


## Microsoft MCSA Certification 70-767 Exam



- Vendor: Microsoft
- Exam Code: 70-767
- Exam Name: Implementing a Data Warehouse using SQL

**Get Complete Version Exam 70-767 Dumps with VCE and PDF Here**



<https://www.passleader.com/70-767.html>

**QUESTION 1**

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**

**QUESTION 2**

You are implementing the indexing strategy for a fact table in a data warehouse. The fact table is named Quotes. The table has no indexes and consists of seven columns:

- [ID]
- [QuoteDate]
- [Open]
- [Close]
- [High]
- [Low]
- [Volume]

Each of the following queries must be able to use a columnstore index:

```
SELECT AVG ([Close]) AS [AverageClose] FROM Quotes WHERE [QuoteDate]
BETWEEN '20100101' AND '20101231'.
SELECT AVG ([High] - [Low]) AS [AverageRange] FROM Quotes WHERE [QuoteDate]
BETWEEN '20100101' AND '20101231'.
SELECT SUM ([Volume]) AS [SumVolume] FROM Quotes WHERE [QuoteDate] BETWEEN
'20100101' AND '20101231'.
```

You need to ensure that the indexing strategy meets the requirements. The strategy must also minimize the number and size of the indexes. What should you do?

- A. Create one columnstore index that contains [ID], [Close], [High], [Low], [Volume], and [QuoteDate].
- B. Create three columnstore indexes:
  - One containing [QuoteDate] and [Close]
  - One containing [QuoteDate], [High], and [Low]
  - One containing [QuoteDate] and [Volume]
- C. Create one columnstore index that contains [QuoteDate], [Close], [High], [Low], and [Volume].
- D. Create two columnstore indexes:
  - One containing [ID], [QuoteDate], [Volume], and [Close]
  - One containing [ID], [QuoteDate], [High], and [Low]

**Answer: C**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/gg492088.aspx>

<http://msdn.microsoft.com/en-us/library/gg492153.aspx>

**QUESTION 3**

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.

**[70-767 Exam Dumps](#) [70-767 Exam Questions](#) [70-767 PDF Dumps](#) [70-767 VCE Dumps](#)**

**[Back to the Source of this PDF & Get More Free Braindumps -- www.microsoftbraindumps.com](#)**

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 dimensions.

**Answer: C**

#### **QUESTION 4**

To facilitate the troubleshooting of SQL Server Integration Services (SSIS) packages, a logging methodology is put in place. The methodology has the following requirements:

- The deployment process must be simplified.
- All the logs must be centralized in SQL Server.
- Log data must be available via reports or T-SQL.
- Log archival must be automated.

You need to configure a logging methodology that meets the requirements while minimizing the amount of deployment and development effort. What should you do?

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

**Answer: J**

#### **Explanation:**

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

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

#### **QUESTION 5**

You are developing a SQL Server Integration Services (SSIS) project that copies a large amount of rows from a SQL Azure database. The project uses the Package Deployment Model. This project is deployed to SQL Server on a test server. You need to ensure that the project is deployed to the SSIS catalog on the production server. What should you do?

- A. Open a command prompt and run the dtexec/dumperror/conn command.
- B. Create a reusable custom logging component and use it in the SSIS project.
- C. Open a command prompt and run the gacutil command.
- D. Add an OnError event handler to the SSIS project.
- E. Open a command prompt and execute the package by using the SQL Log provider and running the dtexecui.exe utility.
- F. Open a command prompt and run the dtexec/rep/conn command.

- G. Open a command prompt and run the dtutil/copy command.
- H. Use an msi file to deploy the package on the server.
- 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 in SQL Server.

**Answer: I**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/hh231102.aspx>

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

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

#### **QUESTION 6**

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:**

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

#### **QUESTION 7**

Drag and Drop Question

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:**

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.

**QUESTION 8**

You develop and deploy a SQL Server Integration Services (SSIS) package. The package is stored in the file system. You need to execute the package without importing it to the SSIS server. What should you use to execute the package? (Each correct answer presents a complete solution. Choose all that apply.)

- A. catalog.start\_package
- B. dtexec
- C. SQL Server Management Studio
- D. SQL Server Agent

**Answer: BD**

**Explanation:**

[http://technet.microsoft.com/en-us/library/ms141708\(v=sql.110\).aspx](http://technet.microsoft.com/en-us/library/ms141708(v=sql.110).aspx)

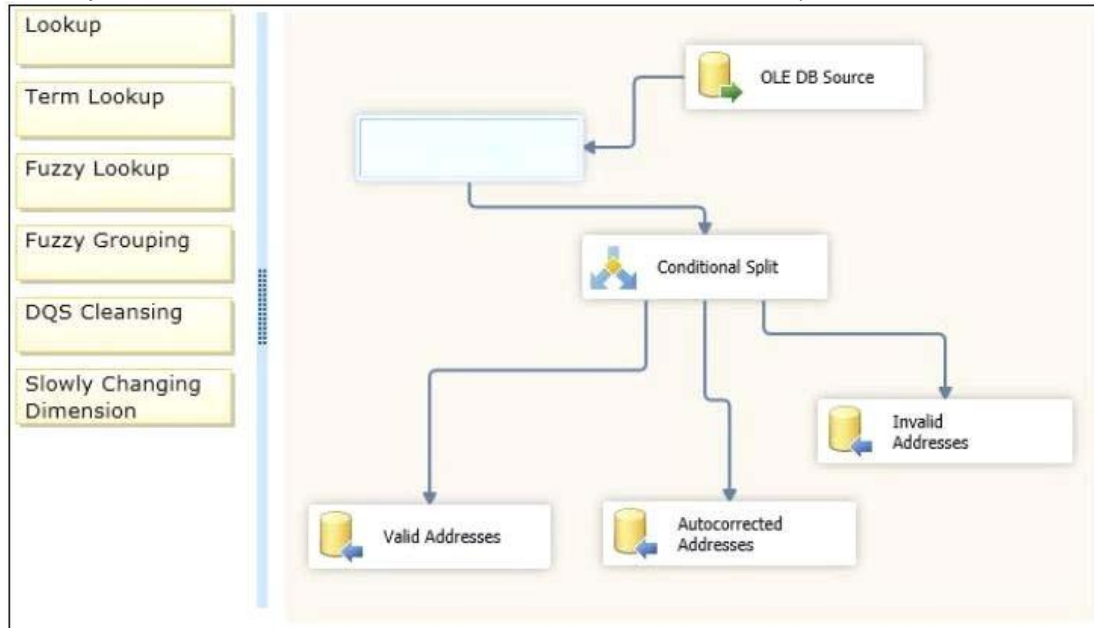
**QUESTION 9**

Drag and Drop Question

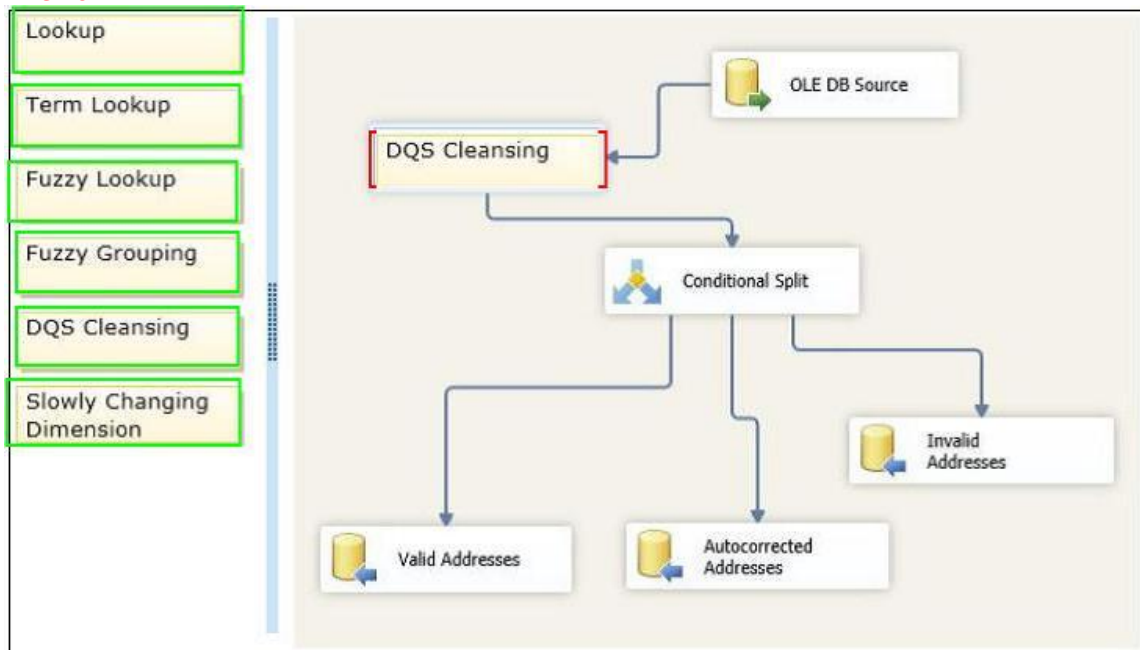
You use SQL Server Integration Services (SSIS) for extract, transformation, load (ETL) processing. Issues concerning addresses are discovered in the data warehouse that you manage. Validation must separate the addresses into three categories:

- Valid addresses
- Autocorrected addresses
- Invalid addresses

You need to enhance the SSIS packages to perform address validation by using an external service. Which transformation should you use? (To answer, drag the appropriate transformation from the list of options to the correct location or locations in the answer area.)



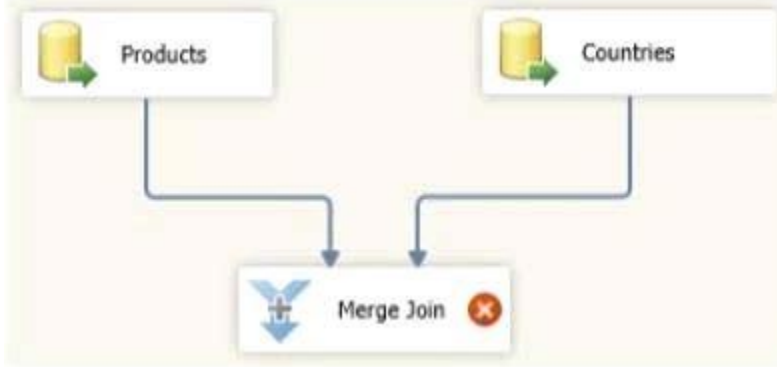
**Answer:**



#### QUESTION 10

You are developing a data flow transformation to merge two data sources. One source contains product data and the other source contains data about the country in which the product was manufactured. Both data sources contain a two-character CountryCode column and both use SQL

Server. Both data sources contain an ORDER BY clause to sort the data by the CountryCode column in ascending order. You use a Merge Join transformation to join the data. You need to ensure that the Merge Join transformation works correctly without additional transformations. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)



- A. Change the ORDER BY clause on the product source to order by ProductName.
- B. Change the Merge Join transformation to a Merge transformation.
- C. Set the appropriate SortKeyPosition properties on the data sources.
- D. Set the IsSorted property on both data sources.

**Answer: CD**

**QUESTION 11**

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:**

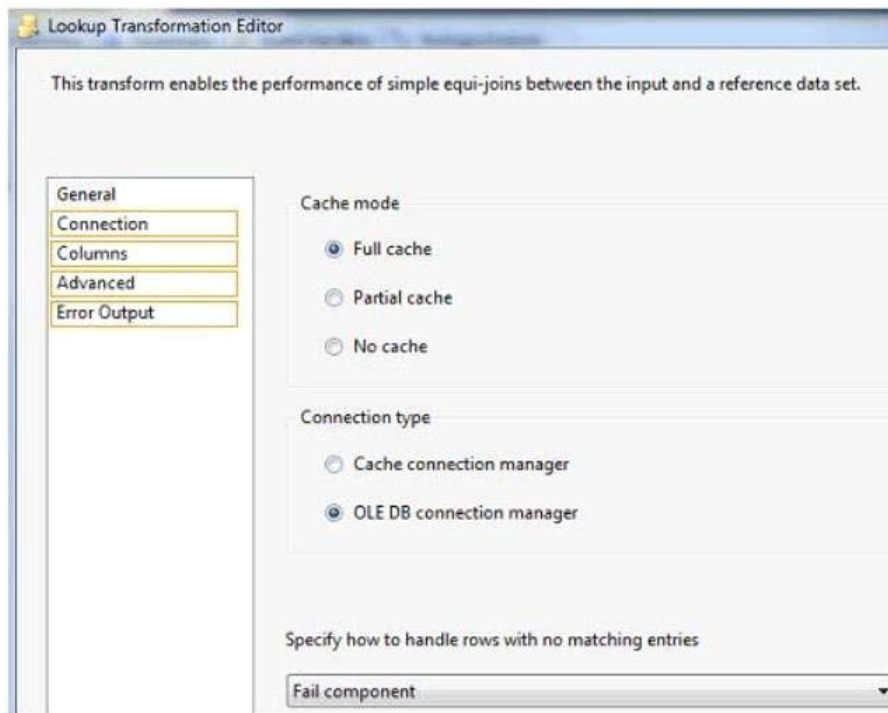
<http://msdn.microsoft.com/en-us/library/hh213137.aspx>

**QUESTION 12**

Hotspot Questions

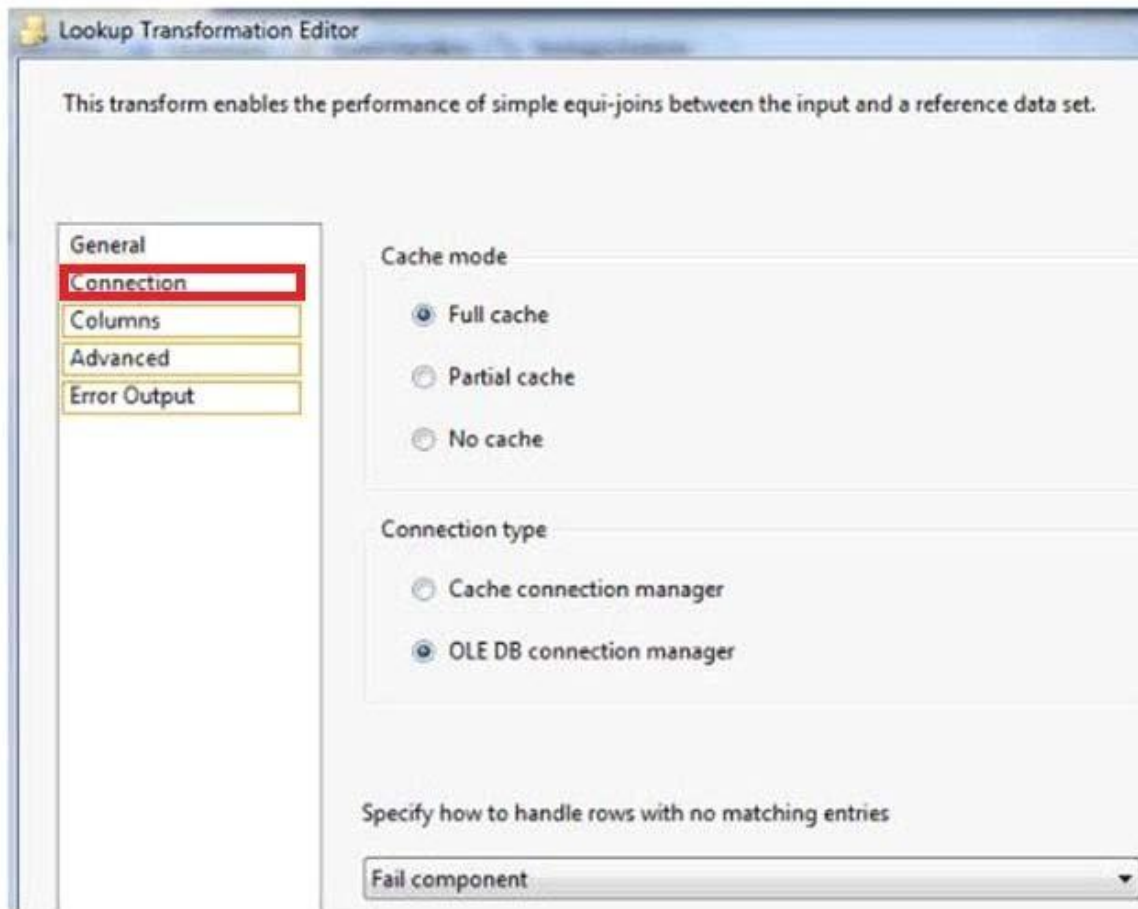
You are developing a data flow to load sales data into a fact table. In the data flow, you configure a Lookup Transformation in full cache mode to look up the product data for the sale. The lookup source for the product data is contained in two tables. You need to set the data source for the lookup to be a query that combines the two tables. Which page of the Lookup Transformation Editor should you select to configure the query? (To answer, select the appropriate page in the answer area.)





**Answer:**





**QUESTION 13**

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 correct answer presents a complete solution. Choose all that apply.)

- A. Multicast
- B. Merge Join
- C. Term Extraction
- D. union All
- E. Merge

**Answer: DE**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms141703.aspx>

<http://msdn.microsoft.com/en-us/library/ms141775.aspx>

<http://msdn.microsoft.com/en-us/library/ms141020.aspx>

<http://msdn.microsoft.com/en-us/library/ms141809.aspx>

<http://msdn.microsoft.com/en-us/library/ms137701.aspx>

**QUESTION 14**

[70-767 Exam Dumps](#) [70-767 Exam Questions](#) [70-767 PDF Dumps](#) [70-767 VCE Dumps](#)

[Back to the Source of this PDF & Get More Free Braindumps -- www.microsoftbraindumps.com](#)

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 flat file. Additional packages in the project must also access this file. You need to define and reuse the flat file connection in all project packages. What should you do?


- A. Convert the package Connection Manager in the first package to a project Connection Manager.
- B. Copy the package Connection Manager and paste it into the second package.
- C. Convert the project to the Package Deployment model.
- D. Set the ProtectionLevel property of the package Connection Manager to DontSaveSensitive to reuse the flat file connection.

**Answer: A**

### QUESTION 15

Drag and Drop Question

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.)

|   |  |  |
|---|--|--|
| Create an <b>OnError</b> event handler for the Data Flow task.  |  |  |
| Define an expression for the <b>ActionAtEvent</b> property.   |  |  |
| Add a Notify Operator task.   |  |  |
| Add a WMI Event Watcher task.   |  |  |
| In the <b>ReadOnlyVariables</b> list, select <b>System::ErrorDescription</b> and add conditional logic to invoke the <b>FireError</b> method. |  |  |
| Add a Script task.  |  |  |

**Answer:**

|   |   |
|---|---|
| Create an <b>OnError</b> event handler for the Data Flow task.  | Create an <b>OnError</b> event handler for the Data Flow task.  |
| Define an expression for the <b>ActionAtEvent</b> property.   |   |
| Add a Notify Operator task.   | Add a Script task.  |
| Add a WMI Event Watcher task.   |   |
| In the <b>ReadOnlyVariables</b> list, select <b>System::ErrorDescription</b> and add conditional logic to invoke the <b>FireError</b> method. | In the <b>ReadOnlyVariables</b> list, select <b>System::ErrorDescription</b> and add conditional logic to invoke the <b>FireError</b> method. |
| Add a Script task.  |   |

### QUESTION 16

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 does not change. You need to reuse the Fuzzy Lookup match index to increase performance and reduce maintenance. What should you do?

- A. Select the GenerateAndPersistNewIndex option in the Fuzzy Lookup Transformation Editor.
- B. Select the GenerateNewIndex option in the Fuzzy Lookup Transformation Editor.
- C. Select the DropExistingMatchIndex option in the Fuzzy Lookup Transformation Editor.
- D. Execute the sp\_FuzzyLookupTableMaintenanceUninstall stored procedure.

E. Execute the sp\_FuzzyLookupTableMaintenanceInvoke stored procedure.

**Answer: A**

**Explanation:**

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

**QUESTION 17**

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**

**QUESTION 18**

Drag and Drop Question

You are developing a SQL Server Integration Services (SSIS) package that imports data into a data warehouse. You are developing the part of the SSIS package that populates the ProjectDates dimension table. The business key of the ProjectDates table is the ProjectName column. The business user has given you the dimensional attribute behavior for each of the four columns in the ProjectDates table:

- ExpectedStartDate - New values should be tracked over time.
- ActualStartDate - New values should not be accepted.
- ExpectedEndDate - New values should replace existing values.
- ActualEndDate - New values should be tracked over time.

You use the SSIS Slowly Changing Dimension Transformation. You must configure the Change Type value for each source column. Which settings should you select? (To answer, select the appropriate setting or settings in the answer area.)

| Dimension Columns | Change Type |
|-------------------|-------------|
| ExpectedStartDate |             |
| ActualStartDate   |             |
| ExpectedEndDate   |             |
| ActualEndDate     |             |

**Answer:**

| Dimension Columns | Change Type          |
|-------------------|----------------------|
| ExpectedStartDate | Historical Attribute |
| ActualStartDate   | Fixed Attribute      |
| ExpectedEndDate   | Changing Attribute   |
| ActualEndDate     | Historical Attribute |

**QUESTION 19**

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:**

<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>

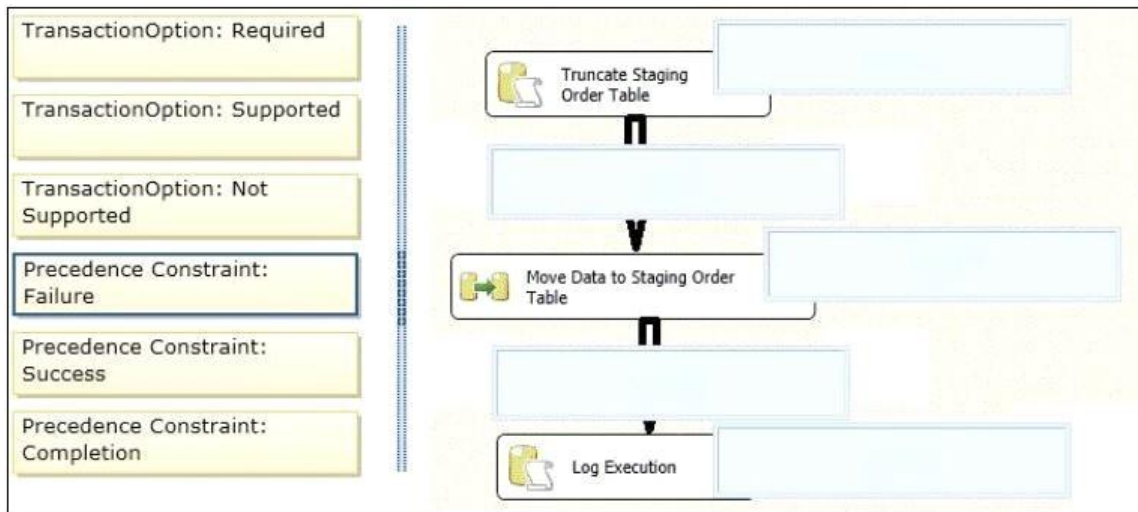
**QUESTION 20**

Drag and Drop Questions

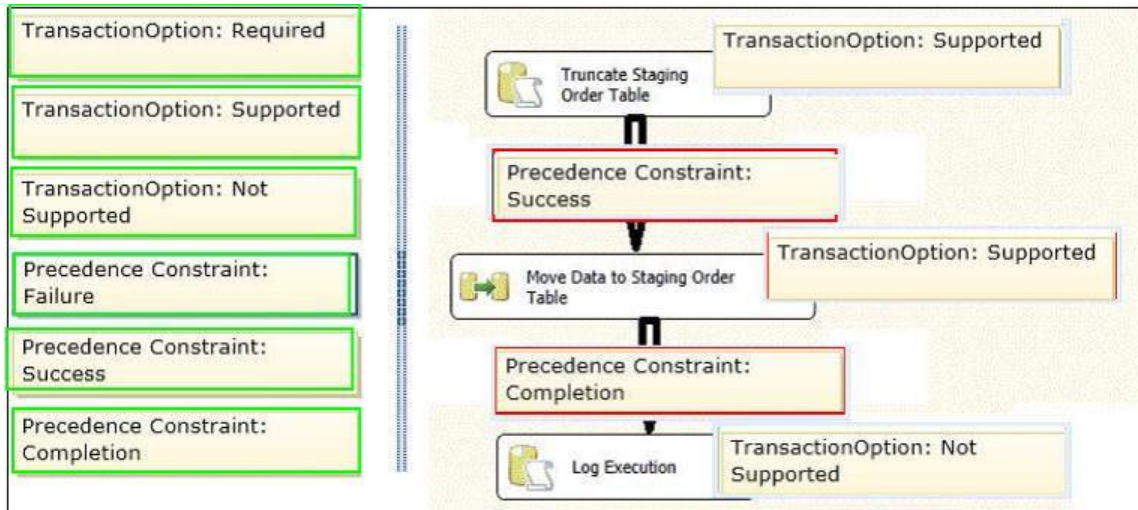
You are designing a SQL Server Integration Services (SSIS) package. The package moves order-related data to a staging table named Order. Every night the staging data is truncated and then all the recent orders from the online store database are inserted into the staging table. Your package must meet the following requirements:

- If the truncate operation fails, the package execution must stop and report an error.
- If the Data Flow task that moves the data to the staging table fails, the entire refresh operation must be rolled back.
- For auditing purposes, a log entry must be entered in a SQL log table after each execution of the Data Flow task.

The TransactionOption property for the package is set to Required. You need to design the package to meet the requirements. How should you design the control flow for the package? (To answer, drag the appropriate setting from the list of settings to the correct location or locations in the answer area.)



**Answer:**



**QUESTION 21**

Drag and Drop Question

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 <b>Foreach File Enumerator</b> option and configure the directory and files to process. |  |
| On the Collection page, select the <b>Foreach ADO Enumerator</b> 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:**

|  |  |
|--|--|
| Open the Control Flow page of the package.   | Open the Control Flow page of the package.   |
| On the Collection page, select the <b>Foreach File Enumerator</b> option and configure the directory and files to process. |  |
| On the Collection page, select the <b>Foreach ADO Enumerator</b> option and configure the ADO object source variable.      | Create a For Loop container.   |
| Create a File System task.   | On the Collection page, select the <b>Foreach File Enumerator</b> option and configure the directory and files to process. |
| Create a For Loop container.   |  |
| Create a Foreach Loop container.   |  |
| Open the Data Flow page of the package.  |  |

**Explanation:**

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

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

## QUESTION 22

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)`

**Answer: D**

**Explanation:**

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

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

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

**QUESTION 23**

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. 1. DimCurrency, DimOrganization in parallel  
2. DimScenario, DimAccount in parallel  
3. FactFinance
- C. 1. DimCurrency, FactFinance in parallel  
2. DimOrganization, DimScenario, DimAccount in parallel
- D. 1. FactFinance  
2. DimOrganization, DimScenario, DimAccount in parallel  
3. DimCurrency
- E. 1. DimCurrency  
2. DimOrganization  
3. DimScenario, DimAccount in parallel  
4. FactFinance

**Answer: A**

**Explanation:**

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

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

**QUESTION 24**

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. OnPostExecute
- B. OnWarning
- C. OnExecStatusChanged
- D. OnVariableValueChanged

**Answer: A**

**Explanation:**

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

**QUESTION 25**

You are editing a SQL Server Integration Services (SSIS) package that contains three Execute SQL tasks and no other tasks. The package and all three Execute SQL tasks have their TransactionOption property set to Supported. You need to ensure that if any of the Execute SQL tasks fail, all three tasks will roll back their changes. What should you do?

- A. Move the three Execute SQL tasks into a Sequence container.
- B. Move the three Execute SQL tasks into a Foreach Loop container.
- C. Change the TransactionOption property of all three Execute SQL tasks to Required.
- D. Change the TransactionOption property of the package to Required.

**Answer: D**

**Explanation:**

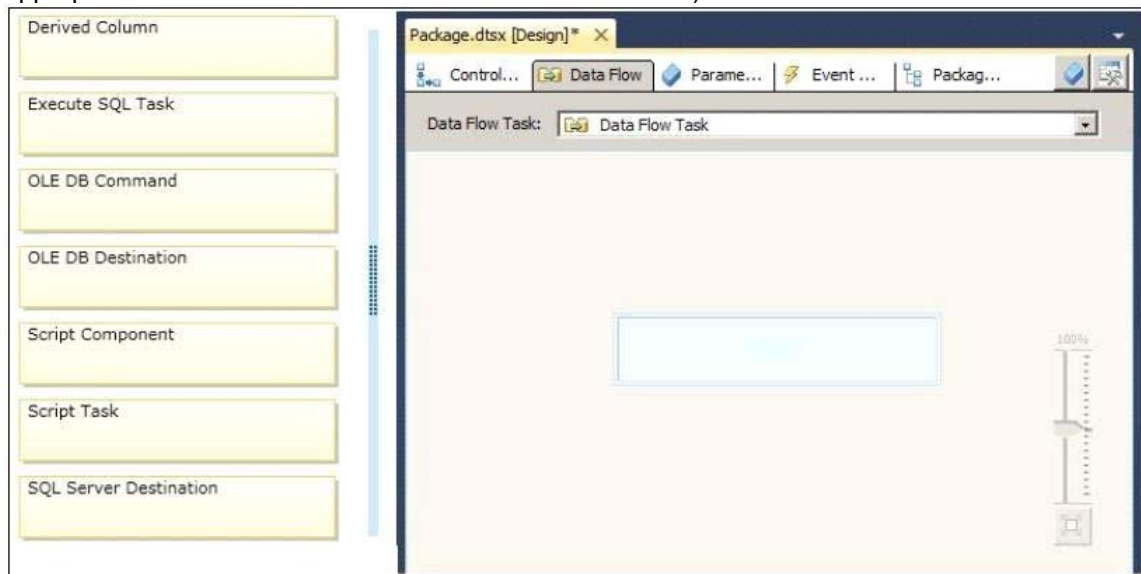
<http://msdn.microsoft.com/en-us/library/ms137749.aspx>

<http://msdn.microsoft.com/en-us/library/microsoft.sqlserver.dts.runtime.dtstransactionoption.aspx>

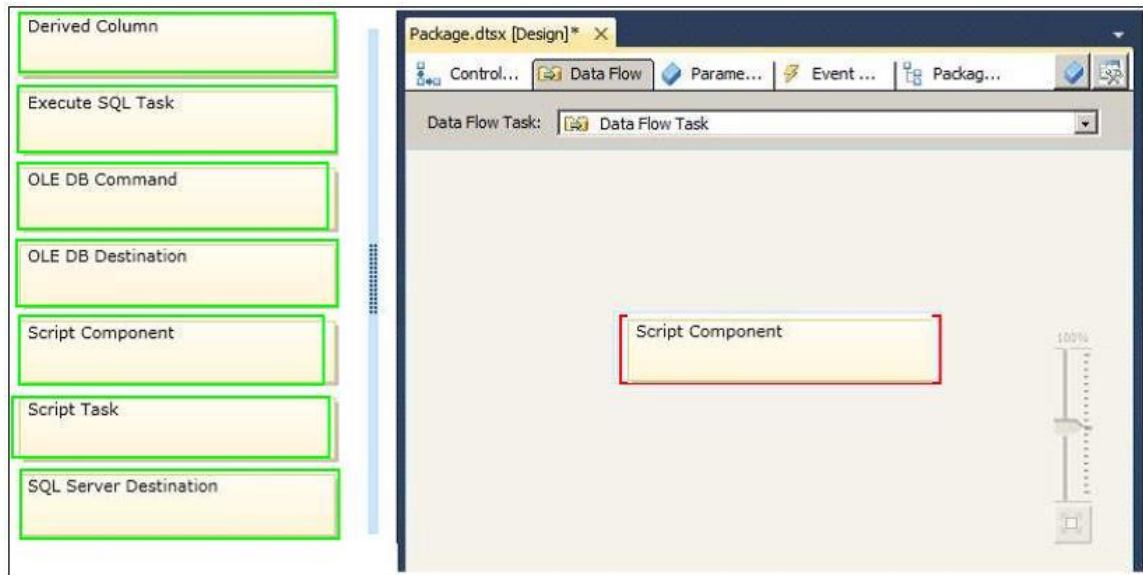
**QUESTION 26**

Drag and Drop Question

You are developing a SQL Server Integration Services (SSIS) package. The package uses custom functionality that accesses a SQL Server database. The custom functionality must be implemented by using Language Integrated Query (LINQ). You need to ensure that the LINQ code can be debugged at design time. What should you select from the SSIS Toolbox? (To answer, drag the appropriate item to the correct location in the answer area.)



**Answer:**



### QUESTION 27

#### Drag and Drop Questions

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 <b>Package A</b> and another SSIS project containing only <b>Package B</b> . | <div style="text-align: center;"> <div style="border: 1px solid #ccc; width: 40px; height: 40px; margin: 0 auto; display: flex; align-items: center; justify-content: center;"> <span style="font-size: 24px;">➤</span> </div> <div style="border: 1px solid #ccc; width: 40px; height: 40px; margin: 0 auto; display: flex; align-items: center; justify-content: center;"> <span style="font-size: 24px;">➤</span> </div> </div> |
| Add a <b>Parent Package Variable</b> configuration in <b>Package B</b> to set the <b>FilePath</b> variable.          |  |
| Add a package parameter named <b>FilePath</b> to <b>Package A</b> and <b>Package B</b> .                             |  |
| Add a configuration file in <b>Package A</b> to set the <b>FilePath</b> value.                                       |  |
| Configure the project to use an Environment that sets the <b>FilePath</b> parameter.                                 |  |
| Deploy the project to the SSIS catalog.  |  |
| Create an SSIS project containing <b>Package A</b> and <b>Package B</b> .  |  |
| Add a project parameter named <b>FilePath</b> .  |  |

**Answer:**



### QUESTION 28

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**

**Explanation:**

Type of Slowly Changing Dimensions Slowly Changing Dimensions are categorized into three types named: Type 1, Type 2, and Type 3. 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.

### QUESTION 29

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:**

<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>

**QUESTION 30**

Drag and Drop Question

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

**Answer:**



### QUESTION 31

You are installing the Data Quality Client on user desktops. You need to ensure that the prerequisite software components are installed. Which components must be present to meet this goal? (Each correct answer presents part of the solution. Choose all that apply.)

- A. SQL Server Management Studio
- B. Internet Explorer 6.0 SP1 or later
- C. Microsoft Silverlight 5
- D. .NET Framework 3.5 SP1
- E. .NET Framework 4.0
- F. Microsoft Silverlight 4
- G. SQL Server Data Tools

**Answer: BE**

#### **Explanation:**

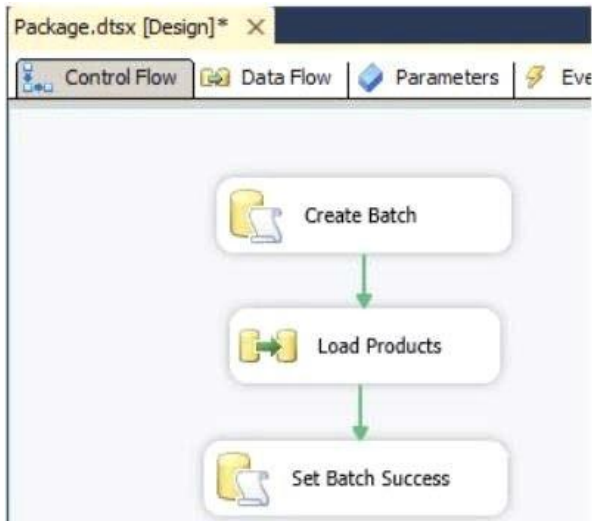
Before installing DQS, make sure that your computer meets the minimum system requirements. Data Quality Client requirement: NET Framework 4.0 (installed during the Data Quality Client installation, if not already installed) and Internet Explorer 6.0 SP1 or later.

### QUESTION 32

Drag and Drop Question

You are building a SQL Server Integration Services (SSIS) package to load product data sourced from a SQL Azure database to a data warehouse. Before the product data is loaded, you create a batch record by using an Execute SQL task named Create Batch. After successfully loading the product data, you use another Execute SQL task named Set Batch Success to mark the batch as successful.





You need to create and execute an Execute SQL task to mark the batch as failed if either the Create Batch or Load Products task fails. 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.)

|  |  |
|--|--|
| Set the <b>Multiple Constraints</b> option to <b>Logical OR</b> .  |  |
| Set the <b>Evaluation Operation</b> option to <b>Expression OR Constraint</b> .  |  |
| Connect the <b>Create Batch</b> and <b>Load Products</b> tasks to the <b>Set Batch Failure</b> task with failure constraints.    |  |
| Create an Execute SQL task named <b>Set Batch Failure</b> to mark the batch as failed.   |  |
| Set the <b>Multiple Constraints</b> option to <b>Logical AND</b> .   |  |
| Connect the <b>Create Batch</b> and <b>Load Products</b> tasks to the <b>Set Batch Failure</b> task with completion constraints. |  |
| Set the <b>ForceExecutionValue</b> property of the <b>Create Batch</b> task to <b>True</b> .                                     |  |

**Answer:**

|  |   |
|--|---|
| Set the <b>Multiple Constraints</b> option to <b>Logical OR</b> .  | Create an Execute SQL task named <b>Set Batch Failure</b> to mark the batch as failed.  |
| Set the <b>Evaluation Operation</b> option to <b>Expression OR Constraint</b> .  |   |
| Connect the <b>Create Batch</b> and <b>Load Products</b> tasks to the <b>Set Batch Failure</b> task with failure constraints.    | Connect the <b>Create Batch</b> and <b>Load Products</b> tasks to the <b>Set Batch Failure</b> task with failure constraints. |
| Create an Execute SQL task named <b>Set Batch Failure</b> to mark the batch as failed.   |   |
| Set the <b>Multiple Constraints</b> option to <b>Logical AND</b> .   | Set the <b>Multiple Constraints</b> option to <b>Logical OR</b> .   |
| Connect the <b>Create Batch</b> and <b>Load Products</b> tasks to the <b>Set Batch Failure</b> task with completion constraints. |   |
| Set the <b>ForceExecutionValue</b> property of the <b>Create Batch</b> task to <b>True</b> .                                     |   |

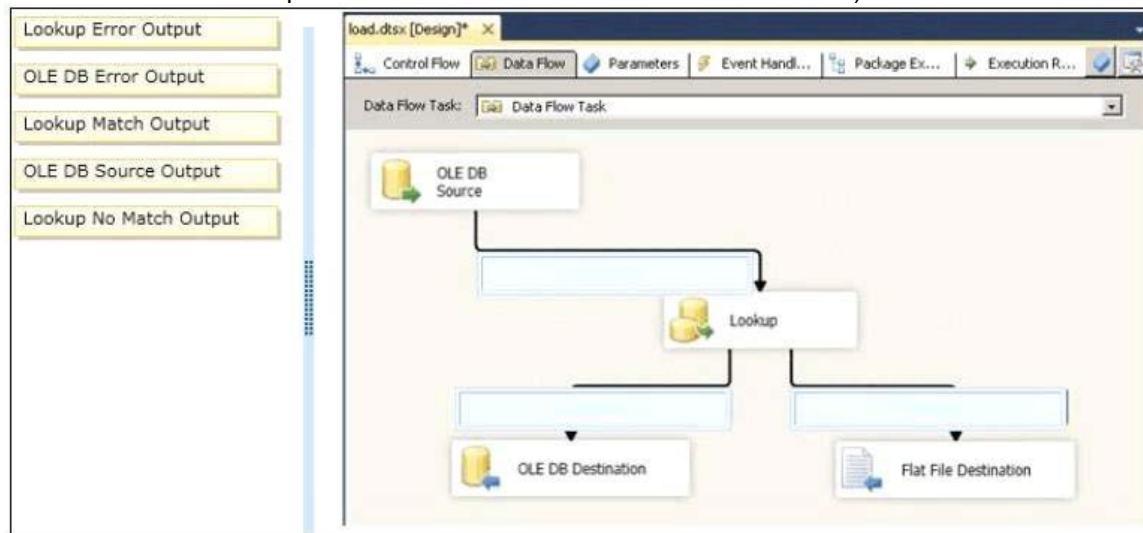
### QUESTION 33

Drag and Drop Questions

You are developing a SQL Server Integration Services (SSIS) package to insert new data into a data mart. The package uses a Lookup transformation to find matches between the source and destination. The data flow has the following requirements:

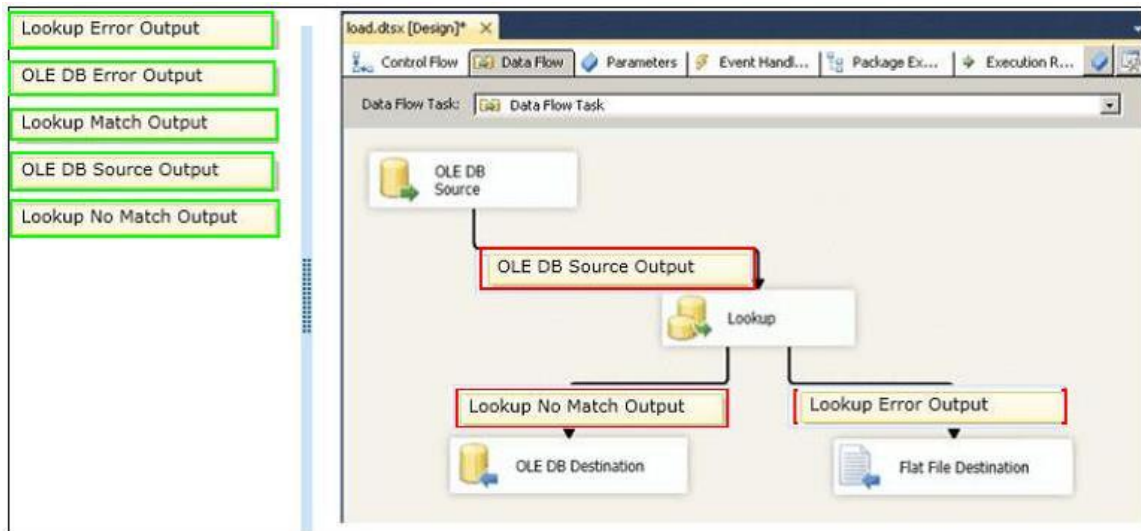
- New rows must be inserted.
- Lookup failures must be written to a flat file.

In the Lookup transformation, the setting for rows with no matching entries is set to Redirect rows to no match output. You need to configure the package to direct data into the correct destinations. How should you design the data flow outputs? (To answer, drag the appropriate transformation from the list of answer options to the correct location in the answer area.)



**Answer:**





**Explanation:**

Records that have no match in the destination are new records and so should be inserted. Any errors should be recorded in the flat file.

<http://msdn.microsoft.com/en-us/library/ms141821.aspx>

<http://msdn.microsoft.com/en-us/library/bb895366.aspx>

<https://www.simple-talk.com/sql/ssis/implementing-lookup-logic-in-sql-server-integration-services/>

**QUESTION 34**

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:**

<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>

**QUESTION 35**

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. Run the DQSInstaller.exe command.

[70-767 Exam Dumps](#) [70-767 Exam Questions](#) [70-767 PDF Dumps](#) [70-767 VCE Dumps](#)

**[Back to the Source of this PDF & Get More Free Braindumps -- www.microsoftbraindumps.com](http://www.microsoftbraindumps.com)**

- B. Install the data providers that are used for data refresh.
- C. Install ADOMD.NET.
- D. Run the dbimpexp.exe command.

**Answer: A**

**Explanation:**

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

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

**QUESTION 36**

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**

**QUESTION 37**

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:**

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

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

<http://msdn.microsoft.com/en-us/library/microsoft.masterdataservices.services.datacontracts.knowledgebase.aspx>

<http://msdn.microsoft.com/en-us/library/bb895263.aspx>

**QUESTION 38**

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 tables.

**Answer: A**

**QUESTION 39**

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

- A. Data Quality Client with administrative credentials
- B. SQL Server Configuration Manager with local administrative credentials
- C. SQL Server Data Tools with local administrative permissions
- D. SQL Server Management Studio with administrative credentials

**Answer: D**

**QUESTION 40**

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 Team Type, does not yet have a domain. You need to complete the mapping of the Team Type column. What should you do?

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

**Answer: D**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ff877917.aspx>  
<http://msdn.microsoft.com/en-us/library/hh213015.aspx>  
<http://msdn.microsoft.com/en-us/library/gg524799.aspx>

**QUESTION 41**

You are maintaining a Data Quality Services (DQS) environment. The production server failed and a new server has been set up. The DQS databases are restored to a new server. All the appropriate permissions are granted. DQS users are experiencing issues connecting to the new Data Quality Server. You need to enable users to connect to the new server. Which Surface Area Configuration property should you enable?

- A. AdHocRemoteQueriesEnabled
- B. SoapEndpointsEnabled
- C. CrlIntegrationEnabled
- D. RemoteDacEnabled
- E. OleAutomationEnabled
- F. XpCmdShellEnabled

**Answer: C**

**QUESTION 42**

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:**

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

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

**QUESTION 43**

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:**

<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>

**QUESTION 44**

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 data files. 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. Convert the parameters to package parameters.
- B. Set the Sensitive property of the parameters to True.
- C. Set the ProtectionLevel property of the package to EncryptSensitiveWithPassword.
- D. Convert the project to the Legacy Deployment model.

**Answer: B**

**QUESTION 45**

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 Log.

**Answer: A**

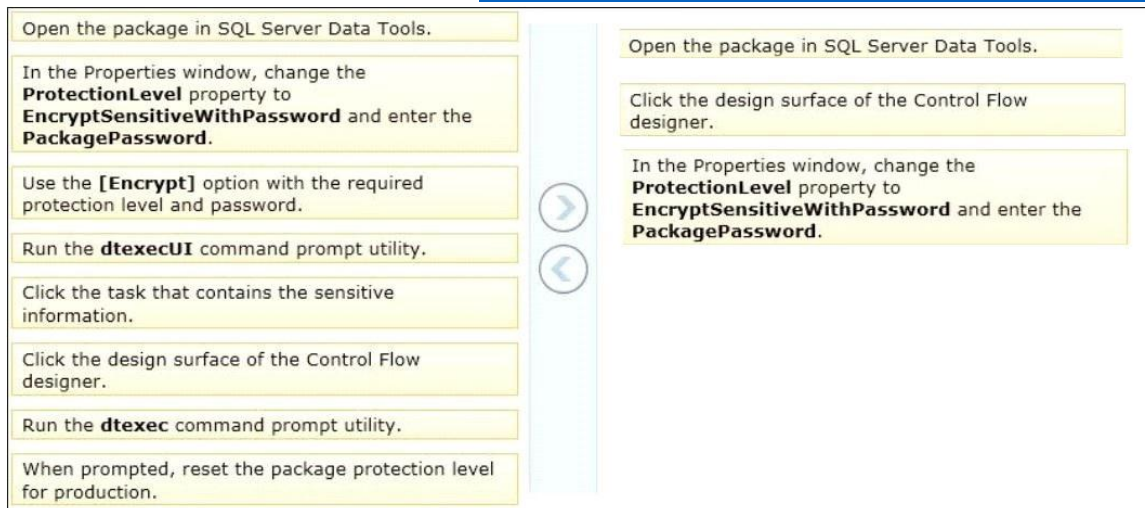
**QUESTION 46**

Drag and Drop Question

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 <b>ProtectionLevel</b> property to <b>EncryptSensitiveWithPassword</b> and enter the <b>PackagePassword</b> . |   |
| Use the <b>[Encrypt]</b> option with the required protection level and password.   |   |
| Run the <b>dtexecUI</b> command prompt utility.  |   |
| Click the task that contains the sensitive information.  |   |
| Click the design surface of the Control Flow designer.   |   |
| Run the <b>dtexec</b> command prompt utility.  |   |
| When prompted, reset the package protection level for production.  |   |

**Answer:**



#### QUESTION 47

You install a SQL Server 2012 database engine instance on a production server. A month later, you install SQL Server 2012 Integration Services (SSIS). You must develop an SSIS project and deploy it to the server by using the Project Deployment model. Operations log records that are outside the configured retention period must be cleaned automatically. You need to create the SSIS catalog on the production server and ensure that the operations log cleaning requirement is met. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Enable CLR Integration.
- B. Enable FILESTREAM with Full Access.
- C. Enable the Resource Governor.
- D. Change the recovery mode of the msdb database to FULL.
- E. Change the Server-wide Default Logging Level in SSISDB to Verbose.
- F. Start the SQL Server Browser service.
- G. Start the SQL Server Agent service.

**Answer: AG**

#### QUESTION 48

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 Configuration Manager
- B. SQL Server Agent
- C. SSIS Upgrade Wizard in SQL Server Management Studio
- D. SQL Server DTEXECUI utility (dtexecui.exe)

**Answer: C**

#### Explanation:

Use the SSIS Package Upgrade Wizard to migrate packages that were developed in SQL Server 2005 Integration Services (SSIS) and SQL Server 2008 Integration Services (SSIS) to the package format that the current release of Integration Services uses.

[https://technet.microsoft.com/en-us/library/cc280629\(v=sql.110\).aspx](https://technet.microsoft.com/en-us/library/cc280629(v=sql.110).aspx)



**QUESTION 49**

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:**

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

**QUESTION 50**

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 Project 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:**

<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>

**QUESTION 51**

To ease the debugging of packages, you standardize the SQL Server Integration Services (SSIS) package logging methodology. The methodology has the following requirements:

- Centralized logging in SQL Server
- Simple deployment



- Availability of log information through reports or T-SQL
- Automatic purge of older log entries
- Configurable log details

You need to configure a logging methodology that meets the requirements while minimizing the amount of deployment and development effort. What should you do?

- A. Deploy the package by using an msi file.
- B. Use the gacutil command.
- C. Create an OnError event handler.
- D. Create a reusable custom logging component.
- E. Use the dtutil/copy command.
- F. Use the Project Deployment Wizard.
- G. Run the package by using the dtexec/rep/conn command.
- H. Add a data tap on the output of a component in the package data flow.
- 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: I**

**Explanation:**

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

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

#### **QUESTION 52**

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:**

<http://technet.microsoft.com/en-us/library/ff487006.aspx>

<http://msdn.microsoft.com/en-us/library/ee633747.aspx>

<http://technet.microsoft.com/en-us/library/ee633759.aspx>

#### **QUESTION 53**

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. Change the Exhaustive property to True.

- 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:**

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

**QUESTION 54**

You are implementing a SQL Server Integration Services (SSIS) package that loads data hosted in a SQL Azure database into a data warehouse. The source system contains redundant or inconsistent data. When the package finds invalid data, the row containing the invalid data must be omitted but it must also be written to a text file for further analysis. You need to establish the best technique to log these invalid rows while keeping the amount of development effort to a minimum. What should you do?

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

**Answer: I**

**Explanation:**

<http://technet.microsoft.com/en-us/library/hh230989.aspx>

<http://www.rafael-salas.com/2012/01/ssis-2012-quick-peek-to-data-taps.html>

<http://msdn.microsoft.com/en-us/library/ms162820.aspx>

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

<http://technet.microsoft.com/en-us/library/ms140223.aspx>

<http://msdn.microsoft.com/en-us/library/jj655339.aspx>

**QUESTION 55**

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 Fuzzy Lookup transform.
- B. Create a Master Data Services (MDS) model.
- C. Create a Data Quality Project.
- D. Create an Integration Services package and use a DQS Cleansing transform.
- E. Create an Integration Services package and use a Fuzzy Grouping transform.

**Answer: C**

**Explanation:**

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

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

[http://msdn.microsoft.com/en-](http://msdn.microsoft.com/en-us/library/microsoft.masterdataservices.services.datacontracts.knowledgebase.aspx)

[us/library/microsoft.masterdataservices.services.datacontracts.knowledgebase.aspx](http://msdn.microsoft.com/en-us/library/microsoft.masterdataservices.services.datacontracts.knowledgebase.aspx)

<http://msdn.microsoft.com/en-us/library/bb895263.aspx>

**QUESTION 56**

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**

**QUESTION 57**

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**

**QUESTION 58**

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 MDS.

**Answer: D**

**QUESTION 59**

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.NET.

**Answer: C**

**Explanation:**

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

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

**QUESTION 60**

You are editing a SQL Server Integration Services (SSIS) package that contains three Execute SQL tasks and no other tasks. The three Execute SQL tasks modify products in staging tables in preparation for a data warehouse load. The package and all three Execute SQL product tasks have their TransactionOption property set to Supported. You need to ensure that if any of the three Execute SQL product tasks fail, all three tasks will roll back their changes. What should you do?

- A. Change the TransactionOption property of the package to Required.
- B. Change the TransactionOption property of all three Execute SQL product tasks to Required.
- C. Move the three Execute SQL product tasks into a Foreach Loop container.
- D. Move the three Execute SQL product tasks into a Sequence container.

**Answer: A**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms137690.aspx>

<http://msdn.microsoft.com/en-us/library/ms141144.aspx>

**QUESTION 61**

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?

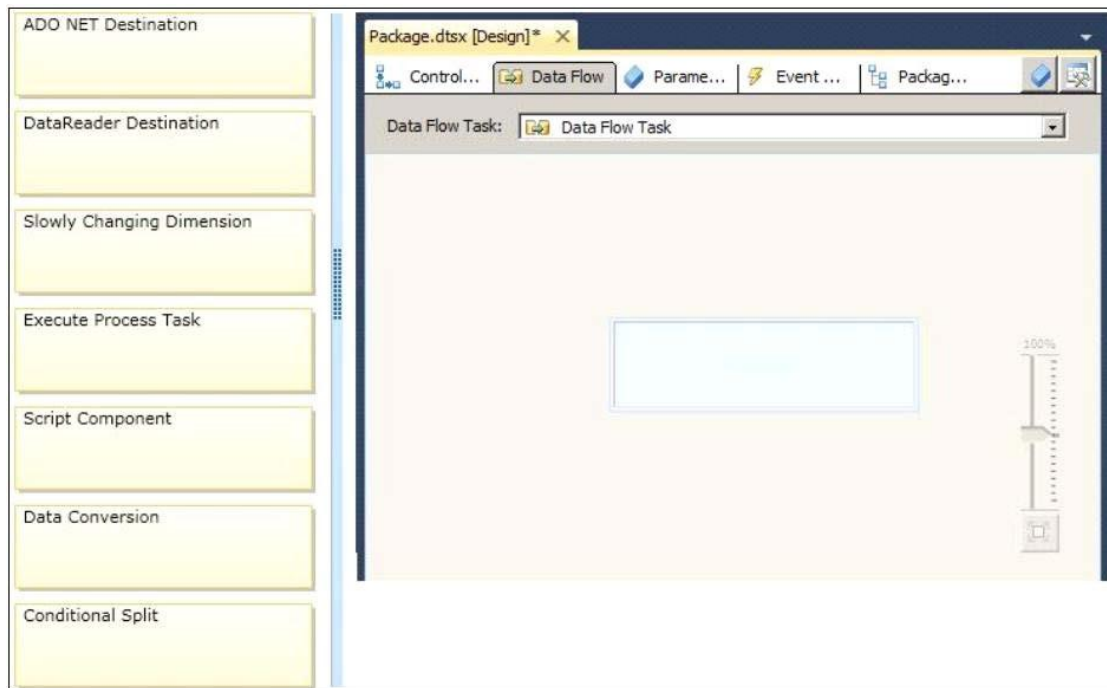
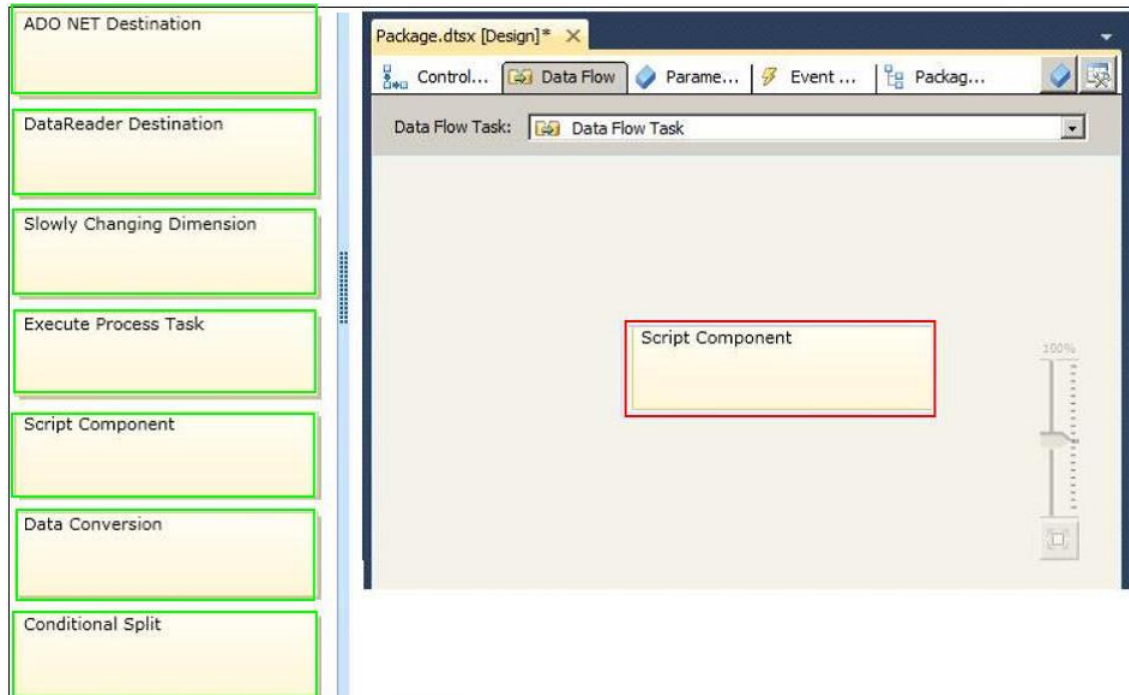
- A. Parent Package Variable
- B. XML Configuration File
- C. Environment Variable
- D. .dtsConfig file
- E. Registry Entry

**Answer: C**

**QUESTION 62**

Drag and Drop Question

You are developing a SQL Server Integration Services (SSIS) package. The package uses custom functionality that accesses a SQL Server database. The custom functionality must be implemented by using Language Integrated Query (LINQ). You need to ensure that the LINQ code can be debugged at design time. What should you select from the SSIS Toolbox? (To answer, drag the appropriate item to the correct location in the answer area.)


**Answer:****QUESTION 63**

Drag and Drop Questions


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.)

|   |   |
|---|---|
| In the SSIS catalog, create an Environment named <b>EnvTest</b> .   |  |
| In the <b>EnvTest</b> Environment, create an Environment Variable. Name the Environment Variable <b>EnvServer</b> . Set the Environment Variable value to <b>SQLTest1</b> . |   |
| In the SQL Agent job, choose the <b>EnvTest</b> Environment in the Configuration tab of the SQL Server Integration Services job step.                                       |   |
| In the SSIS catalog, create a folder named <b>EnvTest</b> .   |   |
| In the project configuration, create an Environment reference to <b>EnvTest</b> . Assign the value of the <b>EnvServer</b> Environment Variable to the project parameter.   |   |
| In the project configuration, change the value of the project parameter to <b>EnvTest</b> .   |   |

**Answer:**

|   |   |   |
|---|---|---|
| In the SSIS catalog, create an Environment named <b>EnvTest</b> .   |  | In the SSIS catalog, create an Environment named <b>EnvTest</b> .   |
| In the <b>EnvTest</b> Environment, create an Environment Variable. Name the Environment Variable <b>EnvServer</b> . Set the Environment Variable value to <b>SQLTest1</b> . |   | In the <b>EnvTest</b> Environment, create an Environment Variable. Name the Environment Variable <b>EnvServer</b> . Set the Environment Variable value to <b>SQLTest1</b> . |
| In the SQL Agent job, choose the <b>EnvTest</b> Environment in the Configuration tab of the SQL Server Integration Services job step.                                       |   | In the project configuration, create an Environment reference to <b>EnvTest</b> . Assign the value of the <b>EnvServer</b> Environment Variable to the project parameter.   |
| In the SSIS catalog, create a folder named <b>EnvTest</b> .   |   | In the SQL Agent job, choose the <b>EnvTest</b> Environment in the Configuration tab of the SQL Server Integration Services job step.                                       |
| In the project configuration, create an Environment reference to <b>EnvTest</b> . Assign the value of the <b>EnvServer</b> Environment Variable to the project parameter.   |   |   |
| In the project configuration, change the value of the project parameter to <b>EnvTest</b> .   |   |   |

#### QUESTION 64

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**

#### QUESTION 65

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 used.
- Default constraints on the destination table must be ignored.
- Batch size must be 100,000 rows.

You need to add a destination and configure it to meet the requirements. Which destination should you use?

- A. OLE DB Destination with Fast Load
- B. SQL Server Destination
- C. ADO NET Destination without Bulk Insert
- D. ADO NET Destination with Bulk Insert
- E. OLE DB Destination without Fast Load

**Answer: A**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms141237.aspx>

<http://msdn.microsoft.com/en-us/library/ms139821.aspx>

<http://msdn.microsoft.com/en-us/library/ms141095.aspx>

#### **QUESTION 66**

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 ($Project::StringVar)`
- C. `INSERT INTO dbo.Table (variablevalue) VALUES (?)`
- D. `INSERT INTO dbo.Table (variablevalue) VALUES ($Package::StringVar)`

**Answer: C**

#### **QUESTION 67**

Drag and Drop Question

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.

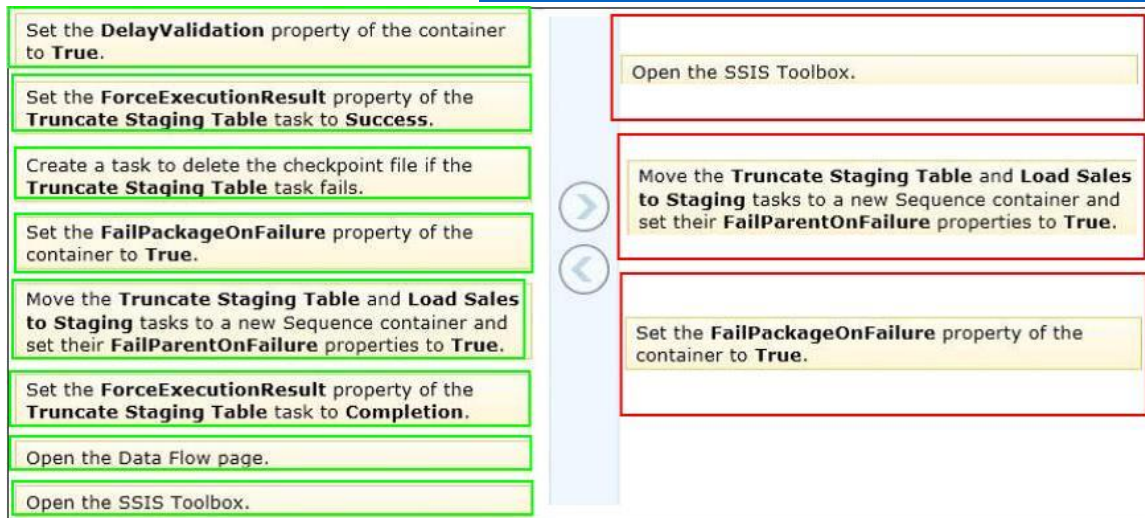




The checkpoints are currently working such that if any of the four steps fail, the package will restart from the failed step the next time it executes. You need to modify the package to ensure that if either the Truncate Staging Table or the Load Sales to Staging task fails, the package will always restart 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.)

- Set the **DelayValidation** property of the container to **True**.
- Set the **ForceExecutionResult** property of the **Truncate Staging Table** task to **Success**.
- Create a task to delete the checkpoint file if the **Truncate Staging Table** task fails.
- Set the **FailPackageOnFailure** property of the container to **True**.
- Move the **Truncate Staging Table** and **Load Sales to Staging** tasks to a new Sequence container and set their **FailParentOnFailure** properties to **True**.
- Set the **ForceExecutionResult** property of the **Truncate Staging Table** task to **Completion**.
- Open the Data Flow page.
- Open the SSIS Toolbox.

**Answer:**



### QUESTION 68

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.

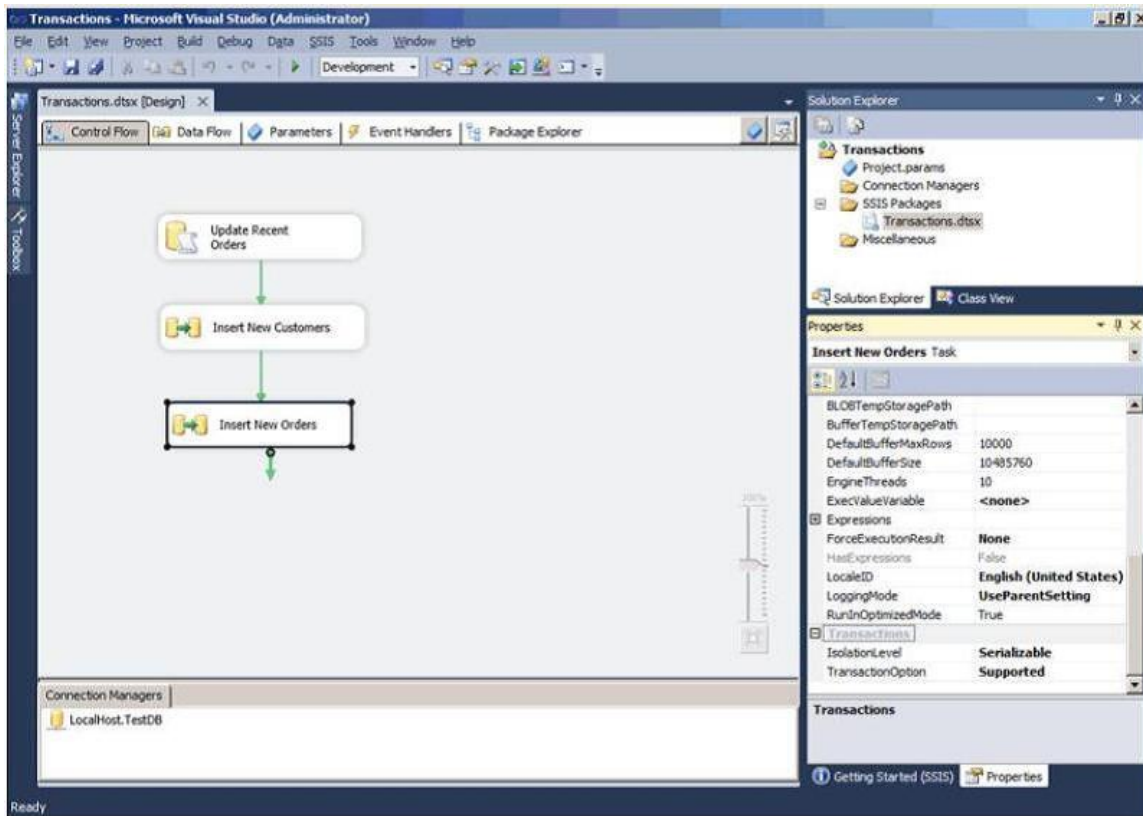
You need to meet the company's policy requirements. Which event handler should you use?

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

**Answer: C**

### QUESTION 69

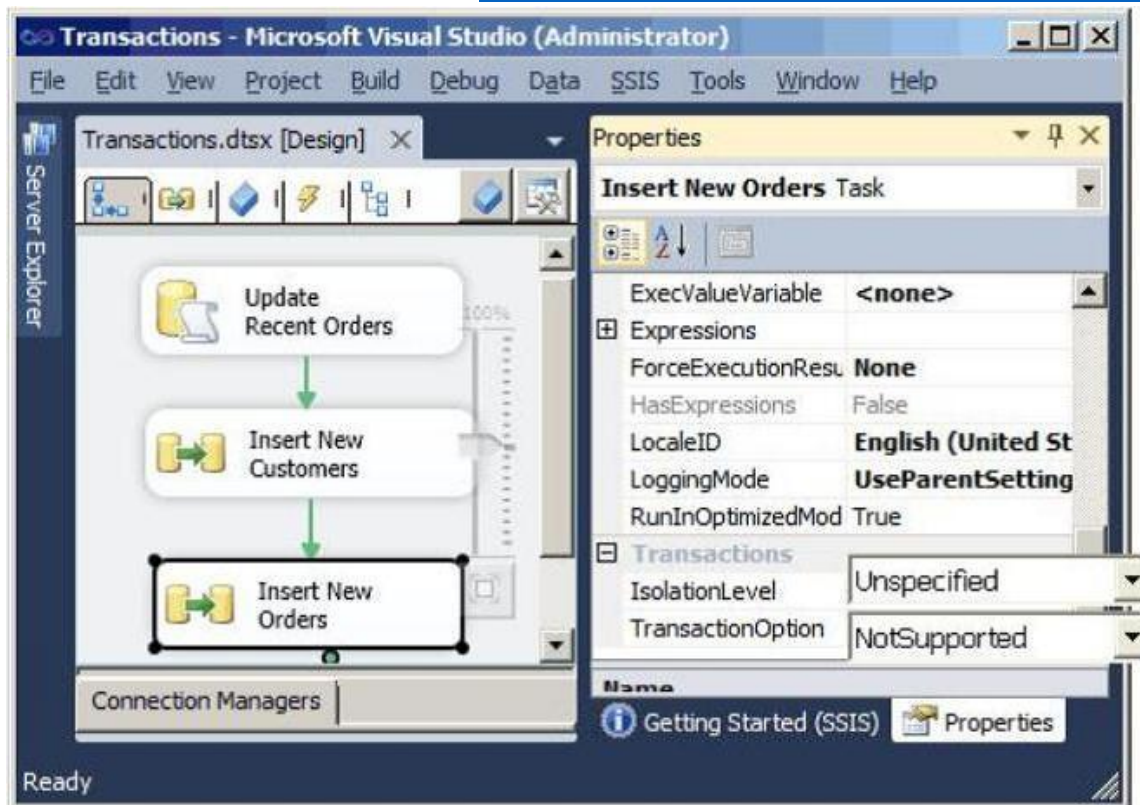
You are designing a package control flow. The package moves sales order data from a SQL Azure transactional database to an on-premise reporting database. The package will run several times a day, while new sales orders are being added to the transactional database. The current design of the package control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



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:**

IsolationLevel = ReadCommitted  
TransactionOption = Required

**QUESTION 70**

Drag and Drop Question

You are building a fact table in a data warehouse. The table must have a columnstore index. The table cannot be partitioned. You need to design the fact table and load it with data. 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.)

|  |  |
|--|--|
| Enable xVelocity on the server.  | <div style="display: flex; align-items: center; justify-content: center;"> <div style="border: 1px solid #ccc; padding: 5px; margin: 5px;"> <div style="border: 1px solid #ccc; padding: 2px; margin-bottom: 2px;"> <div style="background-color: #f0f0f0; padding: 2px;">Enable xVelocity on the server.</div> <div style="background-color: #f0f0f0; padding: 2px;">Create the columnstore index.</div> <div style="background-color: #f0f0f0; padding: 2px;">Load the data.</div> <div style="background-color: #f0f0f0; padding: 2px;">Create the table.</div> <div style="background-color: #f0f0f0; padding: 2px;">Set the <b>Allow Snapshot Isolation</b> database property to <b>True</b>.</div> </div> <div style="margin: 0 10px;"> <div style="border: 1px solid #ccc; border-radius: 50%; width: 30px; height: 30px; display: flex; align-items: center; justify-content: center; margin: 5px;"> <div style="font-size: 12px;">&gt;</div> </div> <div style="border: 1px solid #ccc; border-radius: 50%; width: 30px; height: 30px; display: flex; align-items: center; justify-content: center; margin: 5px;"> <div style="font-size: 12px;">&lt;</div> </div> </div> <div style="border: 1px solid #ccc; width: 150px; height: 150px; margin: 5px;"></div> </div> </div> |
| Create the columnstore index.  |  |
| Load the data.   |  |
| Create the table.  |  |
| Set the <b>Allow Snapshot Isolation</b> database property to <b>True</b> . |  |

**Answer:**

|  |   |                               |
|--|---|-------------------------------|
| Enable xVelocity on the server.  |  | Create the table.             |
| Create the columnstore index.  |   | Load the data.                |
| Load the data.   |   | Create the columnstore index. |
| Create the table.  |   |                               |
| Set the <b>Allow Snapshot Isolation</b> database property to <b>True</b> . |   |                               |

### QUESTION 71

#### Drag and Drop Questions

You are creating a sales data warehouse. When a product exists in the product dimension, you update the product name. When a product does not exist, you insert a new record. In the current implementation, the DimProduct table must be scanned twice, once for the insert and again for the update. As a result, inserts and updates to the DimProduct table take longer than expected. You need to create a solution that uses a single command to perform an update and an insert. How should you use a MERGE T-SQL statement to accomplish this goal? (To answer, drag the appropriate answer choice from the list of options to the correct location or locations in the answer area. You may need to drag the split bar between panes or scroll to view content.)

|  |  |
|--|--|
| Source   |  |
| Target   |  |
| Product  |  |
| Inserted   |  |
| MATCHED THEN   |  |
| NOT MATCHED BY TARGET THEN   |  |
| NOT MATCHED BY PRODUCT AND (Target.ProductID is NULL) THEN                                 |  |
| MATCHED AND (Target.ProductID is NULL) THEN  |  |
| MERGE dbo.DimProduct AS <input type="text"/>   |  |
| USING (SELECT ProductID, ProductName, ProductColor, ProductCategory                        |  |
| FROM dbo.StagingProduct) AS <input type="text"/>   |  |
| ON (Target.ProductID = Source.ProductID)   |  |
| WHEN <input type="text"/>  |  |
| UPDATE SET Target.ProductName = Source.ProductName   |  |
| WHEN <input type="text"/>  |  |
| INSERT (ProductID, ProductName, ProductColor, ProductCategory)                             |  |
| VALUES (Source.ProductID, Source.ProductName, Source.ProductColor, Source.ProductCategory) |  |
| OUTPUT Saction, Inserted.*, Deleted.*;   |  |

**Answer:**



|  |
|--|
| Source   |
| Target   |
| Product  |
| Inserted   |
| MATCHED THEN   |
| NOT MATCHED BY TARGET THEN                                 |
| NOT MATCHED BY PRODUCT AND (Target.ProductID is NULL) THEN |
| MATCHED AND (Target.ProductID is NULL) THEN                |

```

MERGE dbo.DimProduct AS Target
USING (SELECT ProductID, ProductName, ProductColor, ProductCategory
FROM dbo.StagingProduct) AS Source
ON (Target.ProductID = Source.ProductID)
WHEN MATCHED THEN
    UPDATE SET Target.ProductName = Source.ProductName
WHEN NOT MATCHED BY TARGET THEN
    INSERT (ProductID, ProductName, ProductColor, ProductCategory)
    VALUES (Source.ProductID, Source.ProductName, Source.ProductColor, Source.ProductCategory)
OUTPUT $action, Inserted.*, Deleted.*;
```

### QUESTION 72

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. Connect factSales to dimProduct and factProductTarget to dimProductCategory with foreign key constraints. Direct the cube developer to use key granularity attributes.
- B. Create one product table, dimProduct, which contains product detail, category, and subcategory columns. Connect factSales to dimProduct with a foreign key constraint. Direct the cube developer to use a non-key granularity attribute for factProductTarget.
- C. 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 constraints. Direct the cube developer to use key granularity attributes.
- D. Create three product tables, dimProduct, dimProductCategory, and dimProductSubcategory. Connect factSales to all three product tables and connect factProductTarget to dimProductCategory with foreign key constraints. Direct the cube developer to use key granularity attributes.

**Answer: B**

**QUESTION 73**

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.

**QUESTION 74**

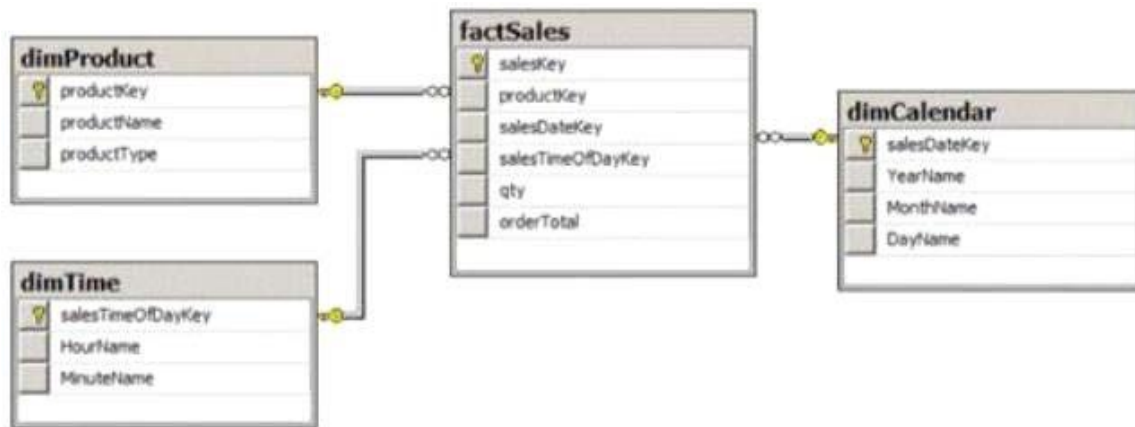
You are designing an enterprise star schema that will consolidate data from three independent data marts. One of the data marts is hosted on SQL Azure. Most of the dimensions have the same structure and content. However, the geography dimension is slightly different in each data mart. You need to design 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 conformed dimension for the geography dimension.
- B. Implement change tracking.
- C. Create a degenerate dimension for the geography dimension.
- D. Create a Type 2 slowly changing dimension for the geography dimension.

**Answer: A**

**QUESTION 75**

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 productKey.

**Answer: C**

#### QUESTION 76

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:

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

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

#### QUESTION 77

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**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms141703.aspx>

<http://msdn.microsoft.com/en-us/library/ms141775.aspx>

<http://msdn.microsoft.com/en-us/library/ms141020.aspx>

<http://msdn.microsoft.com/en-us/library/ms141809.aspx>

<http://msdn.microsoft.com/en-us/library/ms137701.aspx>


**QUESTION 78**

Drag and Drop Question

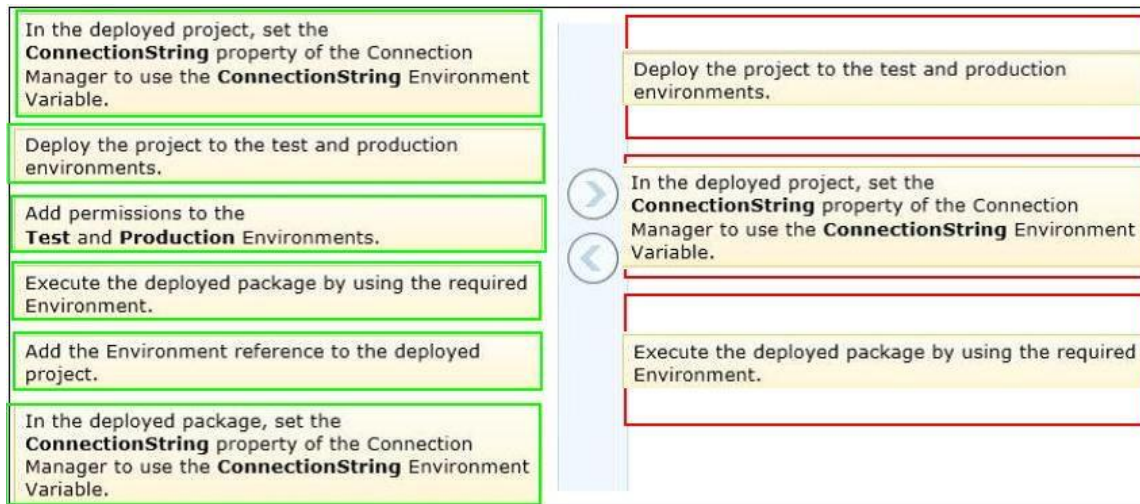
A SQL Server Integration Services (SSIS) project has been deployed to the SSIS catalog. The project includes a project Connection Manager to connect to the data warehouse. The SSIS catalog includes two Environments:

- Test
- Production

Each Environment defines a single Environment Variable named **ConnectionString** of type string. The value of each variable consists of the connection string to the test or production data warehouses. You need 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 <b>ConnectionString</b> property of the Connection Manager to use the <b>ConnectionString</b> Environment Variable. |  |
| Deploy the project to the test and production environments.  |  |
| Add permissions to the <b>Test</b> and <b>Production</b> Environments.   |  |
| Execute the deployed package by using the required Environment.  |  |
| Add the Environment reference to the deployed project.   |  |
| In the deployed package, set the <b>ConnectionString</b> property of the Connection Manager to use the <b>ConnectionString</b> Environment Variable. |  |

**Answer:**



### QUESTION 79

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. Destination Assistant
- D. Script Component

**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>

### QUESTION 80

Drag and Drop Questions

You are developing a SQL Server Integration Services (SSIS) package that imports data into a data warehouse. You are developing the part of the SSIS package that populates the ProjectDates dimension table. The business key of the ProjectDates table is the ProjectName column. The business user has given you the dimensional attribute behavior for each of the four columns in the ProjectDates table.

- ExpectedStartDate - New values should be tracked over time.
- ActualStartDate - New values should not be accepted.
- ExpectedEndDate - New values should replace existing values.
- ActualEndDate - New values should be tracked over time.

You use the SSIS Slowly Changing Dimension Transformation. You must configure the Change Type value for each source column. Which Change Type values should you select? (To answer, drag the appropriate value from the list of values to the correct location or locations in the answer area.)



|                      |                   |             |
|----------------------|-------------------|-------------|
| Fixed Attribute      | Dimension Columns | Change Type |
| Historical Attribute | ExpectedStartDate |             |
| Change Attribute     | ActualStartDate   |             |
|                      | ExpectedEndDate   |             |
|                      | ActualEndDate     |             |

**Answer:**

|                      |                   |                      |
|----------------------|-------------------|----------------------|
| Fixed Attribute      | Dimension Columns | Change Type          |
| Historical Attribute | ExpectedStartDate | Historical Attribute |
| Change Attribute     | ActualStartDate   | Fixed Attribute      |
|                      | ExpectedEndDate   | Change Attribute     |
|                      | ActualEndDate     | Historical Attribute |

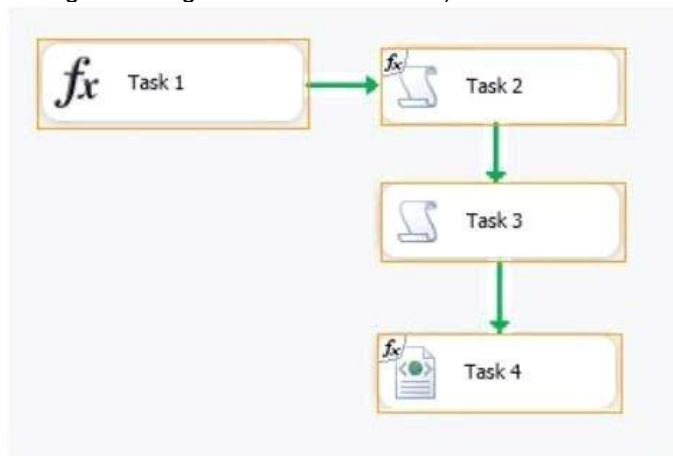
**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms141715.aspx>  
<http://msdn.microsoft.com/en-us/library/ms141662.aspx>

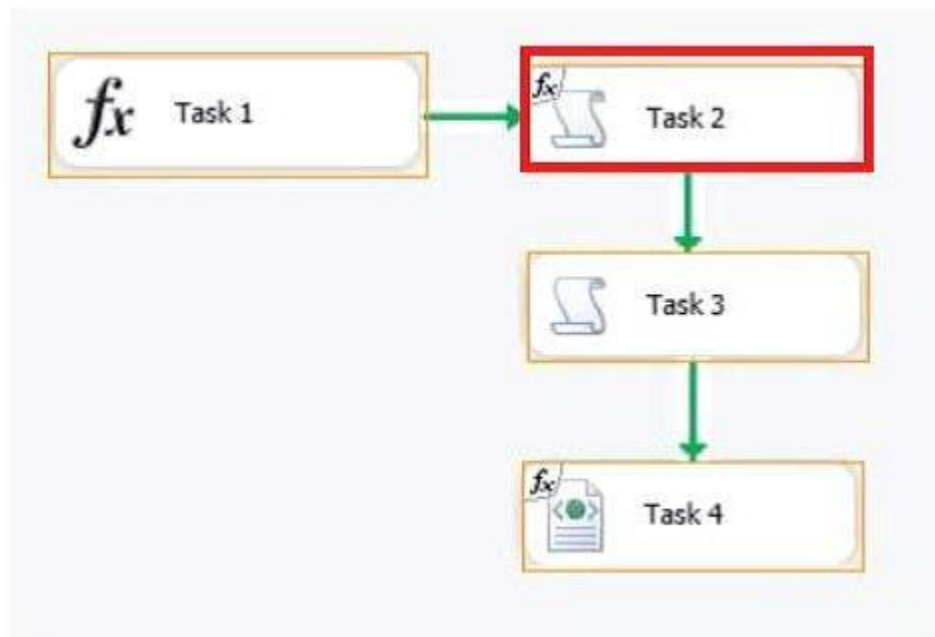
### QUESTION 81

Hotspot Questions

You are editing a SQL Server Integration Services (SSIS) package. You need to edit the Script task that has an expression defined. Which task should you select? (To answer, select the appropriate setting or settings in the answer area.)




**Answer:**



**QUESTION 82**

You are developing a data flow transformation to merge two data sources. One source contains product data and the other source contains data about the country in which the product was manufactured. Both data sources contain a two-character CountryCode column and both use SQL Server. Both data sources contain an ORDER BY clause to sort the data by the CountryCode column in ascending order. You use a Merge Join transformation to join the data. You need to ensure that the Merge Join transformation works correctly without additional transformations. What should you do? (Each answer presents a part of the solution. Choose all that apply.)



- A. Set the appropriate SortKeyPosition properties on the data sources.
- B. Set the ValidateExternalMetaData property on the Merge Join transformation to True.
- C. Set the IsSorted property on both data sources.
- D. Set the MaxBuffersPerInput property on the Merge Join transformation to 2.

**Answer: AC**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms141775.aspx>

<http://msdn.microsoft.com/en-us/library/ms137653.aspx>

<http://siddhumehta.blogspot.com/2009/05/validateexternalmetadata-property.html>

<http://msdn.microsoft.com/en-us/library/ms135950.aspx>

**QUESTION 83**

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.create\_environment\_variable
- B. catalog.create\_environment\_reference
- C. catalog.set\_execution\_parameter\_value
- D. catalog.set\_environment\_variable\_value

**Answer: B**

**QUESTION 84**

You are developing a SQL Server Integration Services (SSIS) package. The package is stored as the myPackage.dtsx file in the root directory of the C drive of the SSIS server. You need to run the package from the command prompt. Which command should you use?

- A. dtexec/sql myPackage
- B. dtexec/d "C:\File System\myPackage.dtsx"
- C. dtexec/f "C:\myPackage.dtsx"
- D. dtexec/com "myPackage.dtsx"

**Answer: C**

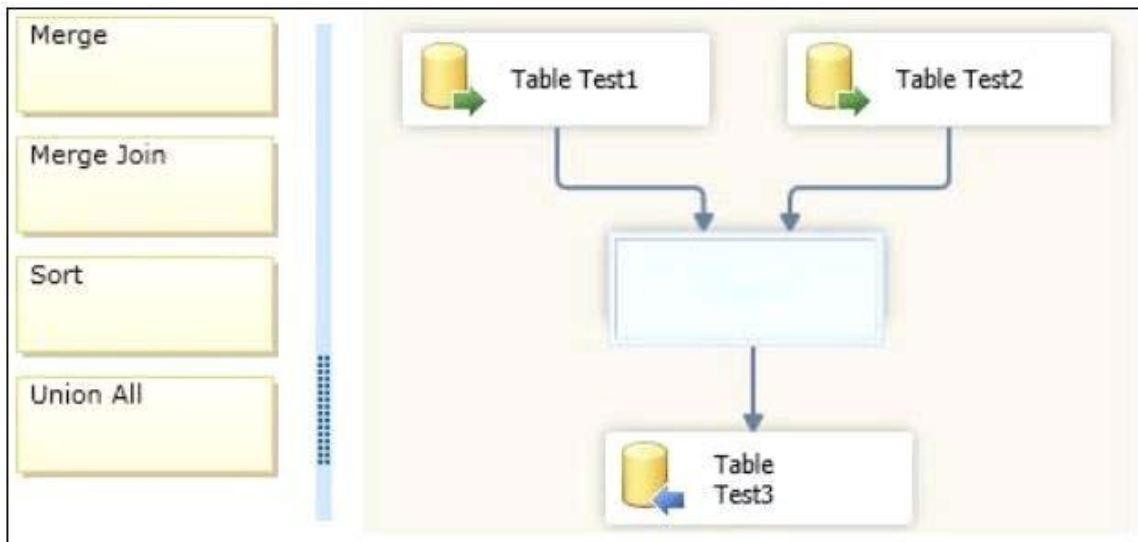
**QUESTION 85**

Drag and Drop Question

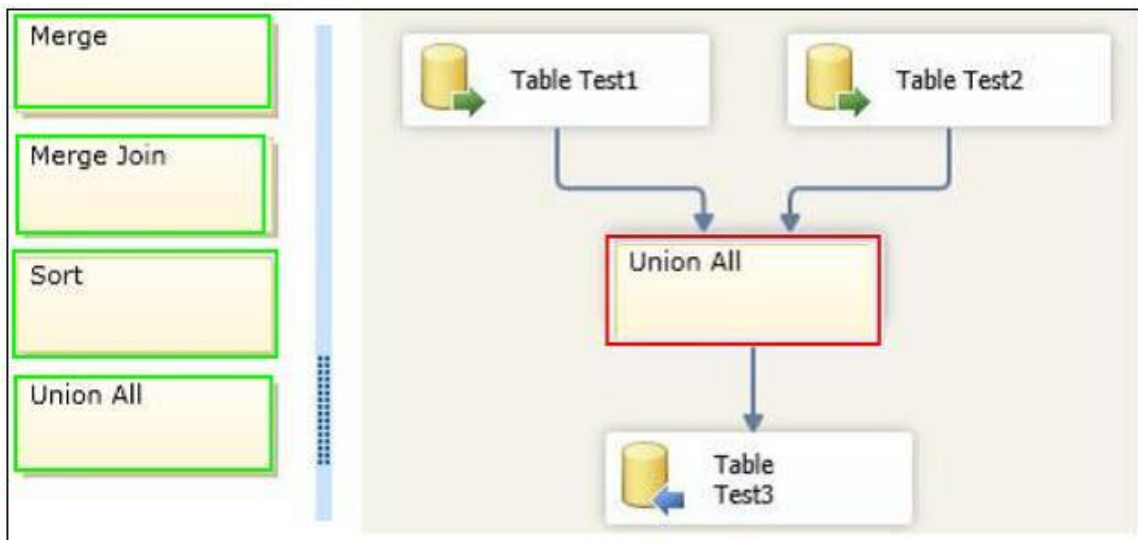
You are developing a SQL Server Integration Services (SSIS) package that imports unsorted data into a data warehouse hosted on SQL Azure. You have the following requirements:

- A destination table must contain all of the data in two source tables.
- Duplicate records must be inserted into the destination table.

You need to develop a data flow that imports the data while meeting the requirements. How should you develop the data flow? (To answer, drag the appropriate transformation from the list of transformations to the correct location in the answer area.)



**Answer:**



**QUESTION 86**

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 tap.

**Answer: C**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms141212.aspx>

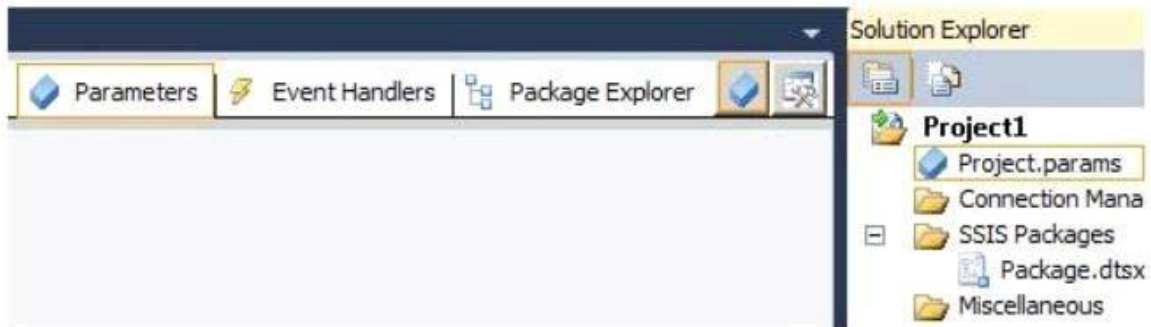
<http://www.mssqltips.com/sqlservertip/2450/ssis-package-deployment-model-in-sql-server-2012-part-1-of-2/>

<http://www.mssqltips.com/sqlservertip/2450/ssis-package-deployment-model-in-sql-server-2012-part-2-of-2/>

### QUESTION 87

Hotspot Questions

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:**



### QUESTION 88

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 model.

**Answer: A**

**QUESTION 89**

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. Change the Exhaustive property to True.
- B. Change the similarity threshold to 0.55.
- C. Change the similarity threshold to 0.40.
- D. Increase the maximum number of matches per lookup.

**Answer: B**

**Explanation:**

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

**QUESTION 90**

You install a SQL Server 2012 database engine instance on a production server. A month later, you install SQL Server 2012 Integration Services (SSIS). You must develop an SSIS project and deploy it to the server by using the Project Deployment model. Operations Log records that are outside the configured retention period must be cleaned automatically. You need to create the SSIS catalog on the production server. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Enable XP Command Shell.
- B. Enable CLR Integration.
- C. Enable OLE Automation.
- D. Start the SQL Server Browser service.
- E. Enable Cross Database Ownership Chaining.
- F. Start the SQL Server Agent service.
- G. Enable Ad Hoc Remote Queries.

**Answer: BF**

**Explanation:**

[http://msdn.microsoft.com/en-us/library/gg471509\(v=sql.110\).aspx](http://msdn.microsoft.com/en-us/library/gg471509(v=sql.110).aspx)

**QUESTION 91**

A SQL Server Integration Services (SSIS) package on a computer is failing periodically in production. The package was created less than one year ago and was deployed to the SSIS catalog. Sometimes the package is started on a SQL Agent schedule; sometimes the package is started manually by an SSIS developer 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 failed in the past. Where can you find this information?

- A. the SQL Server Log
- B. the SSISDB.[catalog].[executions] view

- C. the SSISDB.[catalog].[event\_messages] view
- D. the SQL Agent Job History
- E. the SQL Agent Error Log

**Answer: B**

**QUESTION 92**

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**

**QUESTION 93**

Drag and Drop Question

A SQL Server Integration Services (SSIS) package named DataFeed interacts with an external vendor data feed. The package is executed several times a day, either as part of other packages' control flow or by itself. The external data feed is unreliable because network failures and slow response times are frequent. The package is currently deployed on the file system. To analyze the reliability of the external data feed, you must collect execution data. Every time the DataFeed package is executed, the following information must be logged:

- Start Time
- End Time
- Execution Result
- Execution Duration

You need to design a logging solution that meets the requirements by using the least amount of administrative and development effort. 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 <b>OnError</b> and <b>OnWarning</b> event handlers.   |  |  |
| Query the <b>catalog.executable_statistics</b> view.  |   |  |
| Query the <b>ExecutionLog</b> table.  |   |  |
| Deploy the project that contains the package to the SSIS catalog and execute the package on the server. |   |  |
| Deploy the package to the msdb database.  |   |  |
| Add an Execute SQL task to the event handlers.  |   |  |
| Query the <b>catalog.execution_data_statistics</b> view.  |   |  |
| Create a new project and add the package to the project.  |   |  |

**Answer:**

|   |  |   |
|---|--|---|
| Add <b>OnError</b> and <b>OnWarning</b> event handlers.   | <div> <div>➤</div> <div>➤</div> </div> | Create a new project and add the package to the project.  |
| Query the <b>catalog.executable_statistics</b> view.  |  |   |
| Query the <b>ExecutionLog</b> table.  |  |   |
| Deploy the project that contains the package to the SSIS catalog and execute the package on the server. |  | Deploy the project that contains the package to the SSIS catalog and execute the package on the server. |
| Deploy the package to the msdb database.  |  |   |
| Add an Execute SQL task to the event handlers.  |  |   |
| Query the <b>catalog.execution_data_statistics</b> view.  |  | Query the <b>catalog.executable_statistics</b> view.  |
| Create a new project and add the package to the project.  |  |   |

#### QUESTION 94

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. Create a time dimension that can join to both fact tables at their respective granularity.
- D. Merge the fact tables.

**Answer: C**

**Explanation:**

<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/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/>

#### QUESTION 95

Drag and Drop Question

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.)

|   |  |
|---|--|
| Select the property to be parameterized and set the scope to <b>Project</b> . | <div> <div>➤</div> <div>➤</div> </div> |
| Edit the parameter and set the <b>Sensitive</b> property to <b>True</b> .     |  |
| Convert the project to the Legacy Deployment model.                           |  |
| Select the property to be parameterized and set the scope to <b>Package</b> . |  |
| Right-click the task and choose <b>Parameterize</b> .                         |  |
| Right-click the task and choose <b>Properties</b> .                           |  |

**Answer:**

|   |   |
|---|---|
| Select the property to be parameterized and set the scope to <b>Project</b> . | Right-click the task and choose <b>Parameterize</b> .                         |
| Edit the parameter and set the <b>Sensitive</b> property to <b>True</b> .     | Select the property to be parameterized and set the scope to <b>Package</b> . |
| Convert the project to the Legacy Deployment model.                           | Edit the parameter and set the <b>Sensitive</b> property to <b>True</b> .     |
| Select the property to be parameterized and set the scope to <b>Package</b> . |   |
| Right-click the task and choose <b>Parameterize</b> .                         |   |
| Right-click the task and choose <b>Properties</b> .                           |   |

### QUESTION 96

Drag and Drop Questions

A new SQL Server Integration Services (SSIS) project is deployed to the SSIS catalog. To troubleshoot some data issues, you must output the data streaming through several data flows into text files for further analysis. You have the list of data flow package paths and identification strings of the various task components that must be analyzed. You need to create these output files with the least amount of administrative and development effort. Which three stored procedures should you execute in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

|                               |  |  |
|-------------------------------|--|--|
| catalog.create_folder         |  |  |
| catalog.create_execution_dump |  |  |
| catalog.add_data_tap          |  |  |
| catalog.configure_catalog     |  |  |
| catalog.add_data_tap_by_guid  |  |  |
| catalog.create_execution      |  |  |
| catalog.start_execution       |  |  |

**Answer:**

|                               |  |                          |
|-------------------------------|--|--------------------------|
| catalog.create_folder         |  | catalog.create_execution |
| catalog.create_execution_dump |  |                          |
| catalog.add_data_tap          |  | catalog.add_data_tap     |
| catalog.configure_catalog     |  |                          |
| catalog.add_data_tap_by_guid  |  | catalog.start_execution  |
| catalog.create_execution      |  |                          |
| catalog.start_execution       |  |                          |

**QUESTION 97**

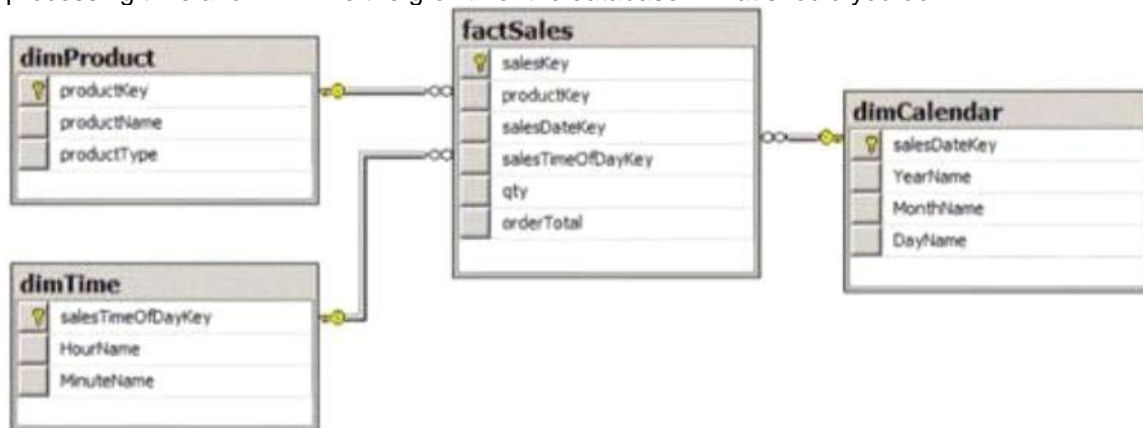
You are designing a data warehouse hosted on SQL Azure. The data warehouse currently includes the dimUser and dimDistrict dimension tables and the factSales fact table. The dimUser table contains records for each user permitted to run reports against the warehouse; and the dimDistrict table contains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

**QUESTION 98**

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**

**QUESTION 99**

You are designing a data warehouse for a software distribution business that stores sales by software title. It stores sales targets by software category. Software titles are classified into subcategories and categories. Each software title is included in only a single software 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 software title
- factTarget, used to record the monthly sales targets by software category

Reports must be developed against the warehouse that reports sales by software title, category and subcategory, and sales targets. You need to design the software title dimension. The solution should use as few tables as possible while supporting all the requirements. What should you do?

- A. Create three software tables, dimSoftware, dimSoftwareCategory, and dimSoftwareSubcategory and a fourth bridge table that joins software titles to their appropriate category and subcategory table records with foreign key constraints.  
Direct the cube developer to use key granularity attributes.
- B. Create three software tables, dimSoftware, dimSoftwareCategory, and dimSoftwareSubcategory.  
Connect factSales to all three tables and connect factTarget to dimSoftwareCategory with foreign key constraints.  
Direct the cube developer to use key granularity attributes.
- C. Create one table, dimSoftware, which contains Software Detail, Category, and Subcategory columns.  
Connect factSales to dimSoftware with a foreign key constraint.  
Direct the cube developer to use a non-key granularity attribute for factTarget.
- D. Create two tables, dimSoftware and dimSoftwareCategory.  
Connect factSales to dimSoftware and fact Target to dimSoftwareCategory with foreign key constraints.  
Direct the cube developer to use key granularity attributes.

**Answer: C**

#### QUESTION 100

You administer a Microsoft SQL Server database. You want to import data from a text file to the database. You need to ensure that the following requirements are met:

- Data import is performed by using a stored procedure.
- Data is loaded as a unit and is minimally logged.

Which data import command and recovery model should you choose? (To answer, drag the appropriate data import command or recovery model to the appropriate location or locations in the answer area. Each data import command or recovery model may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.)

| Command/Recovery Model Name | Command/Recovery Model |
|-----------------------------|------------------------|
| BCP                         | Data import command    |
| BULK INSERT                 | Recovery model         |
| Bulk-logged                 |                        |
| OPENDATASOURCE              |                        |
| Full                        |                        |

**Answer:**



| Command/Recovery Model Name | Command/Recovery Model           |
|-----------------------------|----------------------------------|
| BCP                         | Data import command: BULK INSERT |
| BULK INSERT                 | Recovery model: Bulk-logged      |
| Bulk-logged                 |                                  |
| OPENDATASOURCE              |                                  |
| Full                        |                                  |

#### QUESTION 101

You administer a Microsoft SQL Server 2012 server that has SQL Server Integration Services (SSIS) installed. You plan to deploy new SSIS packages to the server. The SSIS packages use the Project Deployment Model together with parameters and Integration Services environment variables. You need to configure the SQL Server environment to support these packages. What should you do?

- A. Create SSIS configuration files for the packages.
- B. Create an Integration Services catalog.
- C. Install Data Quality Services.
- D. Install Master Data services.

**Answer: B**

**Explanation:**

<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>

#### QUESTION 102

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. Source Assistant
- B. Audit
- C. Destination Assistant
- D. Script Component

**Answer: D**

#### QUESTION 103

Drag and Drop Questions

You plan to deploy a SQL Server Integration Services (SSIS) project by using the project deployment model. You need to monitor control flow tasks to determine whether any of them are running longer than usual. Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

Write a query against the **catalog.operation\_messages** view. Add a calculation to the query to compare durations to the **catalog.executables** view.

Execute the query.

Write a query against the **catalog.execution\_component\_phases** view. Add a calculation to the query to compare durations to the **catalog.executables** view.

Connect to the **SSISDB** database.

Connect to the **msdb** database.

Write a query against the **catalog.execution\_component\_phases** view. Add a calculation to the query to compare durations to the **catalog.executions** view.

Write a query against the **catalog.operation\_messages** view. Add a calculation to the query to compare durations to the **catalog.executions** view.

**Answer:**

Write a query against the **catalog.operation\_messages** view. Add a calculation to the query to compare durations to the **catalog.executables** view.

Execute the query.

Write a query against the **catalog.execution\_component\_phases** view. Add a calculation to the query to compare durations to the **catalog.executables** view.

Connect to the **SSISDB** database.

Connect to the **msdb** database.

Write a query against the **catalog.execution\_component\_phases** view. Add a calculation to the query to compare durations to the **catalog.executions** view.

Write a query against the **catalog.operation\_messages** view. Add a calculation to the query to compare durations to the **catalog.executions** view.

Connect to the **SSISDB** database.

Write a query against the **catalog.execution\_component\_phases** view. Add a calculation to the query to compare durations to the **catalog.executions** view.

Execute the query.

**QUESTION 104**

Drag and Drop Questions

You are validating whether a SQL Server Integration Services (SSIS) package named Master.dtsx in the SSIS catalog is executing correctly. You need to display the number of rows in each buffer passed between each data flow component of the package. Which three actions should you perform in sequence? (To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.)

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

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

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

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

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

**Answer:**

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

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

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

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

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

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

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

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

**QUESTION 105**

**Drag and Drop Questions**

You are designing a SQL Server Integration Services (SSIS) package to execute 12 Transact-SQL (T-SQL) statements on a SQL Azure database. The T-SQL statements may be executed in any order. The T-SQL statements have unpredictable execution times. You have the following requirements:

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

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

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

Add a Sequence container to the control flow.

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

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

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

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



**Answer:**

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

Add a Sequence container to the control flow.

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

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

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

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



Add a Sequence container to the control flow.

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

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

**QUESTION 106**

You work as a senior database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are running a training exercise for Microsoft SQL Server 2012 junior administrators. You are discussing the use of Slowly Changing Dimension Transformation Outputs. One of the output options causes Derived Column transformations to create columns for the expired row and the current row indicators. Which option is the output that causes this?

- A. Unchanged Output
- B. Inferred Member Updates Output
- C. Historical Attributes Inserts Output
- D. Fixed Attribute Output
- E. Changing Attributes Updates Output

**Answer: C**

**QUESTION 107**

Drag and Drop Questions

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

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



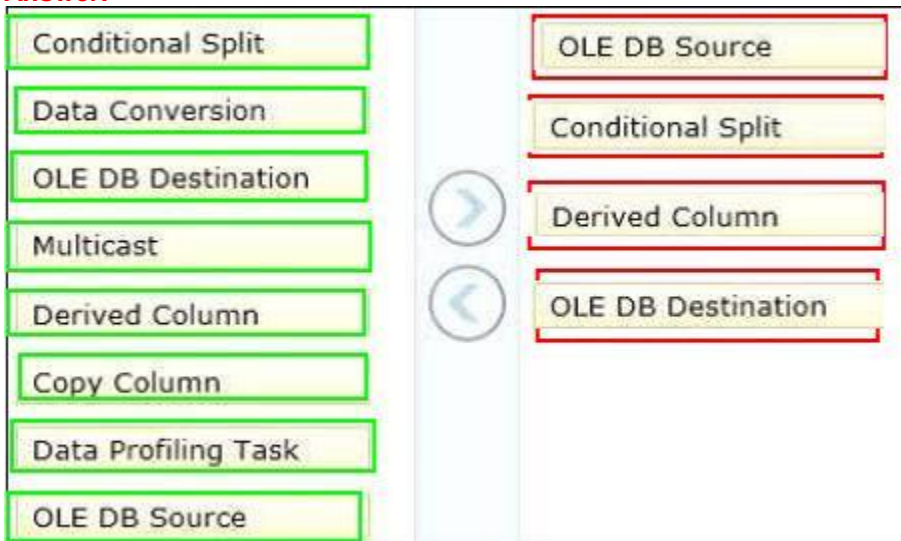
must be discarded.

You need to create the SSIS package in SQL Server Data Tools. In what sequence should you order four of the listed components for the data flow task? (To answer, move the appropriate components from the list of components to the answer area and arrange them in the correct order.)



The image shows the SSIS Data Flow Task component selection interface. On the left, there is a list of components: Conditional Split, Data Conversion, OLE DB Destination, Multicast, Derived Column, Copy Column, Data Profiling Task, and OLE DB Source. In the center, there are two circular arrows, one pointing right and one pointing left. On the right, there is an empty area for the selected components.

**Answer:**



The image shows the SSIS Data Flow Task component selection interface with the answer. On the left, the components are: Conditional Split, Data Conversion, OLE DB Destination, Multicast, Derived Column, Copy Column, Data Profiling Task, and OLE DB Source. In the center, there are two circular arrows, one pointing right and one pointing left. On the right, the components are arranged in the correct order: OLE DB Source, Conditional Split, Derived Column, and OLE DB Destination.

**QUESTION 108**

You work as a senior database developer at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are running a training exercise for Microsoft SQL Server 2012 junior developers. You are discussing the ValidateExternalMetaData property for a data flow. Which of the following represents the data type for this property?

- A. Integer
- B. IDTSExternalMetadataColumn100
- C. String
- D. Boolean



**Answer: D**

**QUESTION 109**

Drag and Drop Question

You administer a Microsoft SQL Server database. Service accounts for SQL Agent are configured to use a local user. A Microsoft SQL Server Integration Services (SSIS) job step has been created within a SQL Server Agent job. The SSIS package accesses a network share when exporting data from a SQL Server database. When you execute the SQL Server Agent job, it fails due to a permissions failure on a share on a remote server. You need to ensure that the SQL Server Agent job can execute the SSIS package. 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.)

| Ordered List Title                  | Answer Choices Title  |
|-------------------------------------|---|
| <div><div>▲▼</div><div></div></div> | <div>Add a proxy that references the local user.</div> <div>Add a proxy that references the credential.</div> <div>Create a local user account and grant local administrator on the SQL Server instance.</div> <div>Create a credential that references the local user.</div> <div>Create a credential that references the domain user.</div> <div>Assign the proxy to the Operating System subsystem.</div> <div>Assign the proxy to the SSIS package execution subsystem.</div> <div>Create a domain user account and grant permissions to the domain user account to access the network share.</div> |
|                                     | <div>&lt;&lt; Move</div> <div>Remove &gt;&gt;</div>   |

**Answer:**

| Ordered List Title   |   | Answer Choices Title  |
|--|---|---|
| <div> <div>▲</div> <div>▼</div> <div>Create a domain user account and grant permissions to the domain user account to access the network share.</div> </div> |   | <div>Add a proxy that references the local user.</div>  |
| <div>Create a credential that references the domain user.</div>  |   | <div>Add a proxy that references the credential.</div>  |
|  | <div>&lt;&lt; Move</div> <div>Remove &gt;&gt;</div> | <div>Create a local user account and grant local administrator on the SQL Server instance.</div>                      |
| <div>Add a proxy that references the credential.</div>   |   | <div>Create a credential that references the local user.</div>  |
|  |   | <div>Create a credential that references the domain user.</div>   |
| <div>Assign the proxy to the SSIS package execution subsystem.</div>   |   | <div>Assign the proxy to the Operating System subsystem.</div>  |
|  |   | <div>Assign the proxy to the SSIS package execution subsystem.</div>  |
|  |   | <div>Create a domain user account and grant permissions to the domain user account to access the network share.</div> |

### QUESTION 110

You work as a senior database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are hosting a training exercise for Microsoft SQL Server 2012 administrators. You are currently discussing Multicast transformation. Which of the following is TRUE with regards to Multicast transformation? (Choose two.)

- A. It prevents a package from creating logical copies of data.
- B. It allows a package to create logical copies of data.
- C. It distributes its input to one or more outputs.
- D. It distributes its input to a single output.

**Answer: BC**

### QUESTION 111

Hotspot Questions

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.
- Force packages to load all settings in the configuration.
- Support Encrypting File System (EFS) formats.

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.)

**Package Configuration Wizard**

**Select Configuration Type**  
You can select the type of configuration to create.

SSIS supports a variety of sources to use for setting the properties of objects.

Configuration type: Environment variable

☒ Specify configuration settings directly

Configuration file name: [Text Box]

Configuration type: Environment variable, XML configuration file, Registry entry, Parent package variable, SQL Server

☐ Configuration location is stored in an environment variable

Environment variable: [Text Box]

**Answer:**

**Package Configuration Wizard**

**Select Configuration Type**  
You can select the type of configuration to create.

SSIS supports a variety of sources to use for setting the properties of objects.

Configuration type: Environment variable

☒ Specify configuration settings directly

Configuration file name: [Text Box]

Configuration type: Environment variable, XML configuration file, Registry entry, Parent package variable, SQL Server

☐ Configuration location is stored in an environment variable

Environment variable: [Text Box]

**QUESTION 112**

You work as a database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are in the process of creating a SQL Server Integration Services (SSIS) package. The package will be configured to make use of the Generate And Persist New Index Fuzzy Lookup transformation option. Which of the following is the purpose of this option?

- A. It allows for a current index to be reused.
- B. It allows for a new index to be created, without saving it.
- C. It allows for a new index to be created and saved, without maintaining it.
- D. It allows for a new index to be created, saved, and maintained.

**Answer: C**

**QUESTION 113**

You work as a database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are in the process of creating a SQL Server Integration Services (SSIS) package. You have decided to make use of Expression Task. Which of the following is TRUE with regards to Expression Task?

- A. It creates and evaluates expressions that set variable values at runtime.
- B. It provides code to perform functions that are not available in the built-in tasks and transformations that SQL Server Integration Services provides.
- C. It runs an application or batch file as part of a SQL Server Integration Services package workflow.
- D. It allows you to run SQL statements during package execution.

**Answer: A**

**QUESTION 114**

You work as a senior database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are running a training exercise for Microsoft SQL Server 2012 junior administrators. You are discussing data quality projects. Which of the following are benefits of data quality projects? (Choose all that apply.)

- A. It enables you to perform data cleansing on your source data by using the knowledge in a DQS knowledge base.
- B. It allows you to perform data matching on your source data by using the knowledge in a DQS knowledge base.
- C. It allows you to perform data matching on your source data by using the matching policy in a knowledge base.
- D. It enables you to perform data cleansing on your source data by using the matching policy in a knowledge base.

**Answer: AC**

**QUESTION 115**

You work as a database developer at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You have been asked to revise a SQL Server Integration Services (SSIS) package. The package only includes Execute SQL tasks. To make sure that, in the event of an Execute SQL task failing, changes made by all of the Execute SQL tasks is rolled back. To do this, you want to alter the TransactionOption property of the package. Which of the following is the value that must be set for the TransactionOption property of the package?

- A. NotSupported.
- B. Supported.
- C. NotRequired.
- D. Required.

**Answer: D**



**QUESTION 116**

You work as a database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You have recently created an SQL Server Integration Services (SSIS) project catalog that contains two Environments. A single Environment Variable, named ConnectionString, of type string is specified by each environment. The project also has a project Connection Manager configured to link up with the data warehouse. You have been instructed to make sure that deployed packages are executed using any of the specified Environments. Which of the following suitably describes an Environment Variable?

- A. An environment variable specifies the data that will be used by a package.
- B. An environment variable can be used to manage Integration Services objects in the catalog.
- C. An environment variable defines a literal value that can be assigned to a parameter during package execution.
- D. An environment variable allows you to use folders to organize your projects and environments.
- E. An environment variable is a container of variables that can be referenced by Integration Services projects.

**Answer: C**

**QUESTION 117**

You work as a senior database developer at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are hosting a training exercise for ABC.com's junior database developers. You are currently discussing an EventHandlerClass property that is used to retrieve or configure a value that indicates whether the container participates in transactions. Which of the following is the EventHandlerClass property that you are discussing?

- A. The TransactionOption property.
- B. The Executables property.
- C. The LoggingOptions property.
- D. The Variables property.

**Answer: A**

**QUESTION 118**

You work as a senior database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are running a training exercise for Microsoft SQL Server 2012 junior administrators. You are discussing the effects of setting IsolationLevel transaction property of a Data Flow task to Chaos. Which of the following suitably describes the effect?

- A. It prevents pending changes from more highly isolated transactions from being overwritten.
- B. It prevents shared locks from being issued and exclusive locks from being honored.
- C. It prevents other users from updating the data.
- D. It reduces blocking by storing a version of data that one application can read while another is modifying the same data.

**Answer: A**

**QUESTION 119**

You work as a database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are in the process of creating a SQL Server Integration Services (SSIS) package. The package will be configured to import data into a SQL Azure hosted data warehouse. You have decided to



make use of the Environment Variable package configuration as part of your solution. Which of the following is TRUE with regards to the Environment Variable package configuration?

- A. It stores the property path inside the package and the value outside the package.
- B. It stores both the property path and the value outside the package.
- C. It is used to fetch a value from a variable in a calling package.
- D. It is the most commonly used configuration type.

**Answer: A**

**QUESTION 120**

You work as a senior database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are hosting a training exercise for Microsoft SQL Server 2012 administrators. You are currently discussing the Fuzzy Lookup transformation. Which of the following is TRUE with regards to the Fuzzy Lookup similarity threshold setting? (Choose all that apply.)

- A. Increasing the threshold can improve the speed of matching.
- B. Increasing the threshold can deteriorate the speed of matching.
- C. The nearer the value is to 1, the nearer the resemblance of the lookup value to the source value must be to qualify as a match.
- D. The nearer the value is to 1, the further the resemblance of the lookup value to the source value must be to qualify as a match.

**Answer: AC**

**QUESTION 121**

You work as a database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You have created a SQL Server Integration Services (SSIS) project that you are preparing to deploy using the Integration Services Deployment Wizard. Which of the following is the third step in the Integration Services Deployment Wizard's process?

- A. Review
- B. Select Destination
- C. Deploy/Results
- D. Select Source

**Answer: A**

**QUESTION 122**

You work as a database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You have received reports regarding a current existing SQL Server Integration Services (SSIS) package not completing its execution, and not transferring data. You have decided to make use of the Package Deployment Model. Which of the following is TRUE with regards to using the Package Deployment Model?

- A. CLR integration is NOT required on the database engine.
- B. Environment-specific parameter values are stored in environment variables.
- C. Events that are created by the package are captured automatically and saved to the catalog during execution.
- D. Events that are created by the package are NOT captured automatically during execution.

**Answer: AD**

**QUESTION 123**

You work as a senior database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are running a training exercise for Microsoft SQL Server 2012 junior administrators. You are discussing the use of the add\_data\_tap stored procedure. Which of the following is TRUE with regards to using add\_data\_tap stored procedure? (Choose all that apply.)

- A. Limiting the number of rows that are captured during the data tap is not advisable.
- B. The add\_data\_tap stored procedure should only be used to troubleshoot data issues.
- C. The add data tap process is validated prior to the package being executed.
- D. You must be an administrator on the machine on which the stored procedure is run to open the file that contains the tapped data.

**Answer: BD**

**QUESTION 124**

You work as a senior database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are running a training exercise for Microsoft SQL Server 2012 junior administrators. You are discussing the use of Slowly Changing Dimension Transformation Outputs. One of the output options are used to update the record in the lookup table. Which of the following is the output option?

- A. Unchanged Output
- B. Inferred Member Updates Output
- C. Historical Attributes Inserts Output
- D. Fixed Attribute Output
- E. Changing Attributes Updates Output

**Answer: E**

**QUESTION 125**

You work as a database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are in the process of creating a SQL Server Integration Services (SSIS) package. The package will be configured upload a file to a table in a SQL Azure database. You need to make sure that the code includes an event handler that makes use of an executable to raise the event when its verification is completed. Which of the following is the event handler that you should use?

- A. The OnPostExecute event handler.
- B. The OnTaskFailed event handler.
- C. The OnPreExecute event handler.
- D. The OnPostValidate event handler.

**Answer: D**

**QUESTION 126**

You work as a senior database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are running a training exercise for Microsoft SQL Server 2012 junior

administrators. You are discussing replication type that is used when subscribers need to receive data, make changes offline, and later synchronize changes with the Publisher and other Subscribers. Which of the following is the replication type being discussed?

- A. Snapshot replication
- B. Transactional replication
- C. Merge replication
- D. Peer-to-peer replication

**Answer: C**

**QUESTION 127**

You work as a senior database developer at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You are running a training exercise for Microsoft SQL Server 2012 junior developers. You are discussing the SortKeyPosition property for a data flow. Which of the following represents the data type for this property?

- A. Integer
- B. IDTSEExternalMetadataColumn100
- C. String
- D. Boolean

**Answer: A**

**QUESTION 128**

You work as a database administrator at ABC.com. The ABC.com network consists of a single domain named ABC.com. ABC.com makes use of Microsoft SQL Server 2012 in their environment. You have received reports regarding a current existing SQL Server Integration Services (SSIS) package not completing its execution, and not transferring data. You have decided to make use of the Project Deployment Model. Which of the following is TRUE with regards to using the Project Deployment Model?

- A. CLR integration is required on the database engine.
- B. CLR integration is not required on the database engine.
- C. Events that are created by the package are captured automatically and saved to the catalog during execution.
- D. Packages are executed using the dtExec and DTEXECUI execution utilities.

**Answer: AC**

**QUESTION 129**

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

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

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

- A. Partition the destination fact table by date.

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

**Answer: B**

**QUESTION 130**

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 columns.  
Place the columnstore index on the dimensional key columns.
- B. Design a fact table with 5 dimensional key columns and 10 measure columns.  
Place the columnstore index on the measure columns.
- C. Design a fact table with 10 dimensional key columns and 5 measure columns.  
Place the columnstore index on the dimensional key columns and the measure columns.
- D. Design a fact table with 10 dimensional key columns and 5 measure columns.  
Place the columnstore index on only the measure columns.

**Answer: C**

**QUESTION 131**

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.  
Use a Data Flow task to extract the XML output of the Data Profiling task into a SQL Server table.  
Query the table to view the exceptions.
- B. Use Data Flow tasks to process the clean data.
- C. Use a Data Profiling task.  
Read the exceptions in Data Profile Viewer.

- D. Design a stored procedure that examines data for common dirty data patterns.  
Use an Execute SQL task.

**Answer: C**

**QUESTION 132**

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

**Answer: A**

**QUESTION 133**

You are designing a partitioning strategy for a large fact table in a data warehouse. Tens of millions of new records are loaded into the data warehouse weekly, outside of business hours. Most queries are generated by reports and by cube processing. Data is frequently queried at the day level and occasionally at the month level. You need to partition the table to maximize the performance of queries. What should you do? (More than one answer choice may achieve the goal. Select the BEST answer.)

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

**Answer: D**

**QUESTION 134**

Hotspot Question

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.)

The screenshot shows the 'SQL Server Import and Export Wizard' window, specifically the 'Choose a Destination' step. The window title is 'SQL Server Import and Export Wizard'. Below the title bar, it says 'Choose a Destination' and 'Specify where to copy data to.' The 'Destination' dropdown menu is set to 'Flat File Destination'. Below this, it says 'Select a file and specify the file properties and the file format.' The 'File name' field contains 'C:\Employee.csv' with a 'Browse...' button to its right. The 'Locale' dropdown is set to 'English (United States)' and the 'Unicode' checkbox is unchecked. The 'Code page' dropdown is set to '1252 (ANSI - Latin I)'. The 'Format' dropdown is currently empty. The 'Text qualifier' field contains '<none>'. The 'Column names in the first data row' checkbox is unchecked. At the bottom, there are buttons for 'Help', '< Back', 'Next >', 'Finish >>', and 'Cancel'.

**Answer:**

Tick the "Unicode" checkbox.

Set the Format option to "Fixed Width".

Untick the "Column names in the first data row" checkbox.

The screenshot shows the 'SQL Server Import and Export Wizard' window, specifically the 'Choose a Destination' step. The window title is 'SQL Server Import and Export Wizard'. Below the title bar, it says 'Choose a Destination' and 'Specify where to copy data to.' The 'Destination' dropdown menu is set to 'Flat File Destination'. Below this, it says 'Select a file and specify the file properties and the file format.' The 'File name' field contains 'C:\Employee.csv' with a 'Browse...' button to its right. The 'Locale' dropdown is set to 'English (United States)' and there is an unchecked 'Unicode' checkbox. The 'Code page' dropdown is set to '1252 (ANSI - Latin I)'. The 'Format' dropdown is highlighted with a green border. The 'Text qualifier' field contains '<none>'. There is an unchecked checkbox for 'Column names in the first data row'. At the bottom, there are buttons for 'Help', '< Back', 'Next >', 'Finish >>', and 'Cancel'.

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms178804.aspx>  
<http://msdn.microsoft.com/en-us/library/ms187828.aspx>

**QUESTION 135**

You are a database developer of a Microsoft SQL Server 2012 database. You are designing a table that will store Customer data from different sources. The table will include a column that contains the CustomerID from the source system and a column that contains the SourceID. A sample of this data is as shown in the following table. You need to ensure that the table has no duplicate CustomerID within a SourceID. You also need to ensure that the data in the table is in the order of SourceID and then CustomerID. Which Transact-SQL statement should you use?

- A. 

```
CREATE TABLE Customer
(SourceID int NOT NULL IDENTITY,
CustomerID int NOT NULL IDENTITY,
CustomerName varchar(255) NOT NULL);
```
- B. 

```
CREATE TABLE Customer
(SourceID int NOT NULL,
CustomerID int NOT NULL PRIMARY KEY CLUSTERED,
CustomerName varchar(255) NOT NULL);
```
- C. 

```
CREATE TABLE Customer
(SourceID int NOT NULL PRIMARY KEY CLUSTERED,
CustomerID int NOT NULL UNIQUE,
CustomerName varchar(255) NOT NULL);
```
- D. 

```
CREATE TABLE Customer
(SourceID int NOT NULL,
CustomerID int NOT NULL,
```

```
CustomerName varchar(255) NOT NULL,  
CONSTRAINT PK_Customer PRIMARY KEY CLUSTERED,  
(SourceID, CustomerID));
```

**Answer: D**

**QUESTION 136**

Hotspot Question

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.)

**Package Configuration Wizard**

**Select Configuration Type**  
You can select the type of configuration to create.

SSIS supports a variety of sources to use for setting the properties of objects.

Configuration type:

☒ Specify configuration settings directly

Configuration file name:

☐ Configuration location is stored in an environment variable

Environment variable:

**Answer:**

**QUESTION 137**

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. 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 create a unique constraint?

| Employee (jek) |                |
|----------------|----------------|
| 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<br>Used throughout the database by all the other tables that reference the Employee table            |
| EmployeeNum    | An alphanumeric value calculated according to company requirements<br>Has to be unique within the Employee table<br>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  |

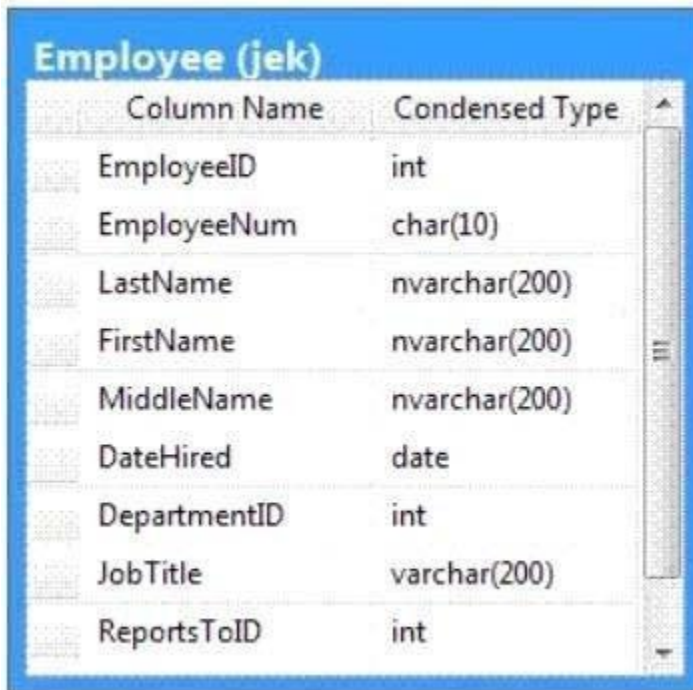
- A. DateHired
- B. DepartmentID
- C. EmployeeID
- D. EmployeeNum
- E. FirstName
- F. JobTitle
- G. LastName
- H. MiddleName
- I. ReportsToID

**Answer: D**

**QUESTION 138**

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. 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?





| 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<br>Used throughout the database by all the other tables that reference the Employee table            |
| EmployeeNum  | An alphanumeric value calculated according to company requirements<br>Has to be unique within the Employee table<br>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  |

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

I. ReportsToID

**Answer: C**

**QUESTION 139**

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

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

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

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

**Answer: A**

**QUESTION 140**

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 bulkimport operation. The value specified for ROWS\_PER\_BATCH should approximately the same as the actual number of rows.

**QUESTION 141**

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.  
After 10 executions, view the emails.
- C. Enable logging to the Application Event Log in the package control flow for the OnInformation event.  
After 10 executions, view the Application Event Log.
- D. Enable logging to the Application Event Log in the package control flow for the OnPostExecute event.  
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.

**QUESTION 142**

Your team is creating SQL Server Integration Services (SSIS) packages that have several dependencies. The packages use parameters for configuration purposes. Your company's IT policies include the following change control requirements:

- After a package is ready for deployment, your team must hand over the process to junior IT personnel.
- The process must guarantee that when a package has been validated, that same package (and all its dependencies) is deployed to production.
- The process must be repeatable and reliable and must be executed with the least administrative and training effort by junior IT personnel.

You need to use the most appropriate deployment unit to satisfy the company policies, while minimizing issues such as incorrect version or configuration. Which type of deployment unit should you use?

- A. dtsx
- B. SSISdeploymentmanifest
- C. msd
- D. ispac

**Answer: D**

**QUESTION 143**

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. Use an msi file to deploy the package on the server.
- B. Open a command prompt and run the dtutil/copy command.
- C. Add an OnError event handler to the SSIS project.
- D. Open a command prompt and run the gacutil command.

**Answer: B**

**QUESTION 144**

You are developing a SQL Server Integration Services (SSIS) project by using the Project Deployment Model. The project is deployed to a single SSIS catalog, and transfers data to and from multiple databases hosted on SQL Server. The project must be configured to be able to export data to and from five different production servers that run SQL Server 2012. Each target server requires different values for connection strings and parameters in the SSIS project. You need to meet the requirements by using the least amount of administrative effort. What should you do?

- A. For each target server, create separate registry entry configurations.  
Select the registry entry at package execution time.
- B. For each target server, create separate Environments in the SSIS catalog of the host SQL Server SSIS instance.  
Select the appropriate Environment at package execution time.
- C. Create one SSIS catalog Environment.  
Change the values of each Environment variable at package execution time.
- D. For each target server, create a separate XML configuration file.  
Select the XML configuration file at package execution time.

**Answer: B**

**QUESTION 145**

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. Depending on the value in a source column, the package redirects rows to one of five different data flow paths. You need to add a data flow transformation to support the package redirection. Which data flow transformation should you use?

- A. Conditional Split
- B. pivot
- C. Multicast
- D. Lookup

**Answer: A**

**QUESTION 146**

You are developing a SQL Server Integration Services (SSIS) project that contains a project Connection Manager and multiple packages. All packages in the project must connect to the same database. The server name for the database must be set by using a parameter named ServerParam when any package in the project is executed. You need to develop this project with the least amount of development effort. What should you do? (Each correct answer presents part of the solution. Choose all that apply.)

- A. Set the Sensitive property of the parameter to True.
- B. Edit each package Connection Manager.  
Set the ServerName property to @\$Project::ServerParam].
- C. Edit the project Connection Manager in Solution Explorer.  
Set the ServerName property to @\$Project::ServerParam].
- D. Create a project parameter named ServerName.
- E. Create a package parameter named ServerName in each package.
- F. Set the Required property of the parameter to True.

**Answer: CDF**

**Explanation:**

C: from question: "The server name for the database must be set by using a parameter named ServerParam when any package in the project is executed."

D: SSIS 2012 has introduced the concept of Project level connection managers. An SSIS project is generally more than one package. To simplify lives, the SSIS team now allows for the sharing of common resources across projects, connection managers being one of those resources.

F: When a parameter is marked as required, a server value or execution value must be specified for that parameter. Otherwise, the corresponding package does not execute. Although the parameter has a default value at design time, it will never be used once the project is deployed.  
<https://docs.microsoft.com/en-us/sql/integration-services/integration-services-ssis-package-and-project-parameters>

#### **QUESTION 147**

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 operations.

**Answer: B**

#### **QUESTION 148**

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 information. Create an error output on each data flow destination that writes OnError event text to the table.
- R. Create a table to store error information. Create an error output on each data flow destination that writes OnTaskFailed event text to the table.

**Answer: M**

**QUESTION 149**

You are developing a SQL Server Integration Services (SSIS) package to load data into a data warehouse. The package consists of several data flow tasks. The package experiences intermittent errors in the data flow tasks. If any data flow task fails, all package error information must be captured and written to a SQL Server table by using an OLE DB connection manager. You need to ensure that the package error information is captured and written to the table. 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 information. Create an error output on each data flow destination that writes OnError event text to the table.
- R. Create a table to store error information. Create an error output on each data flow destination that writes OnTaskFailed event text to the table.

**Answer: I**

**QUESTION 150**

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**

**QUESTION 151**

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 option.

**Answer: C**

**QUESTION 152**

You are developing a SQL Server Integration Services (SSIS) project that contains a project Connection Manager and multiple packages. All packages in the project must connect to the same database. The server name for the database must be set by using a parameter named ParamConnection when any package in the project is executed. You need to develop this project with the least amount of development effort. What should you do? (Each answer presents a part of the solution. Choose all that apply.)

- A. Create a package parameter named ConnectionName in each package.
- B. Edit each package Connection Manager.  
Set the ConnectionName property to @[\$Project::ParamConnection].
- C. Edit the project Connection Manager in Solution Explorer.  
Set the ConnectionName property to @[\$Project::ParamConnection].
- D. Set the Sensitive property of the parameter to True.
- E. Create a project parameter named ConnectionName.
- F. Set the Required property of the parameter to True.

**Answer: CEF**

**QUESTION 153**

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**

**QUESTION 154**

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 in SQL Server.

**Answer: D**

**QUESTION 155**

Drag and Drop Questions

You are maintaining a SQL Server Integration Services (SSIS) package. The package uses custom functionality that is implemented in Microsoft Visual C. The implementation of the custom functionality changes overtime. The design of the package allows you to deploy new releases of the custom functionality without redeploying the entire package. You need to implement and deploy an update to the custom functionality without requiring package redeployment. 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.)

- Redeploy the package.
- Edit the C# code.
- Rebuild and redeploy the DLL file that is generated by the Script task.
- Rebuild and redeploy the DLL file that is generated by the Script component.
- Rebuild and redeploy the custom object.
- Open the Script task of the package that contains the custom functionality.
- Open the custom object that contains the custom functionality.
- Open the Script component of the package that contains the custom functionality.

**Answer:**

|  |  |
|--|--|
| Redeploy the package.  | Open the custom object that contains the custom functionality. |
| Edit the C# code.  | Edit the C# code.  |
| Rebuild and redeploy the DLL file that is generated by the Script task.          | Rebuild and redeploy the custom object.                        |
| Rebuild and redeploy the DLL file that is generated by the Script component.     |  |
| Rebuild and redeploy the custom object.  |  |
| Open the Script task of the package that contains the custom functionality.      |  |
| Open the custom object that contains the custom functionality.                   |  |
| Open the Script component of the package that contains the custom functionality. |  |

**QUESTION 156**

**Hotspot Questions**

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.)

| Knowledge Base Management   | Data Quality Projects  |
|---|--|
| Create or maintain data quality Knowledge Base                                      | Create or maintain Data Quality Project  |
|  |  |
|  |  |

**Answer:**



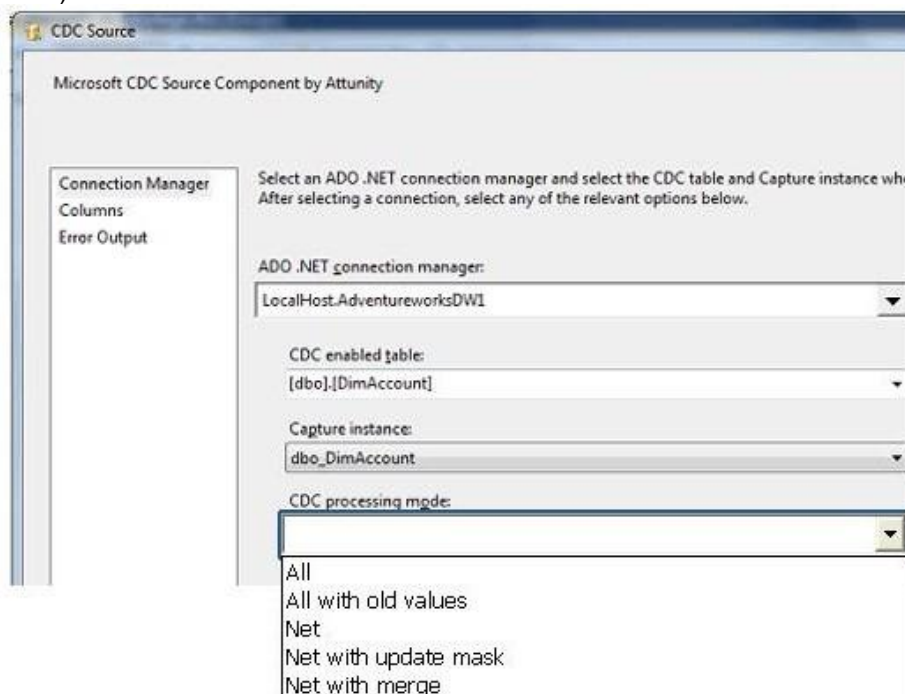
**QUESTION 157**

Hotspot Questions

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:

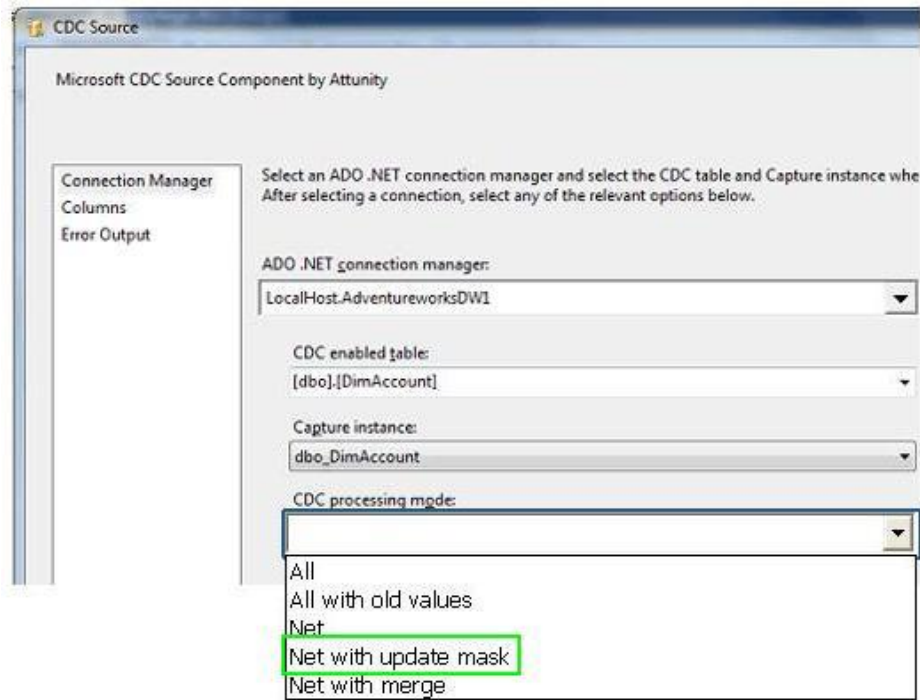
- The output must include metadata columns that indicate which source columns have changed.
- The output must return only one change row per source row that is modified in the current CDC processing range.

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.)





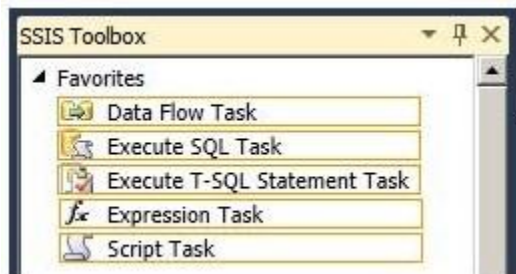
**Answer:**



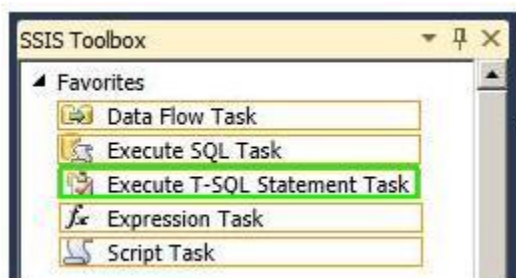
**QUESTION 158**

Hotspot Questions

You are developing a SQL Server Integration Services (SSIS) package. The package must run a parameterized query against a Windows Azure SQL Database. You need to use the least amount of development effort to meet the package requirement. Which task should you use? (To answer, select the appropriate task in the answer area.)



**Answer:**



**QUESTION 159**

Drag and Drop Questions

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:**

- 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.

Open the Control Flow designer of the package.

Add and edit a Foreach Loop container.

Set the enumerator property to **Foreach ADO Enumerator** and select the variable that contains the ADO object source.

**QUESTION 160**

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.

Incorrect:

- Maximum number of matches to output per lookup. Specify the maximum number of matches the transformation can return for each input row. The default is 1.

### QUESTION 161

Hotspot Questions

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.)

☐ Cache connection manager

☒ OLE DB connection manager

Specify how to handle rows with no matching entries

Ignore Failure  
Redirect rows to error output  
Fail component  
Redirect rows to no match output

**Answer:**

☐ Cache connection manager

☒ OLE DB connection manager


Specify how to handle rows with no matching entries

Ignore Failure  
Redirect rows to error output  
Fail component  
Redirect rows to no match output

### QUESTION 162

Hotspot Questions

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 opened 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 command should you use? (To answer, select the appropriate command in the answer area.)



**Answer:**



**QUESTION 163**

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 Integration Services Deployment Wizard
- B. SQL Server Configuration Manager
- C. SSIS Upgrade Wizard in SQL Server Management Studio
- D. SSIS Upgrade Wizard in SQL Server 2008 BIDS

**Answer: C**

**Explanation:**

You can upgrade packages that were created in earlier versions of Integration Services to the Integration Services format that SQL Server 2012 uses. SQL Server provides the SSIS Package Upgrade Wizard to help in this process. Because you can configure the wizard to backup up your original packages, you can continue to use the original packages if you experience upgrade difficulties. You can run the SSIS Package Upgrade Wizard from SQL Server Data Tools (SSDT), from SQL Server Management Studio, or at the command prompt.

**Note:**

When you upgrade an instance of SQL Server 2005 or SQL Server 2008 to the current release of SQL Server, your existing SQL Server 2008 Integration Services (SSIS) packages are not automatically upgraded to the package format that the current release SQL Server Integration Services uses. You will have to select an upgrade method and manually upgrade your packages.

**QUESTION 164**

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?



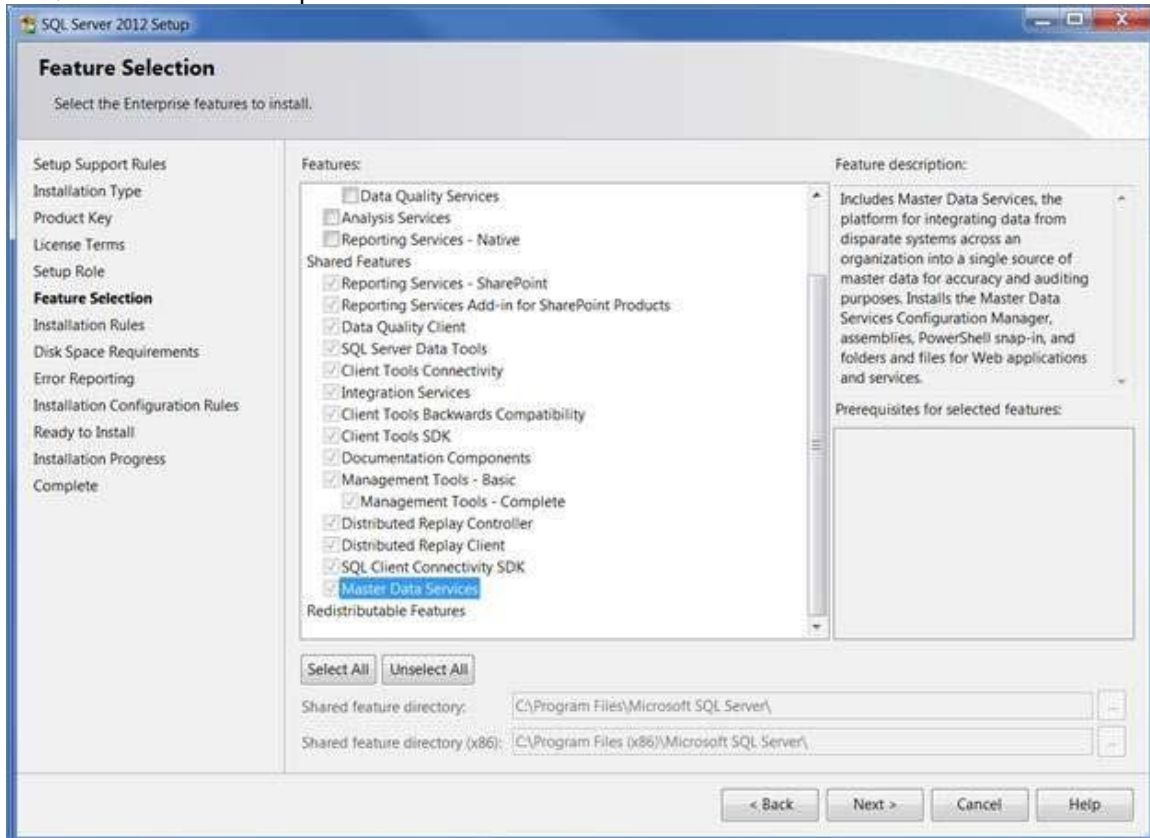
- A. Microsoft SharePoint Foundation 2010 SP1
- B. SQL Server 2012 Enterprise (64-bit) x64 on the database server
- C. SQL Server 2012 Data Center (64-bit) x64 on the database server
- D. SQL Server 2008 Enterprise (64-bit) x64 on the database server

**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.

#### **QUESTION 165**

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 4GB 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: 2GB, recommended: 4GB 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.

#### **QUESTION 166**

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. Install a 64-bit ACE driver and execute the package by using the 64-bit run-time option.
- B. Enable Address Windowing Extensions (AWE) for the local SQL Server instance.
- C. Replace the SSIS Excel source with a SSIS Flat File source.
- D. Install a 64-bit ACE driver and replace the Excel source with an OLE DB source.

**Answer: D**

**Explanation:**

XLSX files, the new file type in Excel 2007/2010 can't be open with Excel Source/Destination in SSIS. To use XLSX in SSIS we need to install ACE driver and use OLE DB Source/Destination to read read/write data in XLSX files. ACE is a set of components that facilitate the transfer of data between existing Microsoft Office files such as Microsoft Office Access (\*.mdb and \*.accdb) files and Microsoft Office Excel 7 (\*.xls, \*.xlsx, and \*.xlsb) files to Microsoft SQL Server. If the SQL Server or your development box does not have ACE driver then install that 1st.

**Note:**

- The most efficient way to link to SQL Server is using an OLEDB connection manager.

Incorrect:

Not B:

- Address Windowing Extensions (AWE) is a Microsoft Windows application programming interface that allows a 32-bit software application to access more physical memory than it has virtual address space.

- Address Windowing Extensions (AWE) is a set of extensions that allows an application to quickly manipulate physical memory greater than 4GB. Certain data-intensive applications, such as database management systems and scientific and engineering software, need access to very large caches of data. In the case of very large data sets, restricting the cache to fit within an application's 2GB of user address space is a severe restriction.

#### **QUESTION 167**

Drag and Drop Questions

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.)

Install the Database Engine, Data Quality Client, and Data Quality Server.

Add the **DataStewards** user to the **dqs\_kb\_operator** role.

Run the Data Quality Server Installer.

Add the **DataStewards** user to the **dqs\_kb\_editor** role.

Add a user named **DataStewards** to the DQS\_MAIN database based on the **DataStewards** login.

**Answer:**

Install the Database Engine, Data Quality Client, and Data Quality Server.

Add the **DataStewards** user to the **dqs\_kb\_operator** role.

Run the Data Quality Server Installer.

Add the **DataStewards** user to the **dqs\_kb\_editor** role.

Add a user named **DataStewards** to the DQS\_MAIN database based on the **DataStewards** login.

Install the Database Engine, Data Quality Client, and Data Quality Server.

Run the Data Quality Server Installer.

Add a user named **DataStewards** to the DQS\_MAIN database based on the **DataStewards** login.

Add the **DataStewards** user to the **dqs\_kb\_editor** role.

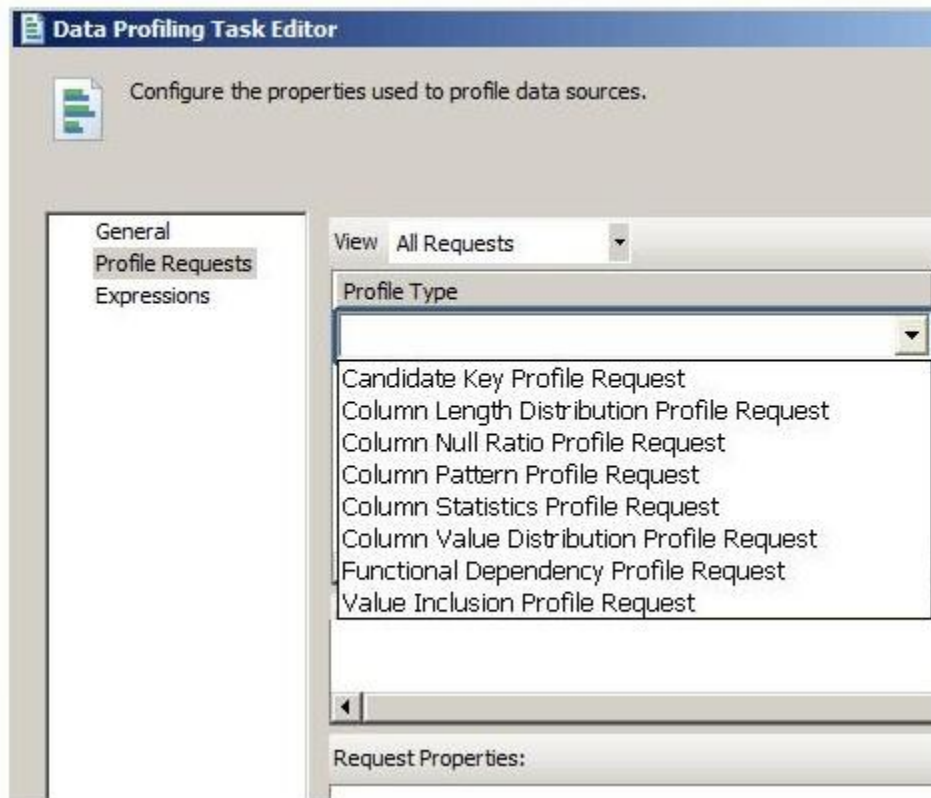
### QUESTION 168

#### Hotspot Questions

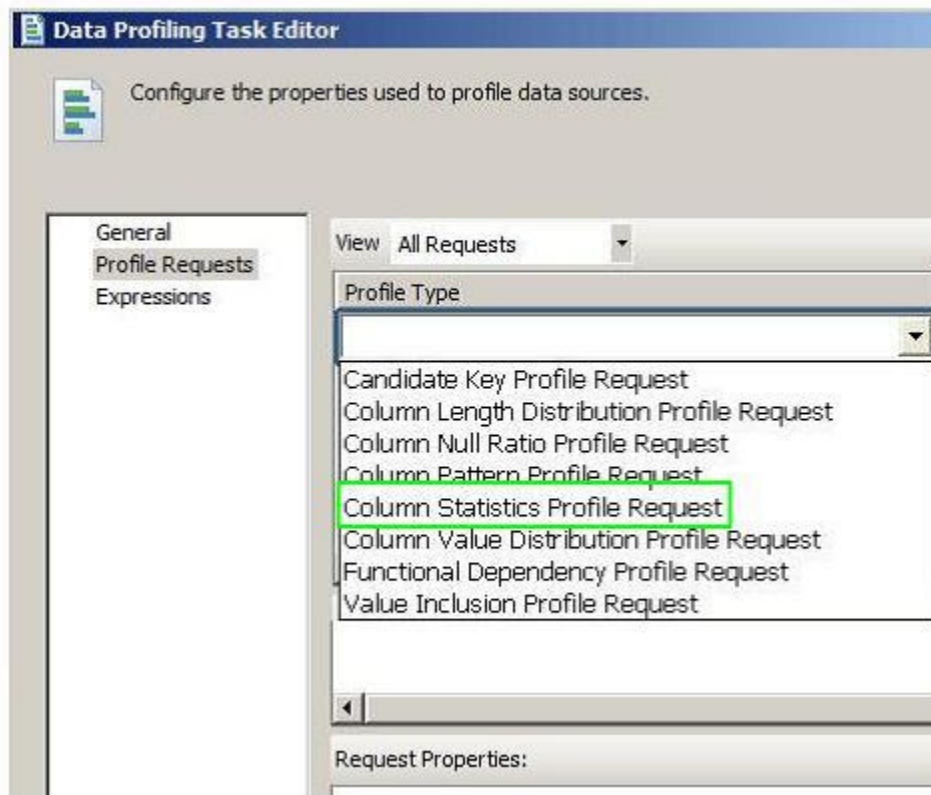
You are developing a SQL Server Integration Services (SSIS) package. You use a Data Profiling task to examine the data from a source system. You need to establish:

- 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:**



**QUESTION 169**

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 handlers.

**Answer: A**

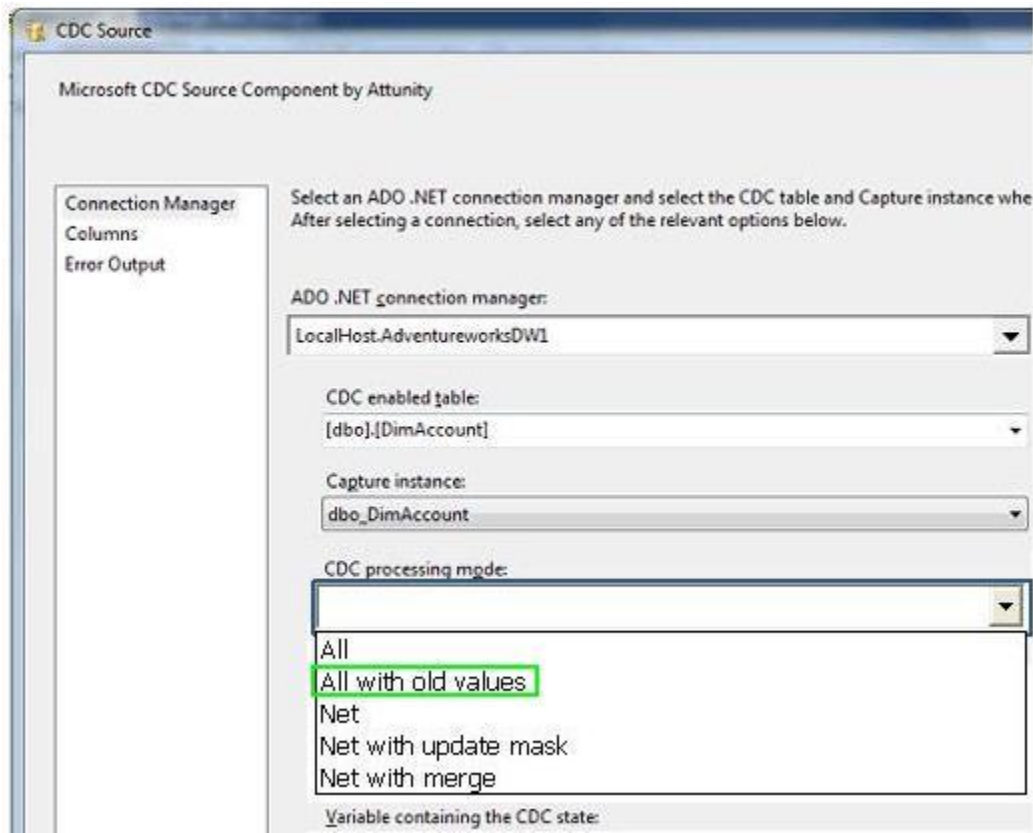
**QUESTION 170**

Hotspot Questions

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.)

The screenshot shows the 'CDC Source' dialog box. The 'CDC processing mode' dropdown is open, displaying the following options: 'All', 'All with old values', 'Net', 'Net with update mask', and 'Net with merge'. The other settings are: 'ADO .NET connection manager' is 'LocalHost.AdventureworksDW1', 'CDC enabled table' is '[dbo].[DimAccount]', and 'Capture instance' is 'dbo\_DimAccount'.

**Answer:**



**QUESTION 171**

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, view the job history for the SQL Server Agent job.
- B. After 10 executions, in SQL Server Management Studio, view the Execution Performance subsection of the All Executions report for the project.
- C. Enable logging to the Application Event Log in the package control flow for the On information event. After 10 executions, view the Application Event Log.
- D. Enable logging to an XML file in the package control flow for the OnPostExecute event. After 10 executions, view the XML file.

**Answer: B**

**QUESTION 172**

Hotspot Questions

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.)



Parameters | Connection Managers | **Advanced**

Property overrides

| Property path | Sensitive | Property value |
|---------------|-----------|----------------|
|---------------|-----------|----------------|

Add Edit Remove

Logging level: None

None  
Basic  
Performance  
Verbose

☐ Dump on errors ☐ 32-bit runtime

**Answer:**

The screenshot shows a software configuration window with three tabs: 'Parameters', 'Connection Managers', and 'Advanced'. The 'Advanced' tab is selected. It contains a section titled 'Property overrides' with a table that has three columns: 'Property path', 'Sensitive', and 'Property value'. The table is currently empty. Below the table are three buttons: 'Add', 'Edit', and 'Remove'. At the bottom of the window, there is a 'Logging level' dropdown menu with four options: 'None', 'Basic', 'Performance' (which is highlighted with a green border), and 'Verbose'. Below the dropdown are two checkboxes: 'Dump on errors' and '32-bit runtime', both of which are currently unchecked.

**QUESTION 173**

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:**

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

**QUESTION 174**

You are designing an extract, transform, and load (ETL) solution that loads data into dimension tables. The ETL process involves many transformation steps. You need to ensure that the design can provide:

- Auditing information for compliance and business user acceptance.
- Tracking and unique identification of records for troubleshooting and error correction.

What should you do?

- A. Create a SQL Common Language Runtime (SQLCLR) component that records all transformation steps in a Microsoft SharePoint list.
- B. Create a version control repository for the transformation steps in Team Foundation Server (TFS).
- C. Develop a custom data lineage solution.
- D. Develop a Data Quality Services (DQS) solution.

**Answer: C**

**QUESTION 175**

To support the implementation of new reports, Active Directory data will be downloaded to a SQL Server database by using a SQL Server Integration Services (SSIS) 2012 package. The following requirements must be met:

- All the user information for a given Active Directory group must be downloaded to a SQL Server table.
- The download process must traverse the Active Directory hierarchy recursively.

You need to configure the package to meet the requirements by using the least development effort. What should you use?

- A. script component
- B. custom component
- C. XML source
- D. script task

**Answer: A**

**QUESTION 176**

Hotspot Questions

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.)



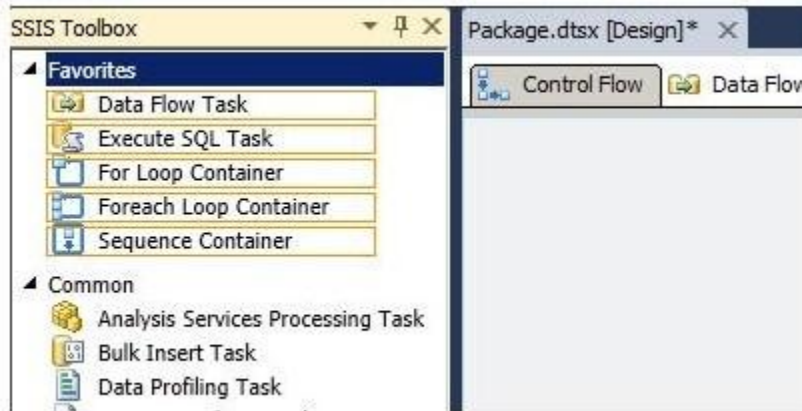
**Answer:**



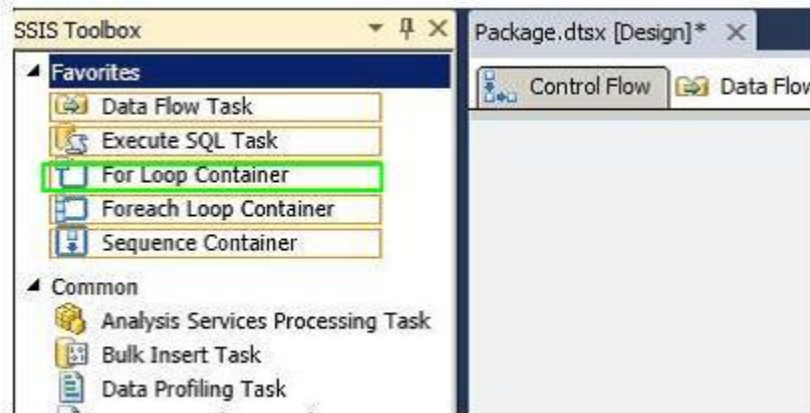
**QUESTION 177**

Hotspot Questions

You are developing a SQL Server Integration Services (SSIS) package. The package contains a user-defined variable named @Queue which has an initial value of 10. The package control flow contains many tasks that must repeat execution until the @Queue variable equals 0. You need to enable the tasks to be grouped together for repeat execution. Which item should you add to the package? (To answer, select the appropriate item in the answer area.)



**Answer:**



**QUESTION 178**

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:**

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

**QUESTION 179**

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:**

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

**QUESTION 180**

A SQL Server Integration Services (SSIS) 2012 package currently downloads sales data from a Windows Azure SQL Database database. To improve sales data accuracy, exchange rates must be downloaded daily from a public HTTP website instead of from a weekly flat file. The public



website hosts a comma separated values (CSV) file that contains one row per currency. You need to download the CSV file to the environment. What should you use to retrieve the document from the website?

- A. a Script component
- B. a Web Service task
- C. a Web Service source
- D. a Script task

**Answer: D**

**QUESTION 181**

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. Source Assistant
- B. Destination Assistant
- C. Audit
- D. Script Component

**Answer: D**

**Explanation:**

Script Component Transformation: The transformation that uses script to extract, transform, or load data.

**Note:**

- SQL Server Integration Services provides three different types of data flow components: sources, transformations, and destinations. Sources extract data from data stores such as tables and views in relational databases, files, and Analysis Services databases. Transformations modify, summarize, and clean data. Destinations load data into data stores or create in-memory datasets. Incorrect:

Not A, Not B: No Data Transforms called Source Assistant or Destination Assistant.

Not C: Not related to auditing.

**QUESTION 182**

You are administering SQL Server Integration Services (SSIS) permissions on a production server that runs SQL Server 2012. Package developers in your company must have permission to perform the following tasks only on their own projects:

- View projects and packages
- View Environments
- Validate packages
- Execute packages

You need to grant rights to the developers without assigning unnecessary privileges. What should you do? (Each correct answer presents part of a solution. Choose all that apply.)

- A. Add developer logins to the db\_issltduser role in the msdb database.
- B. Add developer logins to the db\_issoperator role in the msdb database.
- C. Grant Execute permission in the projects for the developer logins.
- D. Grant Read permission in the SSIS catalog folder, the projects, and the Environments.
- E. Add developer logins to the ssis\_admin role in the SSISDB database.

F. Grant Modify permission in the projects for the developer logins.

**Answer: BD**

**QUESTION 183**

To support the implementation of new reports, Active Directory data will be downloaded to a SQL Server database by using a SQL Server Integration Services (SSIS) 2012 package. The following requirements must be met:

- All the user information for a given Active Directory group must be downloaded to a SQL Server table.
- The download process must traverse the Active Directory hierarchy recursively.

You need to configure the package to meet the requirements by using the least development effort. Which item should you use?

- A. Script task
- B. Script component configured as a transformation
- C. Script component configured as a source
- D. Script component configured as a destination

**Answer: C**

**QUESTION 184**

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. Raw File
- B. ODBC
- C. Bulk Insert
- D. OLE DB

**Answer: D**

**QUESTION 185**

You are implementing a 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 date stored in a control table of the Windows Azure SQL Database database 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. an Execute SQL task
- C. a Script component

D. a Script task

**Answer: A**

**Explanation:**

<http://msdn.microsoft.com/en-us/library/ms141214.aspx>

<http://blog.sqltechie.com/2012/11/expression-task-in-ssis-2012.html>

**QUESTION 186**

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.

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.

**QUESTION 187**

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

**Answer: D**

**QUESTION 188**

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 4GB of RAM.

D. Install Microsoft SharePoint Server 2010 Standard Edition.

**Answer: B**

**QUESTION 189**

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?

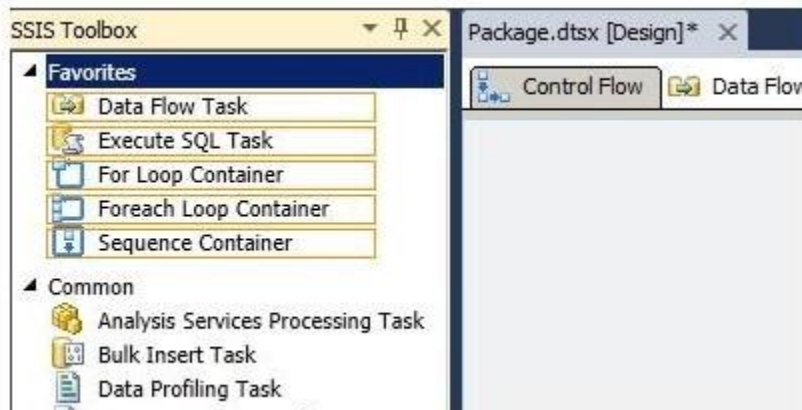
- A. Design two tables.  
One table will hold the original customer surrogate key and attributes.  
The other table will hold the most recent customer surrogate key and attributes.
- B. Design a table as a Type 2 SCD dimension that holds the surrogate key for each customer and its attributes.  
Limit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

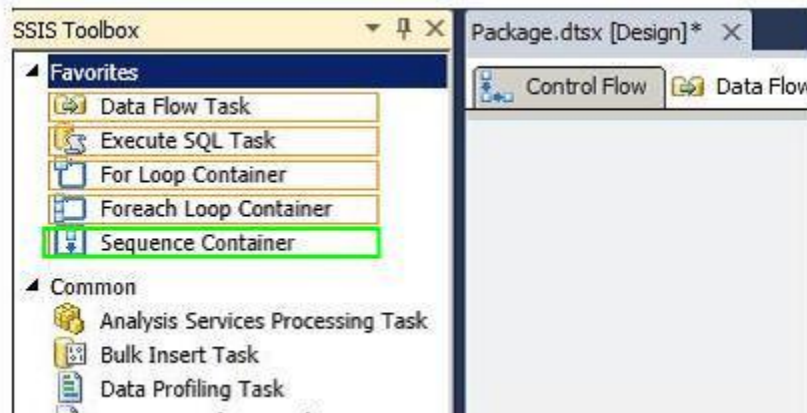
**QUESTION 190**

Hotspot Questions

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:**



### QUESTION 191

Drag and Drop Questions

You administer a large and complex SQL Server Integration Services (SSIS) solution in the SSIS catalog. You are instructed to execute a package by using PowerShell. You need to create the correct PowerShell command. How should you place the four code segments in sequence? (To answer, move the appropriate code segments from the list of code segments to the answer area and arrange them in the correct order.)

|   |   |
|---|---|
| .Folders["SSIS70463"]   |   |
| .Projects["Production"]   |   |
| .Packages["Package.dtsx"].Execute("false", \$ProductionEnvironment) | ) |
| \$SSIS2012.Catalogs["SSISDB"]                                       | ) |

**Answer:**

|   |   |
|---|---|
| .Folders["SSIS70463"]   | \$SSIS2012.Catalogs["SSISDB"]                                       |
| .Projects["Production"]   | .Folders["SSIS70463"]   |
| .Packages["Package.dtsx"].Execute("false", \$ProductionEnvironment) | .Projects["Production"]   |
| \$SSIS2012.Catalogs["SSISDB"]                                       | .Packages["Package.dtsx"].Execute("false", \$ProductionEnvironment) |

### QUESTION 192

Hotspot Questions

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.)



Parameterize - \_\_sql2012

Specify the property to parameterize and the parameter to use. An expression referencing the parameter will be created on the property.

Property:

- ConnectionString
- Description
- InitialCatalog
- Name
- Password
- RetainSameConnection
- ServerName
- UserName

☐ Do not use parameter

☒ Use existing parameter

- \$Project::DatabaseName
- \$Project::ProcessDate
- \$Project::ServerName
- \$Project::StartDate

☐ Create new parameter

Name:

Description:

Value:

Scope:

- ☐ Package
- ☒ Project

**Answer:**

Parameterize - \_\_sql2012

Specify the property to parameterize and the parameter to use. An expression referencing the parameter will be created on the property.

Property:   
ConnectionString  
Description  
InitialCatalog  
Name  
Password  
RetainSameConnection  
**ServerName**  
UserName

☐ Do not use parameter

☒ Use existing parameter

\$Project:::DatabaseName  
\$Project:::ProcessDate  
**\$Project:::ServerName**  
\$Project:::StartDate

☐ Create new parameter

Name: \_\_sql2012\_ConnectionString

Description:

Value: Data Source=.\sql2012;Initial Catalog=testdb;Provider=SQL

Scope:   
☐ Package  
☒ Project

**QUESTION 193**

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**

**QUESTION 194**

Hotspot Questions

You are the Master Data Services (MDS) administrator at your company. An existing user needs to integrate data from other systems into MDS. You need to give the user permissions to integrate data. 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:**



**QUESTION 195**

You are developing a SQL Server Integration Services (SSIS) package. The package uses a data flow task to source data from a SQL Server database for loading into a dimension table in a data warehouse. You need to create a separate data flow path for data that has been modified since it was last processed. Which data flow components should you use to identify modified data? (Each correct answer presents a complete solution. Choose all that apply.)

- A. Multicast
- B. Data Conversion

- C. Lookup
- D. Slowly Changing Dimension
- E. Aggregate

**Answer: AC**

**Explanation:**

A: The transformation that distributes data sets to multiple outputs. The transformation that distributes data sets to multiple outputs.

C: Lookup Transformation. The transformation that looks up values in a reference table using an exact match.

**Note:**

SQL Server Integration Services provides three different types of data flow components: sources, transformations, and destinations. Sources extract data from data stores such as tables and views in relational databases, files, and Analysis Services databases. Transformations modify, summarize, and clean data. Destinations load data into data stores or create in-memory datasets. Incorrect:

Not B: Data Conversion Transformation. The transformation that converts the data type of a column to a different data type.

Not D: Slowly Changing Dimension Transformation. The transformation that configures the updating of a slowly changing dimension.

Not E: The Aggregate transformation applies aggregate functions, such as Average, to column values and copies the results to the transformation output. Besides aggregate functions, the transformation provides the GROUP BY clause, which you can use to specify groups to aggregate across.

#### **QUESTION 196**

Drag and Drop Questions

You are the administrator for a Data Quality Server. You are adding a user who must have permission to:

- Edit and execute a project
- View the activity monitoring data

This user must not be able to:

- Perform any kind of knowledge management
- Create or change a knowledge base

Terminate an activity or perform administrative duties. You need to develop a Transact-SQL (T-SQL) script to meet these requirements. What should you do? (To answer, drag the appropriate code segment or segments to the correct location or locations in the answer area.)

|                     |   |
|---------------------|---|
| [msdb]              | USE [master]  |
| [master]            | GO  |
| [DQS_MAIN]          | CREATE LOGIN [MYDOMAIN\dqsuser] FROM WINDOWS WITH DEFAULT_DATABASE=[ma: |
| [dqs_kb_editor]     | GO  |
| [dqs_kb_operator]   | USE [ ]   |
| [dqs_administrator] | GO  |
|                     | CREATE USER [MYDOMAIN\dqsuser] FOR LOGIN [MYDOMAIN\dqsuser]             |
|                     | GO  |
|                     | ALTER ROLE [ ] ADD MEMBER [MYDOMAIN\dqsuser]                            |
|                     | GO  |

**Answer:**

|  |  |
|--|--|
| <div>[msdb]</div> <div>[master]</div> <div>[DQS_MAIN]</div> <div>[dqs_kb_editor]</div> <div>[dqs_kb_operator]</div> <div>[dqs_administrator]</div> | <pre>USE [master] GO CREATE LOGIN [MYDOMAIN\dqsuser] FROM WINDOWS WITH DEFAULT_DATABASE=[ma: GO USE [DQS_MAIN] GO CREATE USER [MYDOMAIN\dqsuser] FOR LOGIN [MYDOMAIN\dqsuser] GO ALTER ROLE [dqs_kb_operator] ADD MEMBER [MYDOMAIN\dqsuser] GO</pre> |
|--|--|

### QUESTION 197

#### Drag and Drop Questions

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.)

|   |                           |
|---|---------------------------|
| Configure the MsDtsSrvr.ini.xml file to use the msdb database.                        | <div>(</div> <div>(</div> |
| Install Database Engine Services, Integration Services, and Client Tools SDK.         |                           |
| Install Microsoft NET Framework 3.5 SP1.  |                           |
| Install Database Engine Services, Integration Services, and SQL Server Data Tools.    |                           |
| Install Database Engine Services, Integration Services, and Management Tools – Basic. |                           |
| Create the SSIS catalog.  |                           |
| Install Database Engine Services and Integration Services.                            |                           |

#### Answer:

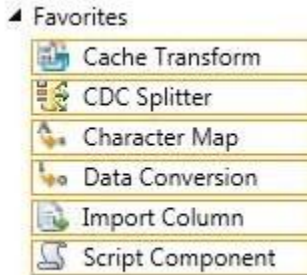
|   |                           |
|---|---------------------------|
| Configure the MsDtsSrvr.ini.xml file to use the msdb database.                        | <div>(</div> <div>(</div> |
| Install Database Engine Services, Integration Services, and Client Tools SDK.         |                           |
| Install Microsoft NET Framework 3.5 SP1.  |                           |
| Install Database Engine Services, Integration Services, and SQL Server Data Tools.    |                           |
| Install Database Engine Services, Integration Services, and Management Tools – Basic. |                           |
| Create the SSIS catalog.  |                           |
| Install Database Engine Services and Integration Services.                            |                           |

### QUESTION 198

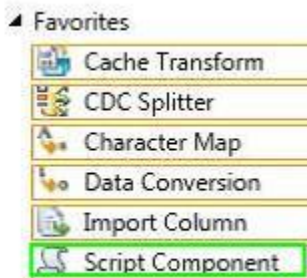
#### Hotspot Questions



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.)




**Answer:**



### QUESTION 199

Hotspot Questions

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:**

| Name      | Data type | Value | Sensitive | Required | Description |
|-----------|-----------|-------|-----------|----------|-------------|
| FinConStr | String    |       | True      | True     |             |

### QUESTION 200

#### Hotspot Questions

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.)

Microsoft SQL Server 2012 Master Data Services

Manage Users Manage Groups

Users

username-BMA\username

Edit

- General
- Membership
- Functions
- Models
- Hierarchy members

**Answer:**



#### QUESTION 201

Drag and Drop Questions

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:**

Prerequisite Choices

Rule Editor

| Domain | Similarity | Weight | Prerequisite |
|--------|------------|--------|--------------|
| Artist | Similar    | 0      | Yes          |
| Song   | Similar    | 100    | No           |

### QUESTION 202

#### Hotspot Questions

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:

Evaluation operation

Constraint  
Expression  
Expression and Constraint  
Expression or Constraint

☐ Multiple constraints

#### Answer:

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:

Evaluation operation

Constraint  
Expression  
Expression and Constraint  
Expression or Constraint

☐ Multiple constraints

### QUESTION 203

#### Drag and Drop Questions

You are developing a SQL Server Integration Services (SSIS) package that downloads data from

a Windows Azure SQL Database database. A stored procedure will be called in an Execute SQL task by using an ODBC connection. This stored procedure has only the @CustomerID parameter of type INT. A project parameter named CustID will be mapped to the stored procedure parameter @CustomerID. You need to ensure that the value of the CustID parameter is passed to the @CustomerID stored procedure parameter. In the Parameter Mapping tab of the Execute SQL task editor, how should you configure the parameter? (To answer, drag the appropriate option or options to the correct location or locations in the answer area.)

|                   |                 |  |
|-------------------|-----------------|--|
| @CustomerID       | Variable Name:  |  |
| \$Project::CustID | Direction:      |  |
| User::CustID      | Data Type:      |  |
| Input             | Parameter Name: |  |
| Output            |                 |  |
| SQL_INTEGER       |                 |  |
| SQL_SMALLINT      |                 |  |
| SQL_VARCHAR       |                 |  |
| 0                 |                 |  |
| 1                 |                 |  |

**Answer:**

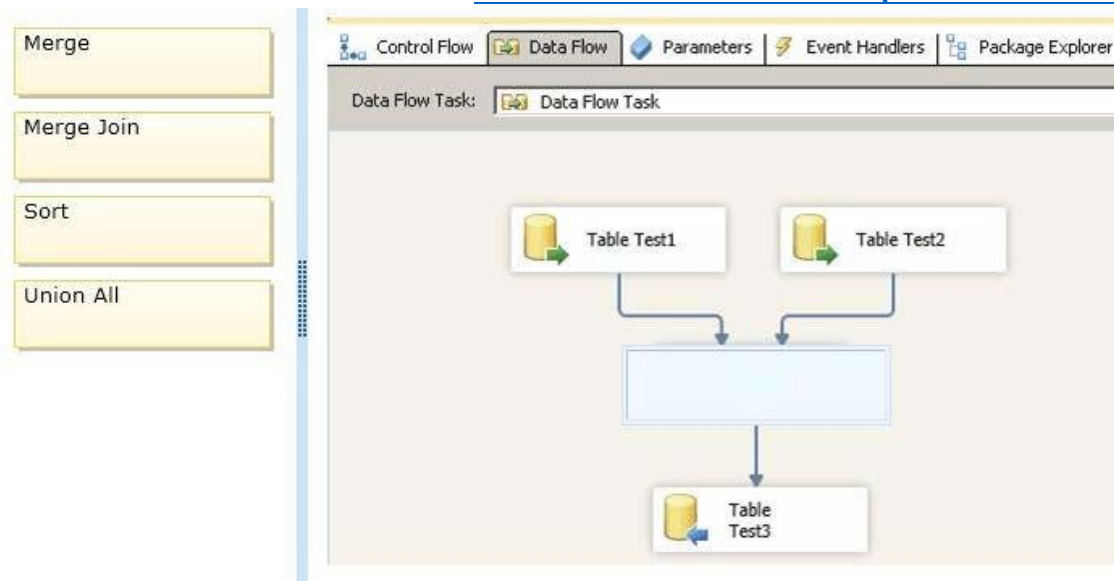
|                   |                 |                   |
|-------------------|-----------------|-------------------|
| @CustomerID       | Variable Name:  | \$Project::CustID |
| \$Project::CustID | Direction:      | Input             |
| User::CustID      | Data Type:      | SQL_INTEGER       |
| Input             | Parameter Name: | 1                 |
| Output            |                 |                   |
| SQL_INTEGER       |                 |                   |
| SQL_SMALLINT      |                 |                   |
| SQL_VARCHAR       |                 |                   |
| 0                 |                 |                   |
| 1                 |                 |                   |

#### QUESTION 204

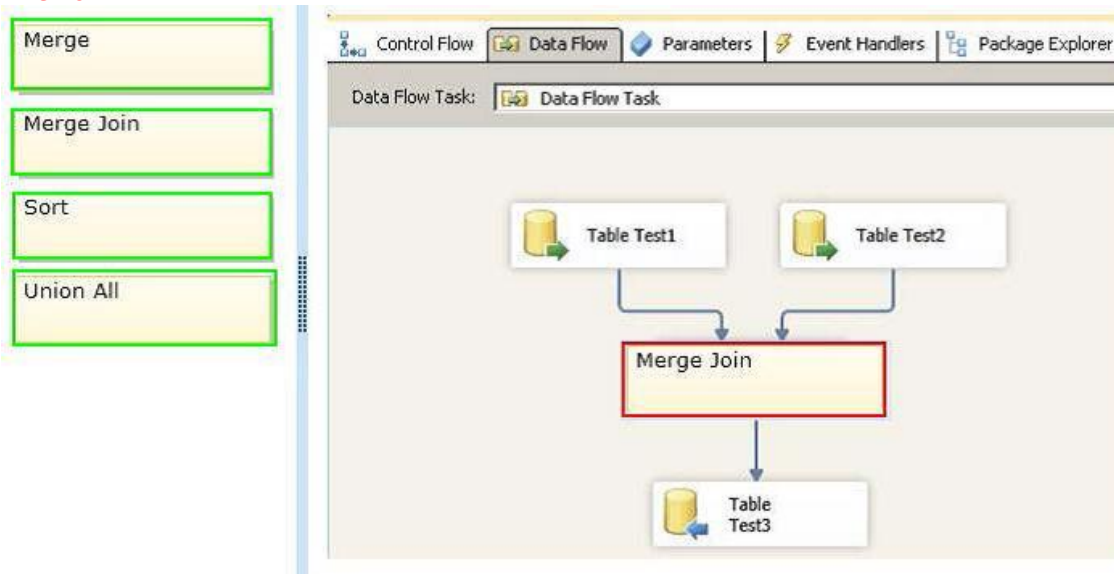
Drag and Drop Questions

You are developing a SQL Server Integration Services (SSIS) package that loads data into a data warehouse hosted on Windows Azure SQL Database. You must combine two data sources together by using the ProductID column to provide complete details for each record. The data retrieved from each data source is sorted in ascending order by the ProductID column. You need to develop a data flow that imports the data while meeting the requirements. How should you develop the data flow? (To answer, drag the appropriate transformation from the list of transformations to the correct location in the answer area.)





**Answer:**



#### QUESTION 205

Drag and Drop Questions

You are developing a SQL Server Integration Services (SSIS) package. The package contains several tasks that must repeat until an expression evaluates to FALSE. You need to add and configure a container to enable this design. 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.)

|  |   |   |
|--|---|---|
| Configure the <b>EvalExpression</b> property.  | ) | ( |
| Configure the <b>Enumerator</b> property.      |   |   |
| Open the Data Flow designer of the package.    |   |   |
| Add and edit a For Loop container.             |   |   |
| Add and edit a Foreach Loop container.         |   |   |
| Open the Control Flow designer of the package. |   |   |

**Answer:**

|  |   |   |
|--|---|---|
| Configure the <b>EvalExpression</b> property.  | ) | ( |
| Configure the <b>Enumerator</b> property.      |   |   |
| Open the Data Flow designer of the package.    |   |   |
| Add and edit a For Loop container.             |   |   |
| Add and edit a Foreach Loop container.         |   |   |
| Open the Control Flow designer of the package. |   |   |

#### QUESTION 206

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 view on the sales table.
- B. Partition the fact tables by day.
- C. Create a surrogate key for the time dimension.
- D. Change the level of granularity in both fact tables to be the same.

**Answer: D**

**Explanation:**

In order to join the facts table in time dimension we cannot have two different time measures (time and day).

#### QUESTION 207

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**

**QUESTION 208**

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.

**QUESTION 209**

You administer a SQL Server Integration Services (SSIS) solution in the SSIS catalog. A SQL Server Agent job is used to execute a package daily with the basic logging level. Recently, the package execution failed because of a primary key violation when the package inserted data into the destination table. You need to identify all previous times that the package execution failed because of a primary key violation. 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 information. Create an error output on each data flow destination that writes OnError

event text to the table.

- R. Create a table to store error information. Create an error output on each data flow destination that writes OnTaskFailed event text to the table.

**Answer: E**

**QUESTION 210**

You are developing a SQL Server Integration Services (SSIS) package to load data into a Windows Azure SQL Database database. The package consists of several data flow tasks. The package has the following auditing requirements:

- If a data flow task fails, a Transact-SQL (T-SQL) script must be executed.
- The T-SQL script must be executed only once per data flow task that fails, regardless of the nature of the error.

You need to ensure that auditing is configured to meet these requirements. 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 information. Create an error output on each data flow destination that writes OnTaskFailed event text to the table.

**Answer: C**

**QUESTION 211**

You are developing a SQL Server Integration Services (SSIS) project with multiple packages to copy data to a Windows Azure SQL Database database. An automated process must validate all related Environment references, parameter data types, package references, and referenced assemblies. The automated process must run on a regular schedule. You need to establish the automated validation process by using the least amount of administrative effort. 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 information. Create an error output on each data flow destination that writes OnError event text to the table.
- R. Create a table to store error information. Create an error output on each data flow destination that writes OnTaskFailed event text to the table.

**Answer: N**

#### **QUESTION 212**

You are developing a SQL Server Integration Services (SSIS) project by using the Project Deployment Model. All packages in the project must log custom messages. You need to produce reports that combine the custom log messages with the system generated log messages. 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 information. Create an error output on each data flow destination that writes OnError event text to the table.
- R. Create a table to store error information. Create an error output on each data flow destination that writes OnTaskFailed event text to the table.

**Answer: G**

#### **QUESTION 213**

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.

<http://msdn.microsoft.com/en-us/library/hh758656.aspx>


#### QUESTION 214

Drag and Drop Questions

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.)

|   |   |  |
|---|---|--|
| Create an <b>OnError</b> event handler for the Data Flow task.  |  |  |
| Define an expression for the <b>ActionAtEvent</b> property.   |   |  |
| Add a Notify Operator task.   |   |  |
| Add a WMI Event Watcher task.   |   |  |
| In the <b>ReadOnlyVariables</b> list, select <b>System::ErrorDescription</b> and add conditional logic to invoke the <b>FireError</b> method. |   |  |
| Add a Script task.  |   |  |

**Answer:**

|   |   |   |
|---|---|---|
| Create an <b>OnError</b> event handler for the Data Flow task.  |  | Create an <b>OnError</b> event handler for the Data Flow task.  |
| Define an expression for the <b>ActionAtEvent</b> property.   |   |   |
| Add a Notify Operator task.   |   | Add a Script task.  |
| Add a WMI Event Watcher task.   |   |   |
| In the <b>ReadOnlyVariables</b> list, select <b>System::ErrorDescription</b> and add conditional logic to invoke the <b>FireError</b> method. |   | In the <b>ReadOnlyVariables</b> list, select <b>System::ErrorDescription</b> and add conditional logic to invoke the <b>FireError</b> method. |
| Add a Script task.  |   |   |

#### QUESTION 215

You need to extract data from delimited text files. What connection manager type would you choose?

- A. A Flat File connection manager
- B. An OLE DB connection manager
- C. An ADO.NET connection manager

D. A File connection manager

**Answer: A**

**QUESTION 216**

Some of the data your company processes is sent in from partners via email. How would you configure an SMTP connection manager to extract files from email messages?

- A. In the SMTP connection manager, configure the OperationMode setting to Send And Receive.
- B. It is not possible to use the SMTP connection manager in this way, because it can only be used by SSIS to send email messages.
- C. The SMTP connection manager supports sending and receiving email messages by default, so no additional configuration is necessary.
- D. It is not possible to use the SMTP connection manager for this, use the IMAP (Internet Message Access Protocol) connection manager instead.

**Answer: B**

**QUESTION 217**

You need to extract data from a table in a SQL Server 2012 database. What connection manager types can you use? (Choose all that apply.)

- A. An ODBC connection manager
- B. An OLE DB connection manager
- C. A File connection manager
- D. An ADO.NET connection manager

**Answer: ABD**

**QUESTION 218**

In your SSIS solution, you need to load a large set of rows into the database as quickly as possible. The rows are stored in a delimited text file, and only one source column needs its data type converted from String (used by the source column) to Decimal (used by the destination column). What control flow task would be most suitable for this operation?

- A. The File System task would be perfect in this case, because it can read data from files and can be configured to handle data type conversions.
- B. The Bulk Insert task would be the most appropriate, because it is the quickest and can handle data type conversions.
- C. The data flow task would have to be used, because the data needs to be transformed before it can be loaded into the table.
- D. No single control flow task can be used for this operation, because the data needs to be extracted from the source file, transformed, and then loaded into the destination table.  
At least three different tasks would have to be used -- the Bulk Insert task to load the data into a staging database, a Data Conversion task to convert the data appropriately, and finally, an Execute SQL task to merge the transformed data with existing destination data.

**Answer: C**

**QUESTION 219**

A part of your data consolidation process involves extracting data from Excel workbooks. Occasionally, the data contains errors that cannot be corrected automatically. How can you handle this problem by using SSIS?

- A. Redirect the failed data flow task to an External Process task, open the problematic Excel file in Excel, and prompt the user to correct the file before continuing the data consolidation process.
  - B. Redirect the failed data flow task to a File System task that moves the erroneous file to a dedicated location where an information worker can correct it later.
  - C. If the error cannot be corrected automatically, there is no way for SSIS to continue with the automated data consolidation process.
  - D. None of the answers above are correct.
- Due to Excel's strict data validation rules, an Excel file cannot ever contain erroneous data.

**Answer: B**

**QUESTION 220**

In your ETL process, there are three external processes that need to be executed in sequence, but you do not want to stop execution if any of them fails. Can this be achieved by using precedence constraints? If so, which precedence constraints can be used?

- A. No, this cannot be achieved just by using precedence constraints.
- B. Yes, this can be achieved by using completion precedence constraints between the first and the second and between the second and the third Execute Process tasks, and by using a success precedence constraint between the third Execute Process task and the following task.
- C. Yes, this can be achieved by using completion precedence constraints between the first and the second, between the second and the third, and between the third Execute Process task and the following task.
- D. Yes, this can be achieved by using failure precedence constraints between the first and the second, and between the second and the third Execute Process tasks, and by using a completion precedence constraint between the third Execute Process task and the following task.

**Answer: B**

**QUESTION 221**

.....

**QUESTION 266**

You have a Microsoft Azure SQL Data Warehouse instance that must be available six months a day for reporting. You need to pause the compute resources when the instance is not being used.

Solution: You use the Azure portal.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: A**

**QUESTION 267**

You have a Microsoft SQL Server data warehouse instance that supports several client applications. The data warehouse includes the following tables: Dimension.SalesTerritory, Dimension.Customer, Dimension.Date, Fact.Ticket, and Fact.Order. The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. The Fact.Order table is optimized for weekly reporting, but the company wants to change it daily. The Fact.Order table is loaded by using an ETL process. Indexes have been added to the table over time, but the presence of these indexes slows data loading. All data in the data warehouse is stored on a shared SAN. All tables are in a database named DB1. You have a second database named DB2 that contains copies of production data for a development environment. The data warehouse has grown and the cost of storage has

increased. Data older than one year is accessed infrequently and is considered historical. You have the following requirements:

- You are not permitted to make changes to the client applications.
- You need to optimize the storage for the data warehouse.

What change should you make?

- A. Partition the Fact.Order table, and move historical data to new filegroups on lower-cost storage.
- B. Create new tables on lower-cost storage, move the historical data to the new tables, and then shrink the database.
- C. Remove the historical data from the database to leave available space for new data.
- D. Move historical data to new tables on lower-cost storage.

**Answer: A**

**Explanation:**

- Create the load staging table in the same filegroup as the partition you are loading.

- Create the unload staging table in the same filegroup as the partition you are deleting.

From scenario: Data older than one year is accessed infrequently and is considered historical.

<https://blogs.msdn.microsoft.com/sqlcat/2013/09/16/top-10-best-practices-for-building-a-large-scale-relational-data-warehouse/>

#### **QUESTION 268**

You have the following line-of-business solutions:

- ERP System
- Online WebStore
- Partner Extranet

One or more Microsoft SQL Server instances support each solution. Each solution has its own product catalog. You have an additional server that hosts SQL Server Integration Services (SSIS) and a data warehouse. You populate the data warehouse with data from each of the line-of-business solutions. The data warehouse does not store primary key values from the individual source tables. The database for each solution has a table named Products that stored product information. The Products table in each database uses a separate and unique key for product records. Each table shares a column named ReferenceNr between the databases. This column is used to create queries that involve more than once solution. You need to load data from the individual solutions into the data warehouse nightly. The following requirements must be met:

- If a change is made to the ReferenceNr column in any of the sources, set the value of IsDisabled to True and create a new row in the Products table.
- If a row is deleted in any of the sources, set the value of IsDisabled to True in the data warehouse.

**Solution:** Perform the following actions:

- Enable the Change Tracking feature for the Products table in the three source databases.
- Query the CHANGETABLE function from the sources for the deleted rows.
- Set the IsDisabled column to True on the data warehouse Products table for the listed rows.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

We must check for updated rows, not just deleted rows.

<https://www.timmitchell.net/post/2016/01/18/getting-started-with-change-tracking-in-sql-server/>

**QUESTION 269**

You have a Microsoft SQL Server data warehouse instance that supports several client applications. The data warehouse includes the following tables: Dimension.SalesTerritory, Dimension.Customer, Dimension.Date, Fact.Ticket, and Fact.Order. The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. The Fact.Order table is optimized for weekly reporting, but the company wants to change it daily. The Fact.Order table is loaded by using an ETL process. Indexes have been added to the table over time, but the presence of these indexes slows data loading. All data in the data warehouse is stored on a shared SAN. All tables are in a database named DB1. You have a second database named DB2 that contains copies of production data for a development environment. The data warehouse has grown and the cost of storage has increased. Data older than one year is accessed infrequently and is considered historical. You have the following requirements:

- Implement table partitioning to improve the manageability of the data warehouse and to avoid the need to repopulate all transactional data each night. Use a partitioning strategy that is as granular as possible.
- Partition the Fact.Order table and retain a total of seven years of data.
- Partition the Fact.Ticket table and retain seven years of data. At the end of each month, the partition structure must apply a sliding window strategy to ensure that a new partition is available for the upcoming month, and that the oldest month of data is archived and removed.
- Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.
- Incrementally load all tables in the database and ensure that all incremental changes are processed.
- Maximize the performance during the data loading process for the Fact.Order partition.
- Ensure that historical data remains online and available for querying.
- Reduce ongoing storage costs while maintaining query performance for current data.

You are not permitted to make changes to the client applications. You need to implement the data partitioning strategy. How should you partition the Fact.Order table?

- A. Create 17,520 partitions.
- B. Use a granularity of two days.
- C. Create 2,557 partitions.
- D. Create 730 partitions.

**Answer: C**

**Explanation:**

We create one partition for each day. 7 years times 365 days is 2,555. Make that 2,557 to provide for leap years.

From scenario: Partition the Fact.Order table and retain a total of seven years of data. Maximize the performance during the data loading process for the Fact.Order partition.

**QUESTION 270**

You have a Microsoft Azure SQL Data Warehouse instance that must be available six months a day for reporting. You need to pause the compute resources when the instance is not being used. Solution: You use SQL Server Configuration Manager.

Does the solution meet the goal?

- A. Yes
- B. No



**Answer: B**

**Explanation:**

To pause a SQL Data Warehouse database, use any of these individual methods:

- Pause compute with Azure portal
- Pause compute with PowerShell
- Pause compute with REST APIs

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-compute-overview>

#### **QUESTION 271**

You have the following line-of-business solutions:

- ERP System
- Online WebStore
- Partner Extranet

One or more Microsoft SQL Server instances support each solution. Each solution has its own product catalog. You have an additional server that hosts SQL Server Integration Services (SSIS) and a data warehouse. You populate the data warehouse with data from each of the line-of-business solutions. The data warehouse does not store primary key values from the individual source tables. The database for each solution has a table named Products that stored product information. The Products table in each database uses a separate and unique key for product records. Each table shares a column named ReferenceNr between the databases. This column is used to create queries that involve more than once solution. You need to load data from the individual solutions into the data warehouse nightly. The following requirements must be met:

- If a change is made to the ReferenceNr column in any of the sources, set the value of IsDisabled to True and create a new row in the Products table.
- If a row is deleted in any of the sources, set the value of IsDisabled to True in the data warehouse.

**Solution:** Perform the following actions:

- Enable the Change Tracking for the Product table in the source databases.
- Query the CHANGETABLE function from the sources for the updated rows.
- Set the IsDisabled column to True for the listed rows that have the old ReferenceNr value.
- Create a new row in the data warehouse Products table with the new ReferenceNr value.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

We must check for deleted rows, not just updates rows.

<https://www.timmitcheil.net/post/2016/01/18/getting-started-with-change-tracking-in-sql-server/>

#### **QUESTION 272**

You have the following line-of-business solutions:

- ERP System
- Online WebStore
- Partner Extranet

One or more Microsoft SQL Server instances support each solution. Each solution has its own product catalog. You have an additional server that hosts SQL Server Integration Services (SSIS) and a data warehouse. You populate the data warehouse with data from each of the line-of-business solutions. The data warehouse does not store primary key values from the individual source tables. The database for each solution has a table named Products that stored product

information. The Products table in each database uses a separate and unique key for product records. Each table shares a column named ReferenceNr between the databases. This column is used to create queries that involve more than once solution. You need to load data from the individual solutions into the data warehouse nightly. The following requirements must be met:

- If a change is made to the ReferenceNr column in any of the sources, set the value of IsDisabled to True and create a new row in the Products table.
- If a row is deleted in any of the sources, set the value of IsDisabled to True in the data warehouse.

**Solution:** Perform the following actions:

- Enable the Change Tracking for the Product table in the source databases.
- Query the `cdc.fn_cdc_get_all_changes_capture_dbo_products` function from the sources for updated rows.
- Set the IsDisabled column to True for rows with the old ReferenceNr value.
- Create a new row in the data warehouse Products table with the new ReferenceNr value.

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

We must also handle the deleted rows, not just the updated rows.

<https://solutioncenter.apexsql.com/enable-use-sql-server-change-data-capture/>

### **QUESTION 273**

You have a data quality project that focuses on the Products catalog for the company. The data includes a product reference number. The product reference should use the following format:

- Two letters followed by an asterisk and then four or five numbers. An example of a valid number is XX\*55522.
- Any reference number that does not conform to the format must be rejected during the data cleansing.

You need to add a Data Quality Services (DQS) domain rule in the Products domain. Which rule should you use?

- A. value matches pattern `ZA*9876[5]`
- B. value matches pattern `AZ[*]1234[5]`
- C. value matches regular expression `AZ[*]1234[5]`
- D. value matches pattern `[a-zA-Z][a-zA-Z]*[0-9][0-9][0-9][0-9]`

**Answer: A**

**Explanation:**

For a pattern matching rule:

- Any letter (A...Z) can be used as a pattern for any letter; case insensitive
- Any digit (0...9) can be used as a pattern for any digit
- Any special character, except a letter or a digit, can be used as a pattern for itself
- Brackets, [], define optional matching

This rule implies that the data will contain three parts: any three letters followed by a colon (:), which is again followed by any four digits.

**QUESTION 274**

You have a Microsoft Azure SQL Data Warehouse instance that must be available six months a day for reporting. You need to pause the compute resources when the instance is not being used. Solution: You use SQL Server Management Studio (SSMS).

Does the solution meet the goal?

- A. Yes
- B. No

**Answer: B**

**Explanation:**

To pause a SQL Data Warehouse database, use any of these individual methods:


- Pause compute with Azure portal
- Pause compute with PowerShell
- Pause compute with REST APIs

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-compute-overview>

**QUESTION 275**

.....

**Get Complete Version Exam 70-767 Dumps with VCE and PDF Here**



<https://www.passleader.com/70-767.html>