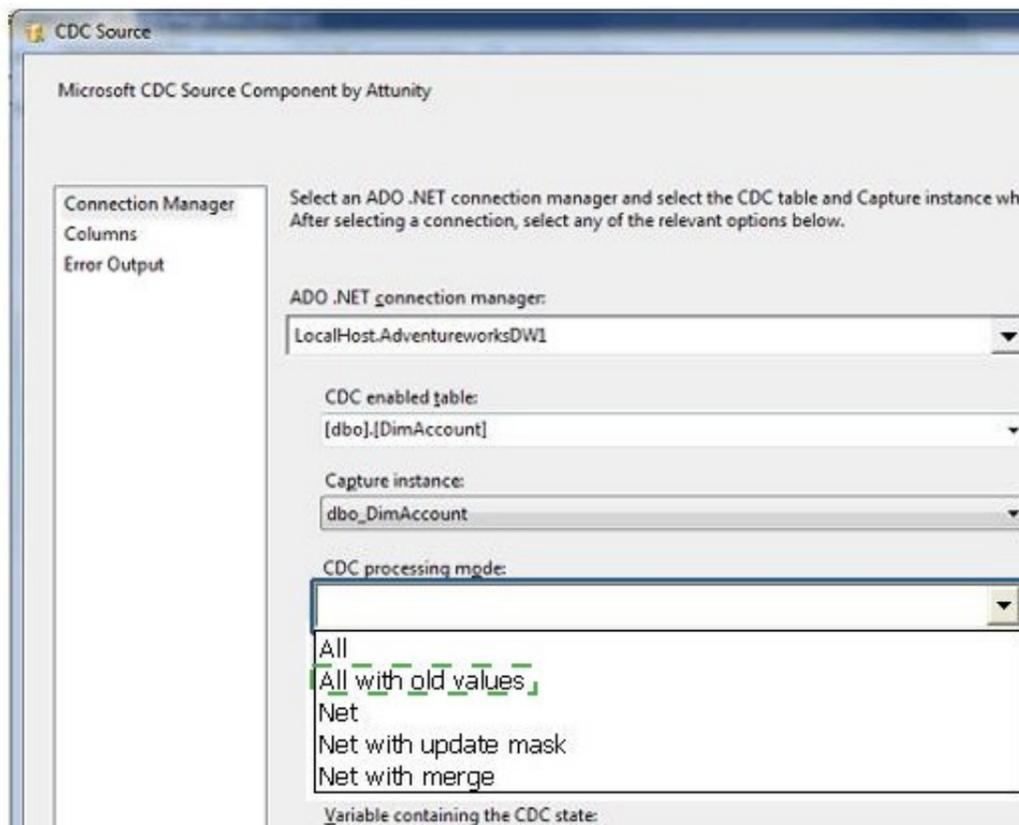
For each captured data change, the before and after values must be included. You need to configure the CDC Source component.

Which CDC processing mode should you select? (To answer, configure the appropriate option in the dialog box in the answer area.)



Answer:

Explanation:



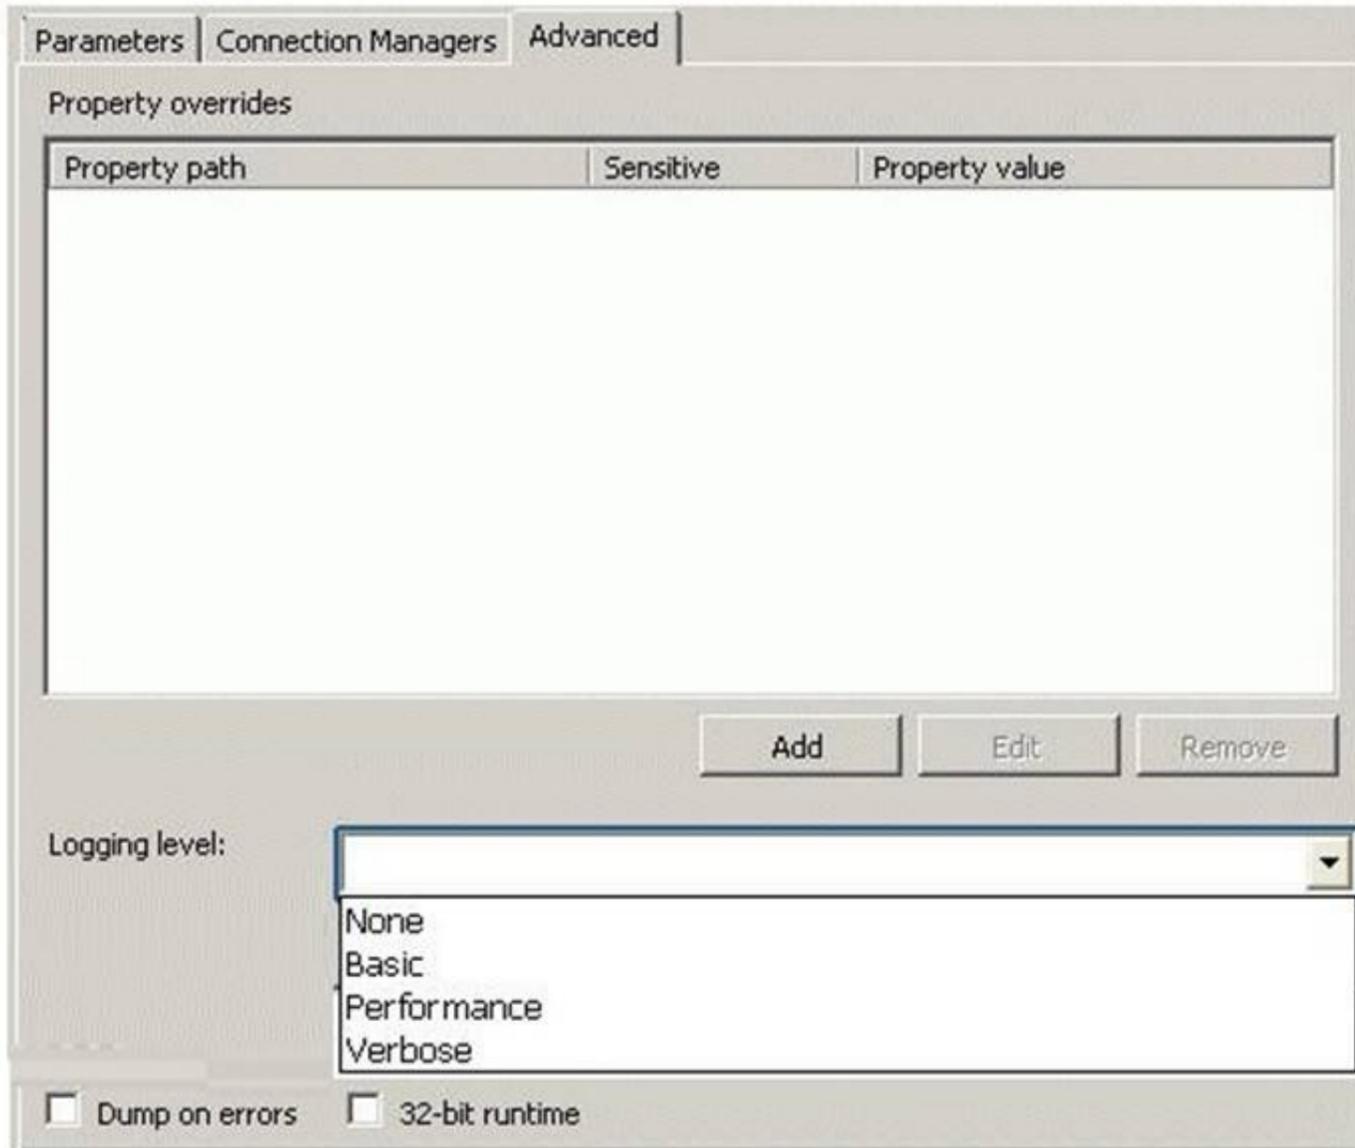
NEW QUESTION 273

HOTSPOT

You are developing a SQL Server Integration Services (SSIS) package.

OnError and OnWarning events must be logged for viewing in the built-in SSIS reports by using SQL Server Management Studio.

You need to execute the package and minimize the number of event types that are logged. Which setting should you use? (To answer, change the appropriate setting in the answer area.)



Answer:

Explanation: Ref: <http://msdn.microsoft.com/en-gb/library/hh231191.aspx>

NEW QUESTION 275

You are performance tuning a SQL Server Integration Services (SSIS) package to load sales data from a source system into a data warehouse that is hosted on Windows Azure SQL Database. The package contains a data flow task that has 10 source-to-destination execution trees. Only five of the source-to-destination execution trees are running in parallel. You need to ensure that all the execution trees run in parallel. What should you do?

- A. Add OnError and OnWarning event handlers.
- B. Create a new project and add the package to the project.
- C. Query the ExecutionLog table.
- D. Add an Execute SQL task to the event handler

Answer: A

NEW QUESTION 279

You are the administrator of a server that hosts Data Quality Server for a large retail company. The server had a hardware failure during business hours. You need to restore the server that hosts Data Quality Server to another server. You have a recent backup of all the required databases. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Restore the DQS_MAIN, DQS_PROJECTS, and DQS_STAGING_DATA databases to another server as soon as possible.
- B. Execute the DQS_MAIN.internal_core.RestoreDQDatabases stored procedure with the appropriate parameter.
- C. Restore only the DQS_MAIN and DQS_STAGING_DATA databases to another server as soon as possible.
- D. Execute the DQS_MAIN.internal_core.InitServer stored procedure with the appropriate parameter.

Answer: AB

Explanation: Steps to restore DQS Databases: Restore DQS_MAIN database.

Restore the DQS_PROJECTS database. Restore the DQS_STAGING_DATA database.

In Object Explorer, right-click the server, and then click New Query.

In the Query Editor window, copy the following SQL statements, and replace <PASSWORD> with the password that you provided during the DQS installation for the database master key:

```
USE [DQS_MAIN] GO
EXECUTE [internal_core].[RestoreDQDatabases] '<PASSWORD>' GO
```

Press F5 to execute the statements. Check the Results pane to verify that the statements have executed successfully.

Note:

* Backup and restore of SQL Server databases are common operations that database administrators perform for preventing loss of data in a case of disaster by recovering data from the backup databases. Data Quality Server is primarily implemented by two SQL Server databases: DQS_MAIN and DQS_PROJECTS. The backup and restore procedures of the Data Quality Services (DQS) databases are similar to any other SQL Server databases.

Ref: [http://msdn.microsoft.com/en-gb/library/hh213068\(v=sql.110\).aspx](http://msdn.microsoft.com/en-gb/library/hh213068(v=sql.110).aspx)

NEW QUESTION 284

HOTSPOT

You are a data warehouse developer responsible for developing data cleansing processes. Duplicate employees exist in an employee dimension. You need to map, discover, and manage domain values based on the employee dimension.

Which Data Quality Services (DQS) option should you use? (To answer, select the appropriate option in the answer area.)



Answer:

Explanation:



NEW QUESTION 288

You are creating a Data Quality Services (DQS) solution. You must provide statistics on the accuracy of the data. You need to use DQS profiling to obtain the required statistics. Which DQS activity should you use?

- A. Cleansing
- B. Matching
- C. Knowledge Discovery
- D. Matching Policy

Answer: A

Explanation: Ref: [http://msdn.microsoft.com/en-gb/library/hh213055\(v=sql.110\).aspx](http://msdn.microsoft.com/en-gb/library/hh213055(v=sql.110).aspx)

NEW QUESTION 292

You are developing a SQL Server Integration Services (SSIS) project to read and write data from a Windows Azure SQL Database database to a server that runs SQL Server 2012.

The connection will be used by data flow tasks in multiple SSIS packages. The address of the target Windows Azure SQL Database database will be provided by a project parameter.

You need to create a solution to meet the requirements by using the least amount of administrative effort.

What should you do?

- A. Add a SQLMOBILE connection manager to each package.
- B. Add an ADO.NET project connection manager.
- C. Add a SQLMOBILE project connection manager.
- D. Add an ADO.NET connection manager to each data flow task.
- E. Add a SQLMOBILE connection manager to each data flow task.
- F. Add an ADO.NET connection manager to each package.

Answer: B

Explanation: Ref: <http://www.databasejournal.com/features/mssql/windows-azure-sql-database-uploading-databy-using-sql-server-integration-services.html>

NEW QUESTION 295

You are implementing a SQL Server Integration Services (SSIS) package that imports Microsoft Excel workbook data into a Windows Azure SQL Database database. The package has been deployed to a production server that runs Windows Server 2008 R2 and SQL Server 2012.

The package fails when executed on the production server.

You need to ensure that the package can load the Excel workbook data without errors. You need to use the least amount of administrative effort to achieve this goal.

What should you do?

- A. Create a custom SSIS source component that encapsulates the 32-bit driver and compile it in 64-bit mode.
- B. Install a 64-bit ACE driver and execute the package by using the 64-bit run-time option.
- C. Execute the package by using the 32-bit run-time option.
- D. Replace the SSIS Excel source with a SSIS Flat File source.

Answer: C

Explanation: * See step 3 below.

To publish an Excel worksheet to Azure SQL Database, your package will contain a Data Flow Task, Excel Source task, and ADO NET Destination.

- 1) Create an SSIS project.
- 2) Drop a Data Flow Task onto the Control Flow design surface, and double click the Data Flow Task.
- 3) Drop an Excel Source onto the Data Flow design surface.

Note When using the Excel Source task on a 64-bit machine, set Run64BitRuntime to False.

*Incorrect:

Not D: The Flat File source reads data from a text file. The text file can be in delimited, fixed width, or mixed format.

NEW QUESTION 298

You are installing the Data Quality Server component of Data Quality Services.

You need to provision the hardware and install the software for the server that runs the Data Quality Server.

You must ensure that the minimum Data Quality Server prerequisites are met. What should you do?

- A. Install Microsoft Internet Explorer 6.0 SP1 or later.
- B. Install SQL Server 2012 Database Engine.
- C. Make sure the server has at least 4 GB of RAM.
- D. Install Microsoft SharePoint Server 2010 Standard Edition.

Answer: B

Explanation: Data Quality Server Minimum System Requirements

* SQL Server 2012 Database Engine.

* Memory (RAM): Minimum: 2 GB Recommended: 4 GB or more

Note: SQL Server Data Quality Services (DQS) is a new feature in SQL Server 2012 that contains the following two components: Data Quality Server and Data Quality Client.

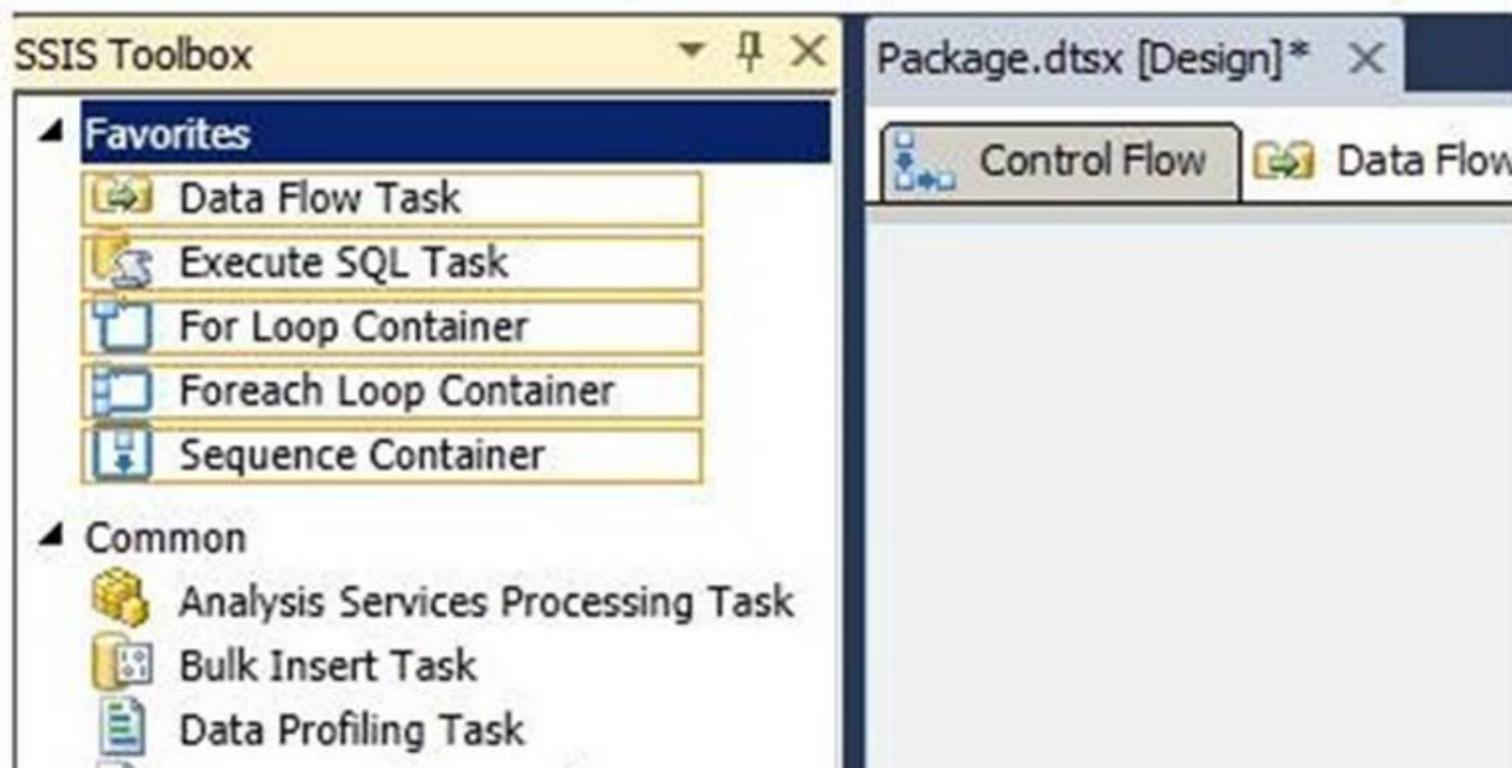
NEW QUESTION 301

HOTSPOT

You are developing a SQL Server Integration Services (SSIS) package.

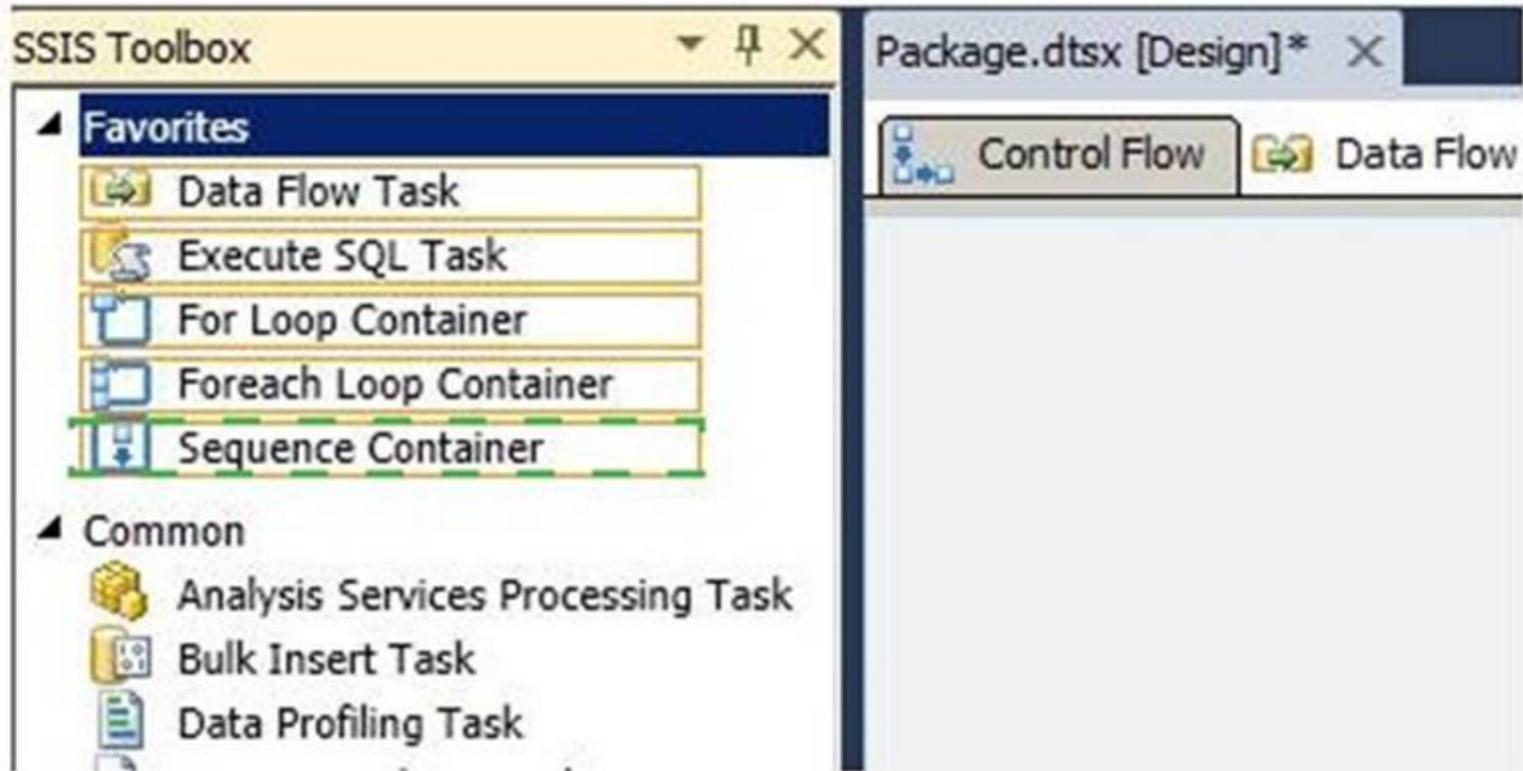
The package control flow will contain many tasks. The tasks will execute consecutively and none will execute more than once. Certain groups of tasks will share variable and transaction scope.

You need to group tasks together while enabling them to be collapsed and expanded as a group. Which item should you use from the SSIS Toolbox? (To answer, select the appropriate item in the answer area.)



Answer:

Explanation:



NEW QUESTION 302

HOTSPOT

You are developing a SQL Server Integration Services (SSIS) package to load data into a SQL Server 2012 database.

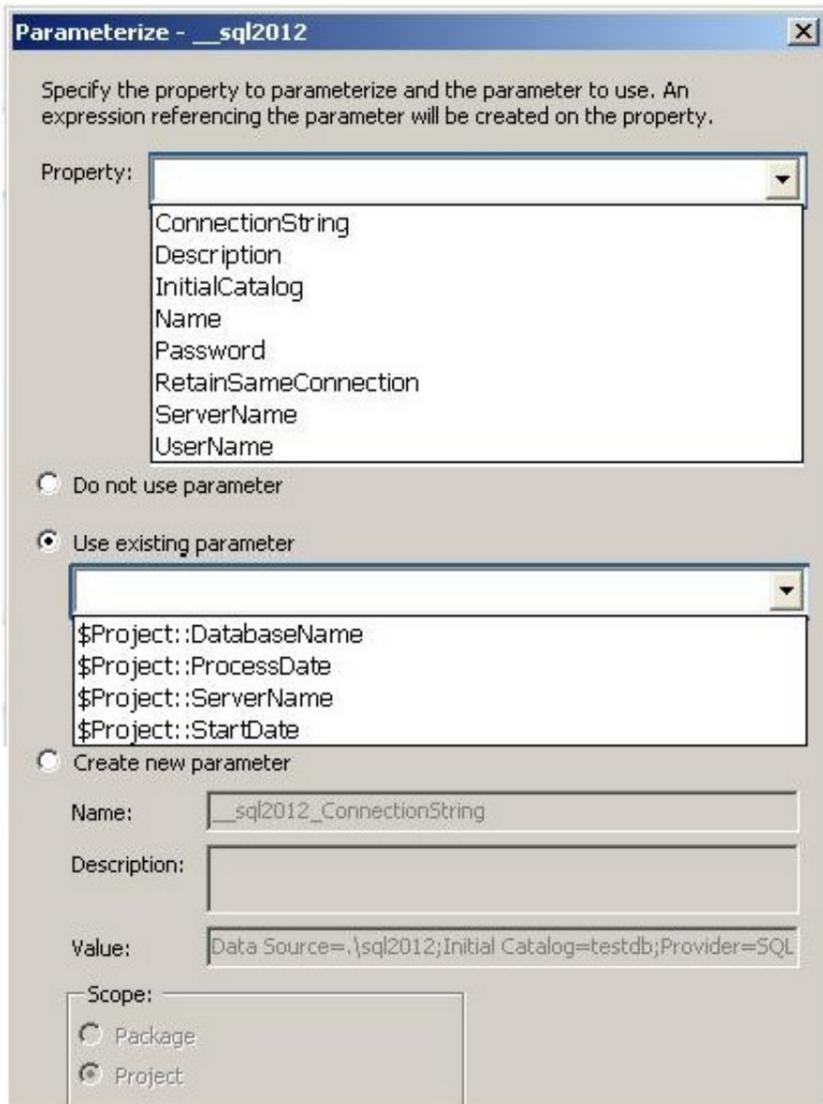
The name of the target instance is passed to the package via an Environment variable.

The Environment variable provides the server name to the Project parameter named ServerName. The OLE DB Project Connection Manager has been parameterized.

You need to configure the Connection Manager property to accept the parameter.

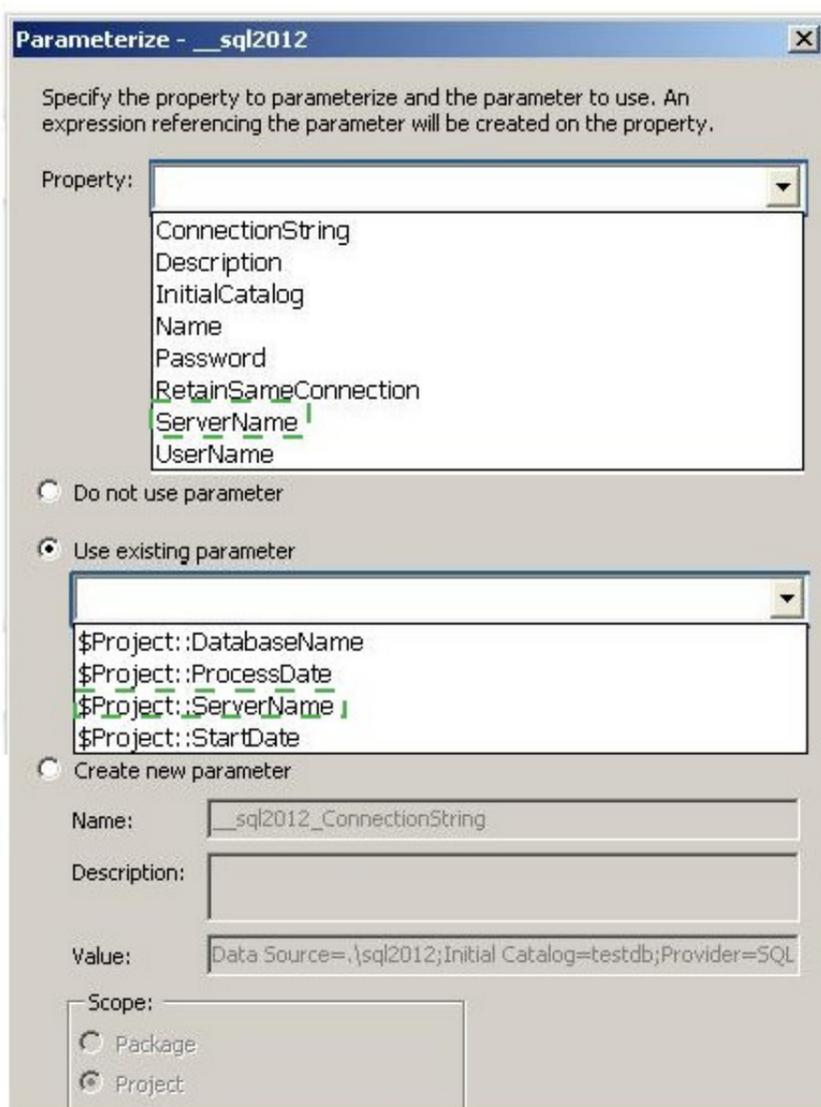
Which property and parameter should you use? (To answer, configure the appropriate option or options in the dialog box in the answer area.)





Answer:

Explanation:



NEW QUESTION 307

You are designing a data warehouse that contains a customer dimension.

The customer dimension contains the original customer attributes as well as the most recent set of updated attributes. When a customer is updated, the original customer attributes must be preserved.

The customer dimension design must:

Support a Type 3 Slowly Changing Dimension (SCD) process Minimize the total storage consumed

You need to design a solution that meets these requirements. What should you do?

- A. Design two table
- B. One table will hold the original customer surrogate key and attribute
- C. The other table will hold the most recent customer surrogate key and attributes.
- D. Design a table as a Type 2 SCD dimension that holds the surrogate key for each customer and its attribute
- E. Limit the number of records per customer within the table to two.
- F. Design a table to hold the surrogate key for each customer and its attribute
- G. Add a column to store an indicator that identifies the record as original data or recent data.
- H. Design a table to hold the surrogate key for each customer and its attributes as well as additional columns to store the original attributes.

Answer: D

NEW QUESTION 311

You are designing a data warehouse hosted on Windows Azure SQL Database. The data warehouse currently includes the dimUser and dimRegion dimension tables and the factSales fact table. The dimUser table contains records for each user permitted to run reports against the warehouse, and the dimRegion table contains information about sales regions.

The system is accessed by users from certain regions, as well as by area supervisors and users from the corporate headquarters.

You need to design a table structure to ensure that certain users can see sales data for only certain regions. Some users must be permitted to see sales data from multiple regions.

What should you do?

- A. For each region, create a view of the factSales table that includes a WHERE clause for the region.
- B. Create a userRegion table that contains primary key columns from the dimUser and dimRegion tables.
- C. Add a region column to the dimUser table.
- D. Partition the factSales table on the region column

Answer: B

NEW QUESTION 313

You are designing a data warehouse that uses SQL Server 2012.

The data warehouse contains a table named factSales that stores product sales. The table has a clustered index on the primary key, four foreign keys to dimension tables, and an aggregate column for sales totals. All key columns use the INT data type, and the aggregate column uses the MONEY data type.

You need to increase the speed of data retrieval from the factSales table. Which index type should you add to the table?

- A. Full text
- B. Spatial
- C. Non-clustered
- D. Clustered

Answer: C

NEW QUESTION 317

You are designing a data warehouse with two fact tables. The first table contains sales per month and the second table contains orders per day.

Referential integrity must be enforced declaratively.

You need to design a solution that can join a single time dimension to both fact tables. What should you do?

- A. Create a time mapping table.
- B. Partition the fact tables by day.
- C. Create a time dimension that can join to both fact tables at their respective granularity.
- D. Join the two fact table

Answer: C

NEW QUESTION 321

You are implementing a new SQL Server Integration Services (SSIS) 2012 package that loads data from various flat files and a Windows Azure SQL Database database.

Daily transactions must be loaded into a staging database. All the SSIS tasks will use the CurrentDate variable as the transaction date.

You need to set the CurrentDate variable to the current execution date when the package starts. You need to achieve this goal by using the least amount of development effort.

What should you use to set the variable?

- A. An Expression task
- B. A Script component
- C. A Script task
- D. An Execute SQL task

Answer: A

Explanation: Ref: <http://msdn.microsoft.com/en-us/library/ms141214.aspx> <http://blog.sqltechie.com/2012/11/expression-task-in-ssis-2012.html>

NEW QUESTION 323

You are developing a SQL Server Integration Services (SSIS) package.

You need to design a package to change a variable value during package execution by using the least amount of development effort.

What should you use?

- A. Expression task

- B. Data Cleansing transformation
- C. Fuzzy Lookup transformation
- D. Term Lookup transformation
- E. Data Profiling task

Answer: A

NEW QUESTION 324

You are designing a SQL Server Integration Services (SSIS) project that uses the Project Deployment Model. The project will contain 10 packages. All the packages must connect to the same data source and destination. You need to define and reuse the connection managers in all the packages by using the least development effort. What should you do?

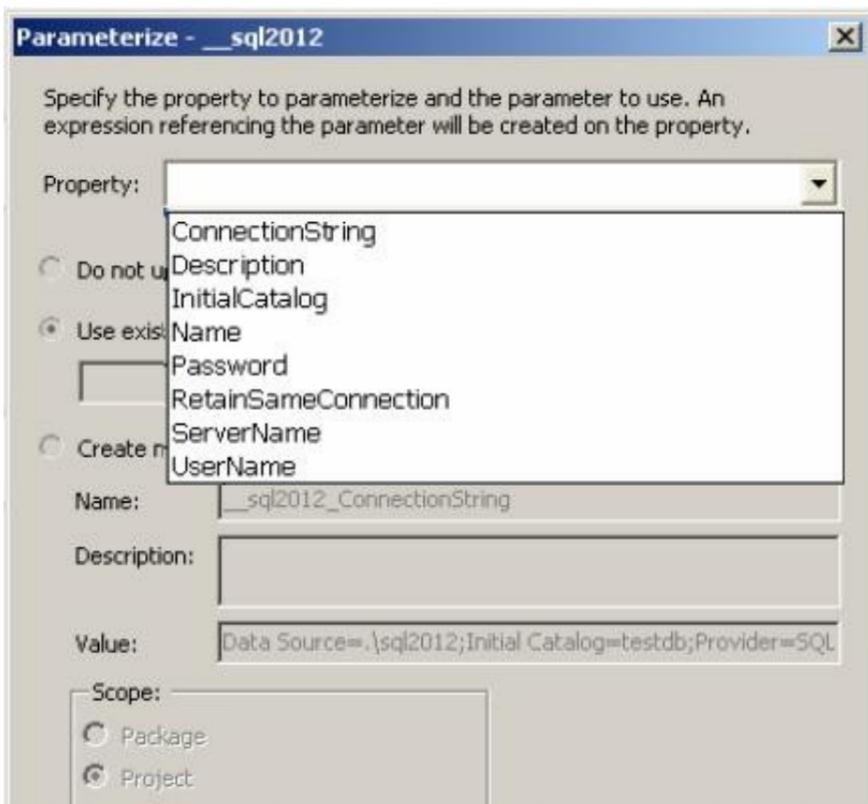
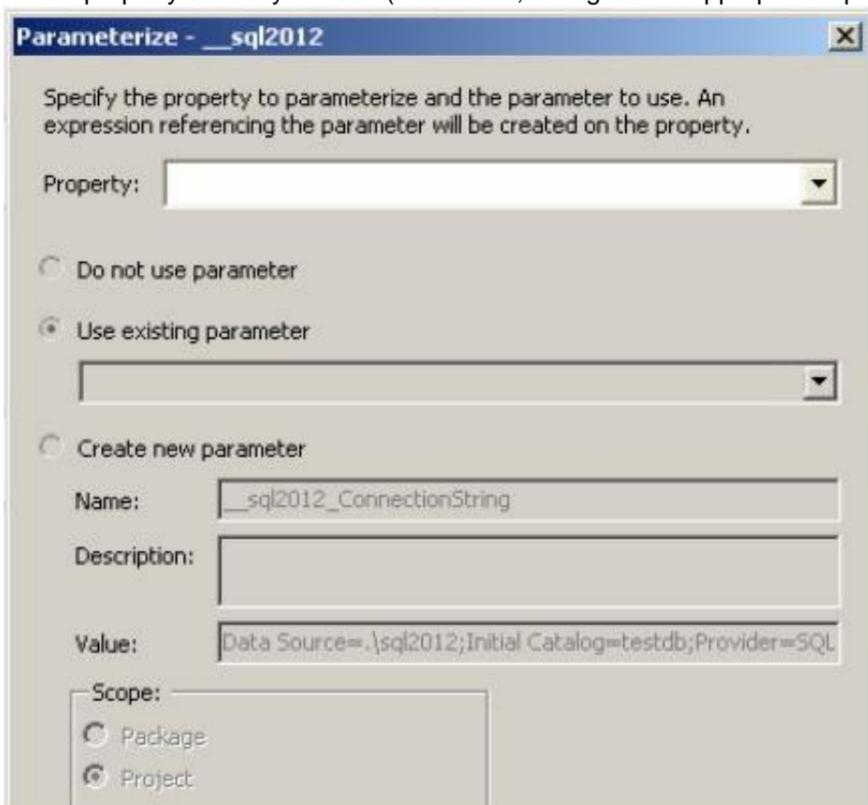
- A. Implement parent package variables.
- B. Copy and paste connection managers from one package to the others.
- C. Implement project connection managers.
- D. Implement package connection manager

Answer: C

NEW QUESTION 327

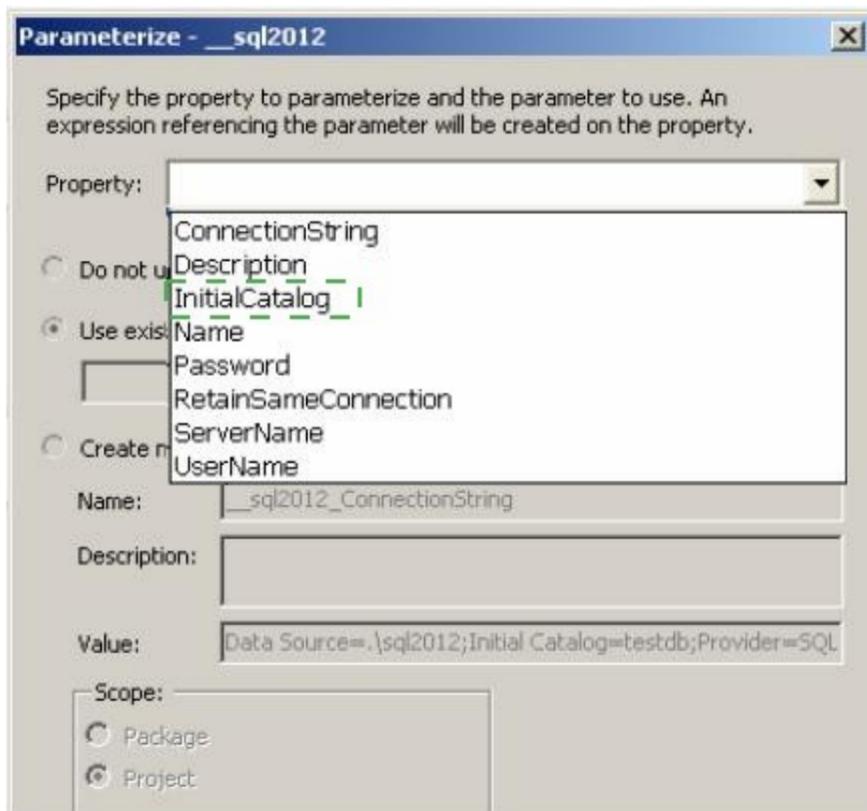
HOTSPOT

You are developing a SQL Server Integration Services (SSIS) package to load data into a SQL Server 2012 database. The package is allowed to connect to only one database. An Environment variable contains the name of the database. The OLE DB project connection manager has been parameterized. You need to configure the connection manager property to accept the value of the Environment variable. Which property should you use? (To answer, configure the appropriate option or options in the dialog box in the answer area.)



Answer:

Explanation:



NEW QUESTION 329

You are designing a SQL Server Integration Services (SSIS) data flow to load sales transactions from a source system into a data warehouse hosted on Windows Azure SQL Database. One of the columns in the data source is named ProductCode. Some of the data to be loaded will reference products that need special processing logic in the data flow. You need to enable separate processing streams for a subset of rows based on the source product code. Which Data Flow transformation should you use?

- A. Multicast
- B. Conditional Split
- C. Script Task
- D. Data Conversion

Answer: B

Explanation: We use Conditional Split to split the source data into separate processing streams. A Script Component (Script Component is the answer to another version of this question) could be used but this is not the same as a Script Task.

NEW QUESTION 333

You are designing a SQL Server Integration Services (SSIS) 2012 package that imports data from a Windows Azure SQL Database database into a SQL Server database. The SSIS package has the following requirements: Every night, a very large amount of data is imported into the staging database. Package processing time must be minimized. The package must run on its own dedicated server when it is deployed to production. Transaction log activity when data is imported must be minimized. You need to design the package to meet the requirements. Which destination component should you use?

- A. DataReader
- B. OLE DB
- C. ADO.NET
- D. SQL Server

Answer: B

NEW QUESTION 336

You are designing a data warehouse that uses SQL Server 2012. The data warehouse contains a table named factSales that stores product sales. The table has a clustered index on the primary key, four foreign keys to dimension tables, and an aggregate column for sales totals. All key columns use the int data type and the aggregate column uses the money data type. You need to increase the speed of data retrieval from the factSales table. Which index type should you add to the table?

- A. Clustered
- B. Semantic search
- C. Nonclustered
- D. XML

Answer: C

NEW QUESTION 337

You are designing a complex SQL Server Integration Services (SSIS) project that uses the Project Deployment model. The project will contain between 15 and 20 packages. All the packages must connect to the same data source and destination.

You need to define and reuse the connection managers in all the packages by using the least development effort. What should you do?

- A. Copy and paste the connection manager details into each package.
- B. Implement project connection managers.
- C. Implement package connection managers.
- D. Implement parent package variables in all package

Answer: B

NEW QUESTION 340

You are designing a SQL Server Integration Services (SSIS) data flow to load sales transactions from a source system into a data warehouse hosted on Windows Azure SQL Database. One of the columns in the data source is named ProductCode. Some of the data to be loaded will reference products that need special processing logic in the data flow. You need to enable separate processing streams for a subset of rows based on the source product code. Which Data Flow transformation should you use?

- A. Script Component
- B. Audit
- C. Destination Assistant
- D. Data Conversion

Answer: A

Explanation: The Script component can be used as a source, a transformation, or a destination. This component supports one input and multiple outputs. Depending on how the component is used, it supports either an input or outputs or both. The script is invoked by every row in the input or output. If used as a source, the Script component supports multiple outputs. If used as a transformation, the Script component supports one input and multiple outputs. If used as a destination, the Script component supports one input.

NEW QUESTION 343

You develop a SQL Server Integration Services (SSIS) package in a project by using the Project Deployment Model. It is regularly executed within a multi-step SQL Server Agent job. You make changes to the package that should improve performance. You need to establish if there is a trend in the durations of the next 10 successful executions of the package. You need to use the least amount of administrative effort to achieve this goal. What should you do?

- A. Enable logging to an XML file in the package control flow for the OnInformation even
- B. After 10 executions, view the XML file.
- C. After 10 executions, view the job history for the SQL Server Agent job.
- D. After 10 executions, in SQL Server Management Studio, view the Execution Performance subsection of the All Executions report for the package.
- E. Enable logging to an XML file in the package control flow for the OnPostExecute even
- F. After 10 executions, view the XML file.

Answer: C

NEW QUESTION 345

You are designing a SQL Server Integration Services (SSIS) package that uses the Fuzzy Lookup transformation. The reference data to be used in the transformation changes with every package execution. Common language runtime (CLR) integration cannot be enabled on the SQL Server database where the reference table is located. You need to configure the Fuzzy Lookup transformation in the most efficient manner. What should you do?

- A. Execute the sp_FuzzyLookupTableMaintenanceInvoke stored procedure.
- B. Select the GenerateNewIndex option in the Fuzzy Lookup Transformation Editor.
- C. Execute the sp_FuzzyLookupTableMaintenanceUninstall stored procedure.
- D. Select the GenerateAndPersistNewIndex option in the Fuzzy Lookup Transformation Editor.

Answer: B

NEW QUESTION 349

You are designing a data warehouse that uses SQL Server 2012. You are preparing to update the contents of a fact table that contains a non-clustered columnstore index. You need to run an update statement against the table. What should you do first?

- A. Pause the columnstore index.
- B. Change the recovery model of the database to Bulk-logged.
- C. Change the non-clustered columnstore index to be a clustered columnstore index.
- D. Drop the columnstore index

Answer: D

NEW QUESTION 351

You are using SQL Server Data Tools to develop a SQL Server Integration Services (SSIS) project. The first package that you create in this project contains a package connection that accesses a Microsoft Excel file. Additional packages in the project must also access this file. You need to define and reuse the Excel file connection in all project packages. What should you do?

- A. Copy the package Connection Manager and paste it in the second package.

- B. Set the RetainSameConnection property of the package Connection Manager to True.
- C. Convert the package Connection Manager in the first package to a project Connection Manager.
- D. Convert the project to the Package Deployment mode

Answer: C

NEW QUESTION 352

You have Master Data Services (MDS) in a SQL Server deployment. You need to deploy a model that contains data to the MDS repository. What should you use?

- A. the Model Deployment Wizard
- B. the MDSModelDeploy tool
- C. the Master Data Services Configuration Manager tool
- D. the Integration Services Deployment Wizard

Answer: B

Explanation: <https://msdn.microsoft.com/en-us/library/ff486956.aspx>

NEW QUESTION 353

Your company has several line-of-business applications. The applications use a server that has SQL Server installed and contains several databases. You need that the business applications can access curated and validated data from the databases. Which features should you deploy to the server?

- A. Data Quality Services (DQS)
- B. SQL Server Analysis Services (SSAS) dimension hierarchies.
- C. SQL Server Integration Services (SSIS)
- D. Master Data Services (MDS) subscription

Answer: C

NEW QUESTION 358

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-463 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-463 Product From:

<https://www.2passeasy.com/dumps/70-463/>

Money Back Guarantee

70-463 Practice Exam Features:

- * 70-463 Questions and Answers Updated Frequently
- * 70-463 Practice Questions Verified by Expert Senior Certified Staff
- * 70-463 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 70-463 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year