Microsoft

70-767 Exam

Microsoft Implementing a SQL Data Warehouse Exam

Questions & Answers
Demo

Version: 14.0

Question: 1

DRAG DROP
Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.
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 partitioning for the Fact.Ticket table.
Which three actions should you perform in sequence? To answer, drag the appropriate actions to the correct locations. Each action 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.
NOTE: More than one combination of answer choices is correct. You will receive credit for any of the correct combinations you select.

Actions

INSERT SELECT

MERGE

SWITCH

DELETE

SPLIT

Answer area

First action   **Second action**

Action

Action   Action

Answer:

Answer area

First action   **Second action**

SWITCH

MERGE   SPLIT

From scenario: - 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.

The detailed steps for the recurring partition maintenance tasks are:

References: https://docs.microsoft.com/en-us/sql/relational-databases/tables/manage-retention-of-historical-data-in-system-versioned-temporal-tables

Question: 2

DRAG DROP

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in this series.

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 configure the Fact.Order table.

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.

Actions

Recreate the Fact.Order table on the partition scheme.
Execute an ALTER TABLE command to specify the partition function.
Create a partiotion scheme based on the partition function.
Execute an ALTER TABLE command to specify the partition scheme.
Recreate the Fact.Order table on the partition function.
Create a partition function.

Answer Area

Answer:

Answer Area

Create a partition function.

Create a partiotion scheme based on the partition function.

Execute an ALTER TABLE command to specify the partition function.

Explanation:
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.
Step 1: Create a partition function.
Using CREATE PARTITION FUNCTION is the first step in creating a partitioned table or index.
Step 2: Create a partition scheme based on the partition function.
To migrate SQL Server partition definitions to SQL Data Warehouse simply:
Step 3: Execute an ALTER TABLE command to specify the partition function.
References: https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-partition

Question: 3

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are a database administrator for an e-commerce company that runs an online store. The company has three databases as described in the following table.

Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. **DB2** contains a table named OnlineOrder that is partitioned in hourly increments. The LOCK_ESCALATION option is set to **AUTO**. The data flow contains 24 OLE DB destinations, one for each partition.
DB3	This database runs Master Data Services (MDS).

You plan to load at least one million rows of data each night from DB1 into the OnlineOrder table.
You must load data into the correct partitions using a parallel process.
You create 24 Data Flow tasks. You must place the tasks into a component to allow parallel load. After all of the load processes compete, the process must proceed to the next task.
You need to load the data for the OnlineOrder table.
What should you use?

A. Lookup transformation
B. Merge transformation

C. Merge Join transformation
D. MERGE statement
E. Union All transformation
F. Balanced Data Distributor transformation
G. Sequential container
H. Foreach Loop container

<div align="right">

Answer: H

</div>

Explanation:
The Parallel Loop Task is an SSIS Control Flow task, which can execute multiple iterations of the standard Foreach Loop Container concurrently.
References: http://www.cozyroc.com/ssis/parallel-loop-task

Question: 4

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.
You are a database administrator for an e-commerce company that runs an online store. The company has the databases described in the following table.

Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. **DB2** contains a table named OnlineOrder that is partitioned in hourly increments. The LOCK_ESCALATION option is set to **AUTO**. The data flow contains 24 OLE DB destinations, one for each partition.
DB3	This database runs Master Data Services (MDS).

Each day, you publish a Microsoft Excel workbook that contains a list of product names and current prices to an external website. Suppliers update pricing information in the workbook. Each supplier saves the workbook with a unique name.
Each night, the Products table is deleted and refreshed from MDS by using a Microsoft SQL Server Integration Services (SSIS) package. All files must be loaded in sequence.
You need to add a data flow in an SSIS package to perform the Excel files import in the data warehouse.
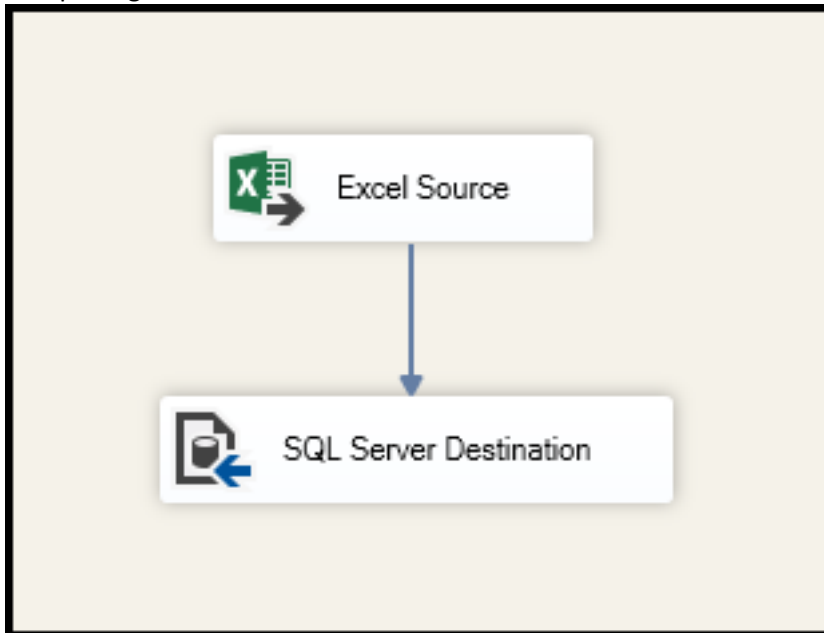What should you use?

A. Lookup transformation
B. Merge transformation
C. Merge Join transformation
D. MERGE statement
E. Union All transformation
F. Balanced Data Distributor transformation
G. Sequential container
H. Foreach Loop container

<div align="right">

Answer: A

</div>

Explanation:
If you're familiar with SSIS and don't want to run the SQL Server Import and Export Wizard, create an SSIS package that uses the Excel Source and the SQL Server Destination in the data flow.

References: https://docs.microsoft.com/en-us/sql/integration-services/import-export-data/import-data-from-excel-to-sql

Question: 5

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are a database administrator for an e-commerce company that runs an online store. The company has the databases described in the following table.

Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. **DB2** contains a table named OnlineOrder that is partitioned in hourly increments. The LOCK_ESCALATION option is set to **AUTO**. The data flow contains 24 OLE DB destinations, one for each partition.
DB3	This database runs Master Data Services (MDS).

Each week, you import a product catalog from a partner company to a staging table in DB2.
You need to create a stored procedure that will update the staging table by inserting new products and deleting discontinued products.
What should you use?

A. Lookup transformation
B. Merge transformation
C. Merge Join transformation
D. MERGE statement
E. Union All transformation

F. Balanced Data Distributor transformation
G. Sequential container
H. Foreach Loop container

Answer: G

Question: 6

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are a database administrator for an e-commerce company that runs an online store. The company has the databases described in the following table.

Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. **DB2** contains a table named OnlineOrder that is partitioned in hourly increments. The LOCK_ESCALATION option is set to **AUTO**. The data flow contains 24 OLE DB destinations, one for each partition.
DB3	This database runs Master Data Services (MDS).

Each day, data from the table OnlineOrder in DB2 must be exported by partition. The tables must not be locked during the process.
You need to write a Microsoft SQL Server Integration Services (SSIS) package that performs the data export.
What should you use?

A. Lookup transformation
B. Merge transformation
C. Merge Join transformation
D. MERGE statement
E. Union All transformation
F. Balanced Data Distributor transformation
G. Sequential container
H. Foreach Loop container

Answer: E

Explanation:
The Union All transformation combines multiple inputs into one output. For example, the outputs from five different Flat File sources can be inputs to the Union All transformation and combined into one output.
References: https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/union-all-transformation

Question: 7

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.

You are a database administrator for an e-commerce company that runs an online store. The company has the databases described in the following table.

Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. **DB2** contains a table named OnlineOrder that is partitioned in hourly increments. The LOCK_ESCALATION option is set to **AUTO**. The data flow contains 24 OLE DB destinations, one for each partition.
DB3	This database runs Master Data Services (MDS).

Product prices are updated and are stored in a table named Products on DB1. The Products table is deleted and refreshed each night from MDS by using a Microsoft SQL Server Integration Services (SSIS) package. None of the data sources are sorted.

You need to update the SSIS package to add current prices to the Products table.

What should you use?

A. Lookup transformation
B. Merge transformation
C. Merge Join transformation
D. MERGE statement
E. Union All transformation
F. Balanced Data Distributor transformation
G. Sequential container
H. Foreach Loop container

Answer: D

Explanation:
In the current release of SQL Server Integration Services, the SQL statement in an Execute SQL task can contain a MERGE statement. This MERGE statement enables you to accomplish multiple INSERT, UPDATE, and DELETE operations in a single statement.
References: https://docs.microsoft.com/en-us/sql/integration-services/control-flow/merge-in-integration-services-packages
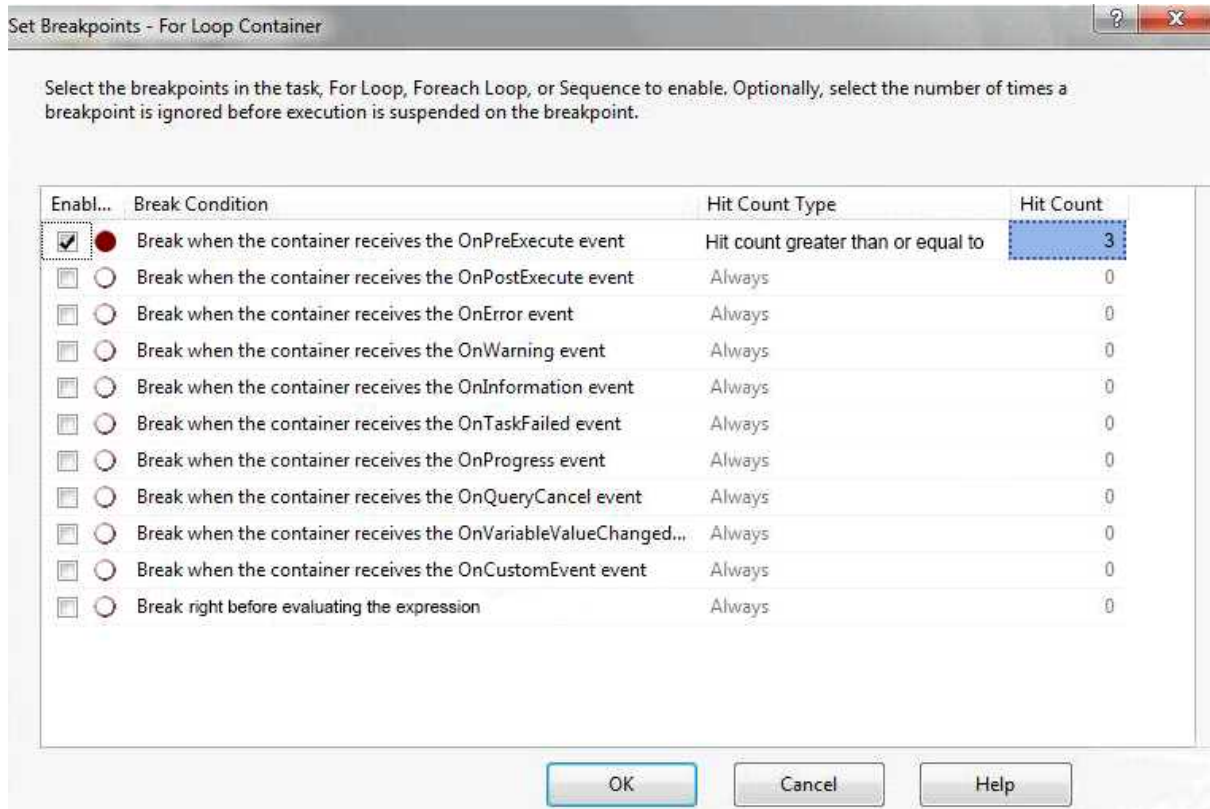
Question: 8

HOTSPOT
You have the Microsoft SQL Server Integration Services (SSIS) package shown in the Control flow exhibit. (Click the Exhibit button.)
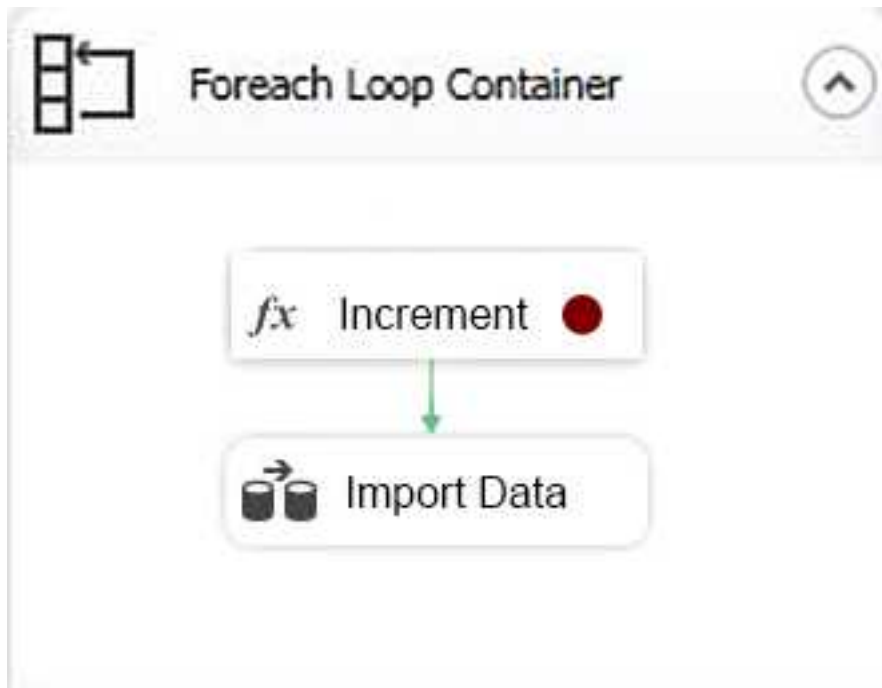
The package iterates over 100 files in a local folder. For each iteration, the package increments a variable named loop as shown in the Expression task exhibit. (Click the Exhibit button) and then imports a file. The initial value of the variable loop is 0.
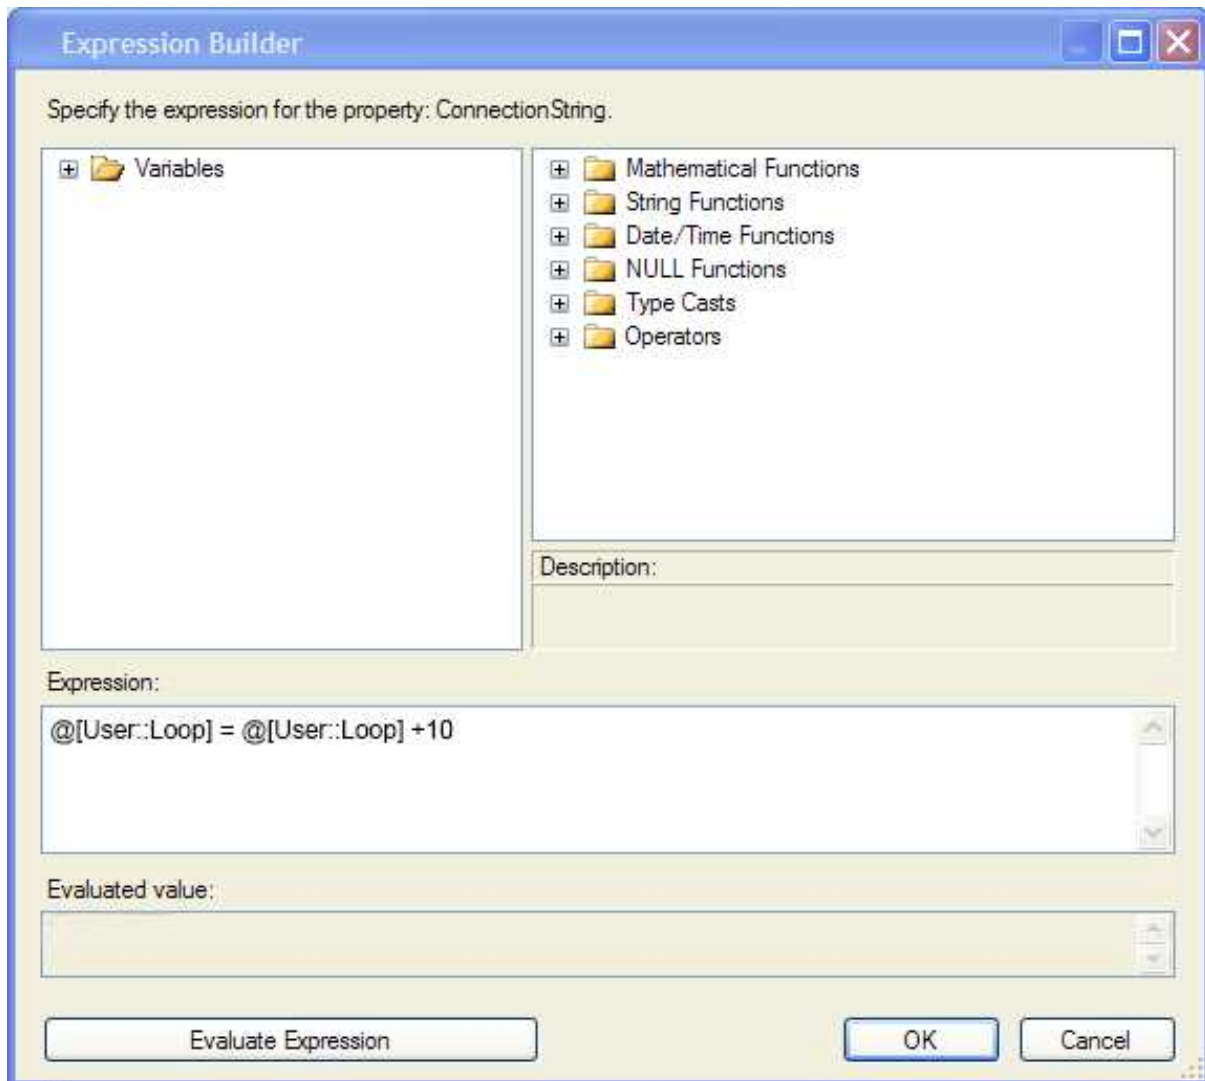
You suspect that there may be an issue with the variable value during the loop. You define a breakpoint on the Expression task as shown in the BreakPoint exhibit. (Click the Exhibit button.)

Set Breakpoints - For Loop Container

Select the breakpoints in the task, For Loop, Foreach Loop, or Sequence to enable. Optionally, select the number of times a breakpoint is ignored before execution is suspended on the breakpoint.

Enabl...		Break Condition	Hit Count Type	Hit Count
☑	●	Break when the container receives the OnPreExecute event	Hit count greater than or equal to	3
☐	○	Break when the container receives the OnPostExecute event	Always	0
☐	○	Break when the container receives the OnError event	Always	0
☐	○	Break when the container receives the OnWarning event	Always	0
☐	○	Break when the container receives the OnInformation event	Always	0
☐	○	Break when the container receives the OnTaskFailed event	Always	0
☐	○	Break when the container receives the OnProgress event	Always	0
☐	○	Break when the container receives the OnQueryCancel event	Always	0
☐	○	Break when the container receives the OnVariableValueChanged...	Always	0
☐	○	Break when the container receives the OnCustomEvent event	Always	0
☐	○	Break right before evaluating the expression	Always	0

OK   Cancel   Help

You need to check the value of the loop variable value.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Answer Area

	Yes	No
The value of the loop variable is **20** after the breakpoint is reached for the first time.	○	○
The loop variable resets to **0** when the breakpoint is reached.	○	○
When the code stops at a breakpoint, you can change the value of the loop variable.	○	○

Answer:

Answer Area

	Yes	No
The value of the loop variable is **20** after the breakpoint is reached for the first time.	◉	○
The loop variable resets to **0** when the breakpoint is reached.	○	◉
When the code stops at a breakpoint, you can change the value of the loop variable.	◉	○

Explanation:
Break condition: When the task or container receives the OnPreExecute event.
Called when a task is about to execute. This event is raised by a task or a container immediately before it runs.
The loop variable does not reset.
With the debugger, you can break, or suspend, execution of your program to examine your code, evaluate and edit variables in your program, etc.

Question: 9

You have a Microsoft SQL Server Integration Services (SSIS) package that includes the control flow shown in the following diagram.
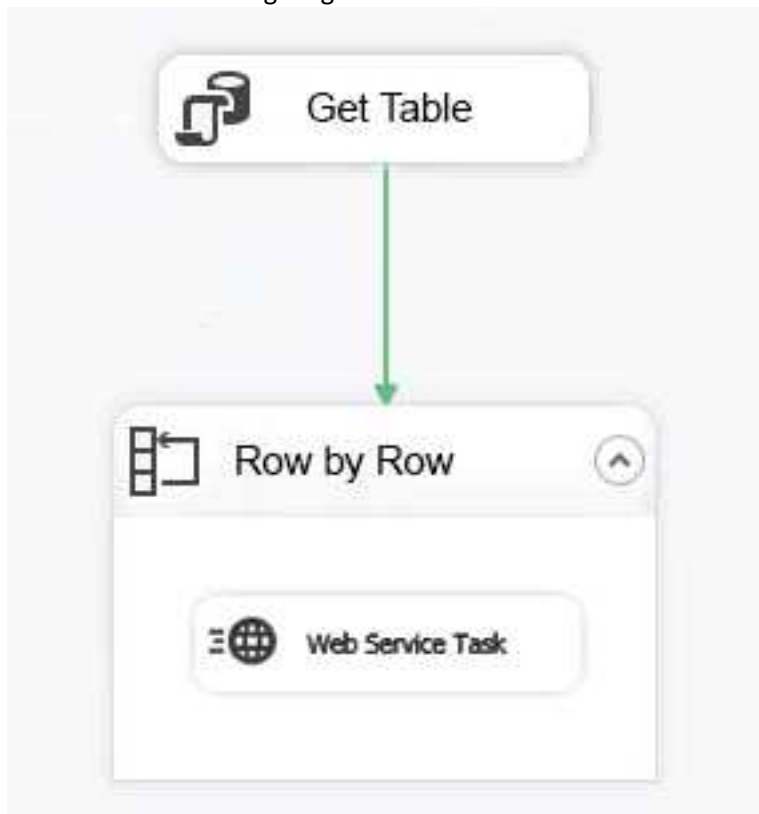
You need to choose the enumerator for the Foreach Loop container.
Which enumerator should you use?

A. Foreach SMO Enumerator
B. Foreach Azure Blob Enumerator
C. Foreach NodeList Enumerator
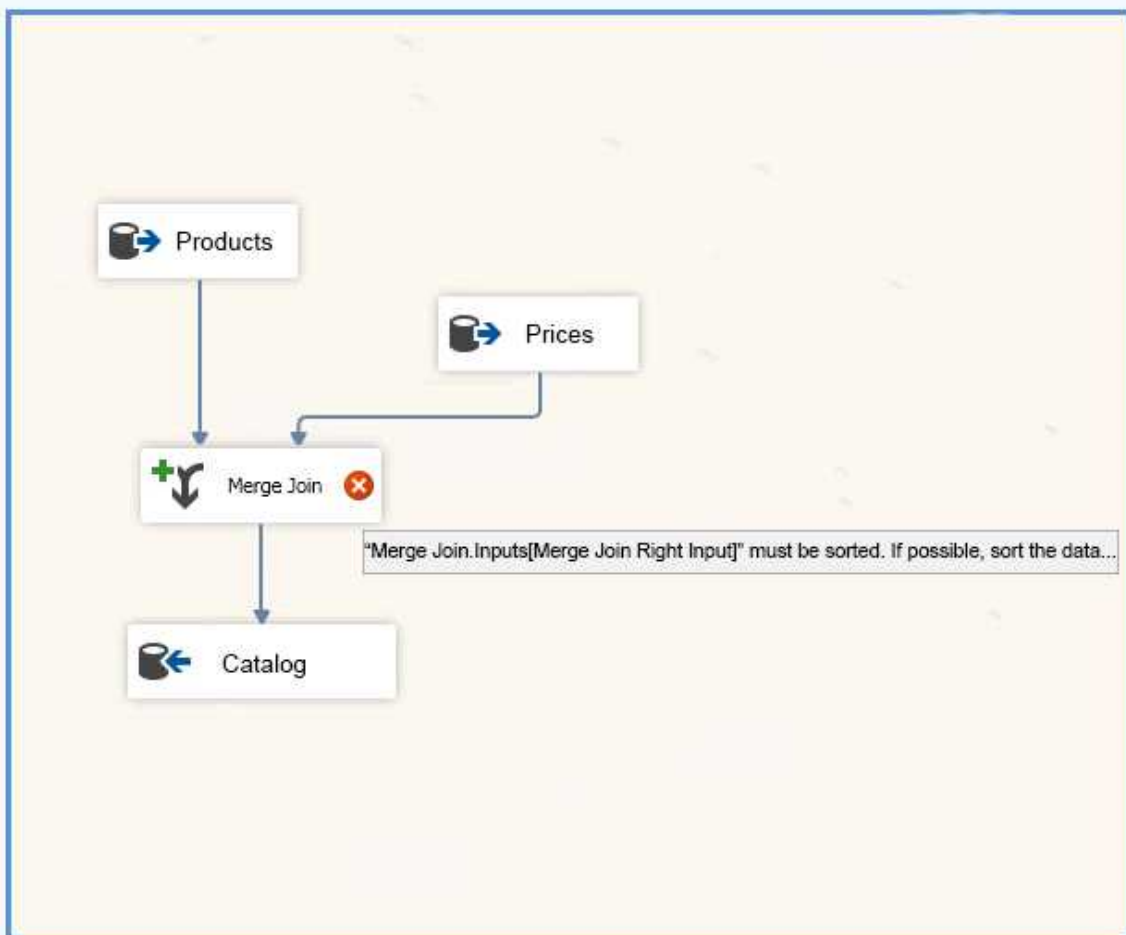D. Foreach ADO Enumerator

Answer: D

Explanation:
Use the Foreach ADO enumerator to enumerate rows in tables. For example, you can get the rows in an ADO recordset.

Question: 10

HOTSPOT
You create a Microsoft SQL Server Integration Services (SSIS) package as shown in the SSIS Package exhibit. (Click the Exhibit button.)

The package uses data from the Products table and the Prices table. Properties of the Prices source are shown in the OLE DB Source Editor exhibit (Click the Exhibit Button.) and the Advanced Editor for Prices exhibit (Click the Exhibit button.)
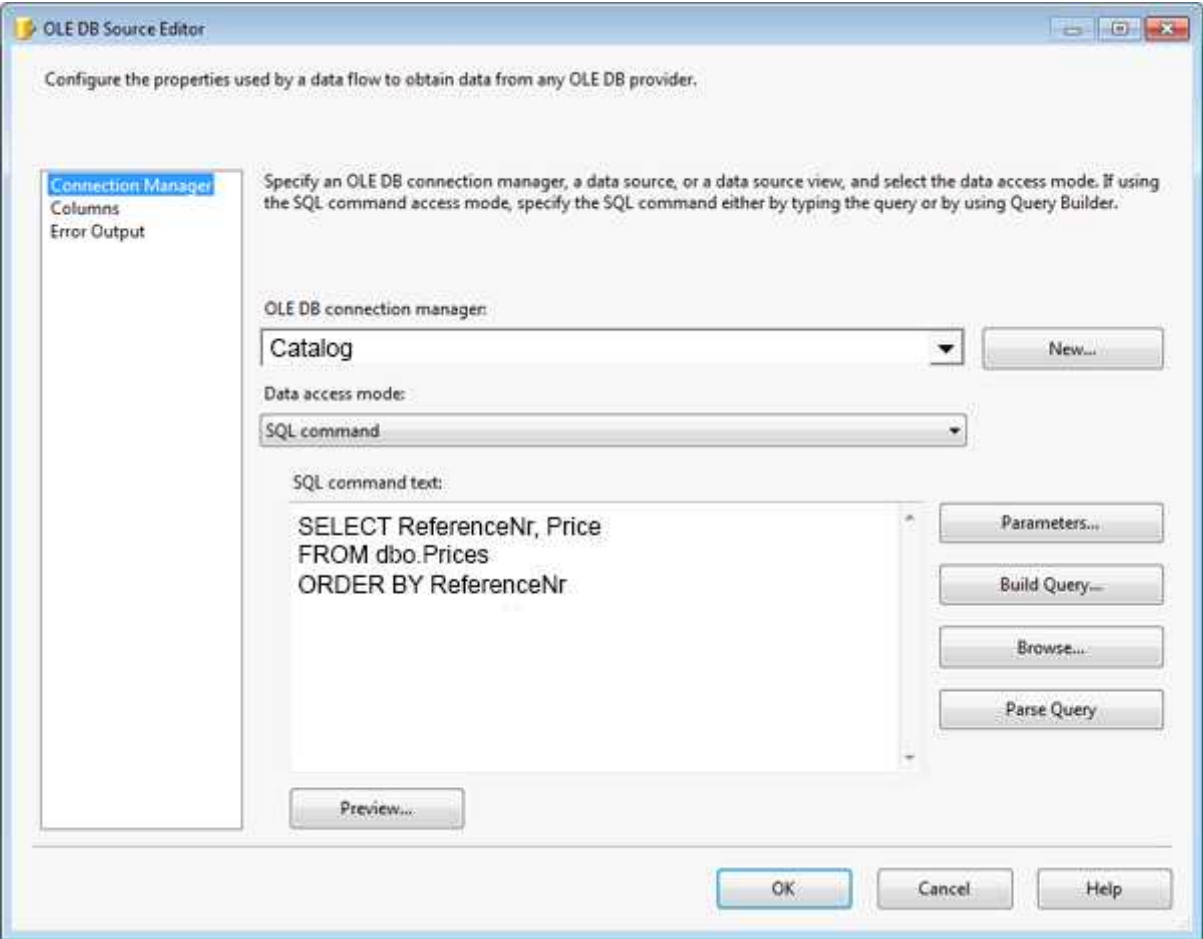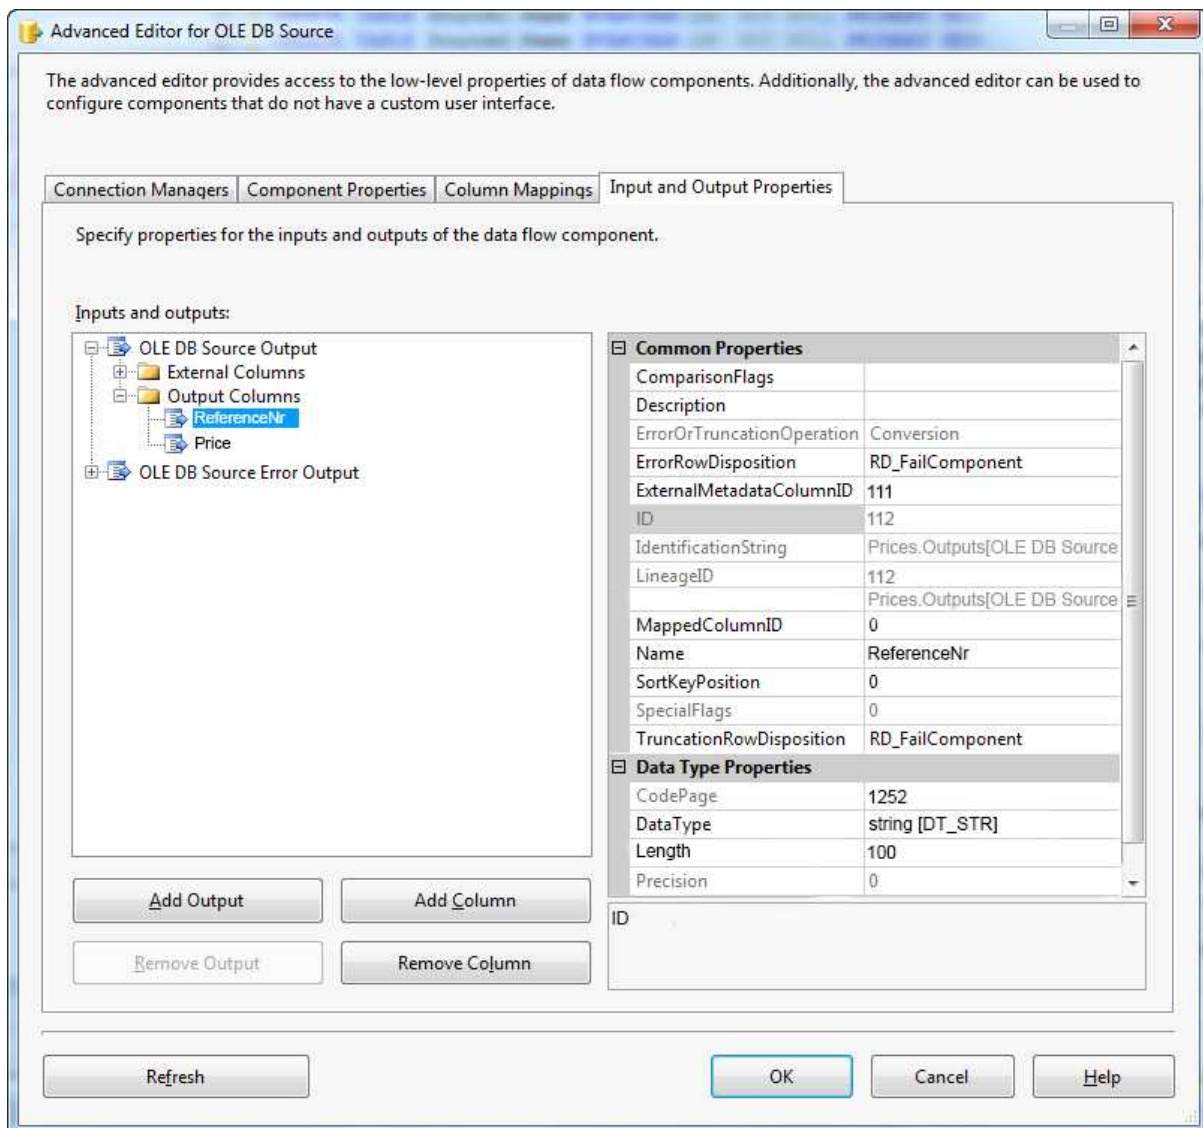
You join the Products and Prices tables by using the ReferenceNr column.

You need to resolve the error with the package.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

	Yes	No
You can resolve the error by adding a Sort transform between the OLE DB source and the Merge Join transform.	○	○
You can resolve the error by changing the SortKeyPosition setting for the ReferenceNr column and the value of the IsSorted setting for the OLE DB Source Output.	○	○
You can resolve the error by adding an Aggregate transform between the OLE DB source and the Merge Join transform.	○	○
You can resolve the error by replacing the Merge Join transform with a Lookup transform.	○	○

Answer:

Answer Area

	Yes	No
You can resolve the error by adding a Sort transform between the OLE DB source and the Merge Join transform.	**⊙**	○
You can resolve the error by changing the SortKeyPosition setting for the ReferenceNr column and the value of the IsSorted setting for the OLE DB Source Output.	**⊙**	○
You can resolve the error by adding an Aggregate transform between the OLE DB source and the Merge Join transform.	○	**⊙**
You can resolve the error by replacing the Merge Join transform with a Lookup transform.	○	**⊙**

Explanation:
There are two important sort properties that must be set for the source or upstream transformation that supplies data to the Merge and Merge Join transformations:
The Merge Join Transformation requires sorted data for its inputs.
If you do not use a Sort transformation to sort the data, you must set these sort properties manually on the source or the upstream transformation.
References: https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/sort-data-for-the-merge-and-merge-join-transformations