

Exam Questions 70-767

Implementing a SQL Data Warehouse (beta)

<https://www.2passeasy.com/dumps/70-767/>



NEW QUESTION 1

You deploy a Microsoft Azure SQL Data Warehouse instance. The instance must be available eight hours each day. You need to pause Azure resources when they are not in use to reduce costs. What will be the impact of pausing resources? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

What will happen to existing queries that are running?

<input type="checkbox"/> The data warehouse instance pauses when existing queries have completed. No new queries are permitted.
<input type="checkbox"/> The existing queries will be immediately terminated.
<input type="checkbox"/> The existing queries will be paused until the data warehouse instance is resumed.

What will happen to the charges for the data warehouse instance?

<input type="checkbox"/> You will stop being charged for compute resources but will continue to be charged for storage.
<input type="checkbox"/> You will continue to be charged for both compute resources and storage.
<input type="checkbox"/> You are no longer charged for storage but continue to pay for the assigned data warehouse instance units.

Answer:

Explanation: To save costs, you can pause and resume compute resources on-demand. For example, if you won't be using the database during the night and on weekends, you can pause it during those times, and resume it during the day. You won't be charged for DWUs while the database is paused.

When you pause a database:

Compute and memory resources are returned to the pool of available resources in the data center Data Warehouse Unit (DWU) costs are zero for the duration of the pause.

Data storage is not affected and your data stays intact.

SQL Data Warehouse cancels all running or queued operations. When you resume a database:

SQL Data Warehouse acquires compute and memory resources for your DWU setting. Compute charges for your DWUs resume.

Your data will be available.

You will need to restart your workload queries. References:

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

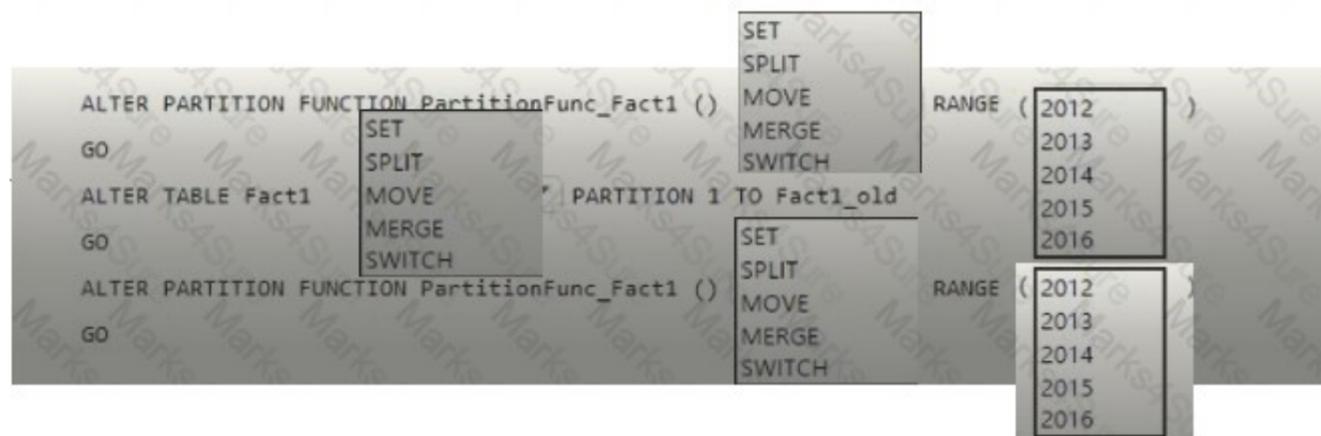
NEW QUESTION 2

You are the administrator of a database that hosts tables for a data warehouse. The table named Fact1 has data from the start of calendar year 2011 through the end of 2017. The table contains at least 20 million rows of data for each year.

You create the table by running the following Transact-SQL statement:

```
CREATE PARTITION FUNCTION PartitionFunc_Fact1(SMALLINT) AS RANGE LEFT FOR VALUES(2012,2013,2014,2015)
```

You need to modify the partition function so that rows for each calendar year are in a separate partition. You must also move all data prior to 2014 to another table named Fact1.old.



Answer:

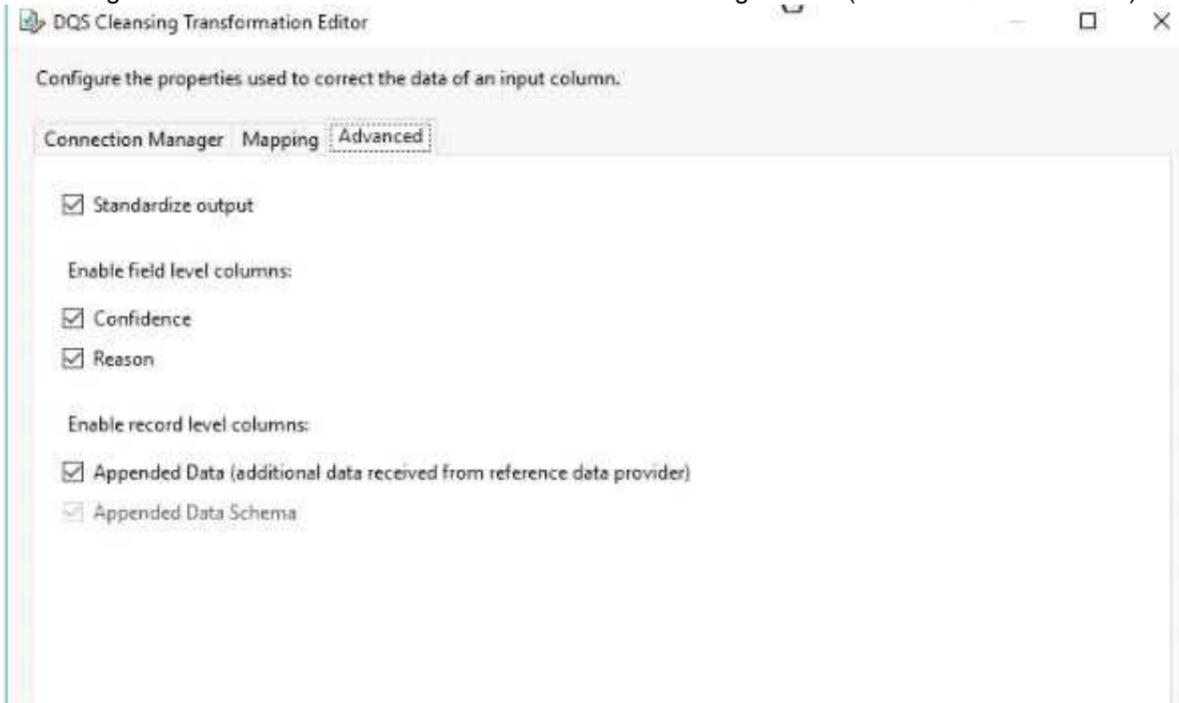
Explanation: SET 2102
Merge Split 2016

NEW QUESTION 3

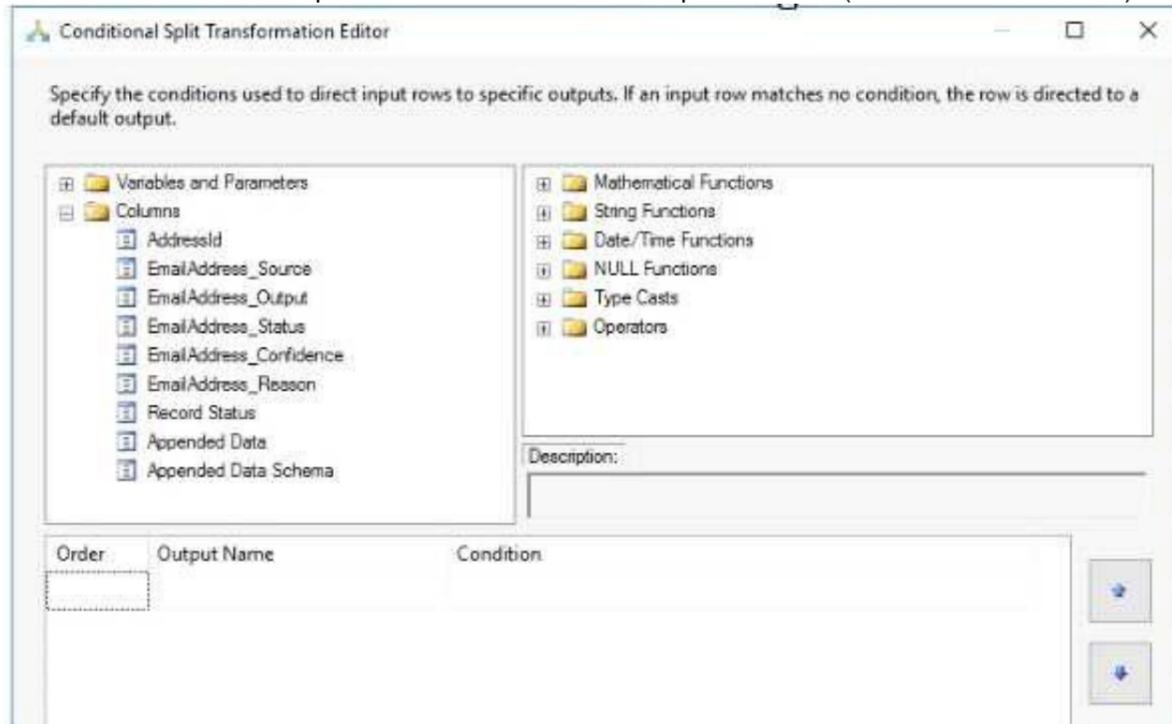
You have a Microsoft SQL Server Integration Services (SSIS) package that contains a Data Flow task as shown in the Data Flow exhibit. (Click the Exhibit button.)



You install Data Quality Services (DQS) on the same server that hosts SSIS and deploy a knowledge base to manage customer email addresses. You add a DQS Cleansing transform to the Data Flow as shown in the Cleansing exhibit. (Click the Exhibit button.)



You create a Conditional Split transform as shown in the Splitter exhibit. (Click the Exhibit button.)



You need to split the output of the DQ5 Cleansing task to obtain only Correct values from the EmailAddress column. For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer area

- | | Yes | No |
|---|-----------------------|-----------------------|
| You can use the EmailAddress_Output column to split the output. | <input type="radio"/> | <input type="radio"/> |
| You can use the EmailAddress_Status column to split the output. | <input type="radio"/> | <input type="radio"/> |

Answer:

Explanation:

You can use the EmailAddress_Output column to split the output.

Yes No

You can use the EmailAddress_Status column to split the output.

NEW QUESTION 4

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You configure a new matching policy Master Data Services (MDS) as shown in the following exhibit.

Rule Details		Rule Editor			
Rule name:	Customer attributes	Domain	Similarity	Weight	Prerequisite
Description:		Gender	Exact		<input checked="" type="checkbox"/>
Min. matching score:	80 %	City	Exact		<input checked="" type="checkbox"/>
		State	Exact		<input checked="" type="checkbox"/>
		Country	Exact		<input checked="" type="checkbox"/>
		Zip	Exact		<input checked="" type="checkbox"/>
		Birth Date	Similar D 3 M 0 Y 0	34 %	<input type="checkbox"/>
		Address Line 1	Similar	10 %	<input type="checkbox"/>
		First Name	Similar	33 %	<input type="checkbox"/>
		Last Name	Similar	23 %	<input type="checkbox"/>

You review the Matching Results of the policy and find that the number of new values matches the new values.

You verify that the data contains multiple records that have similar address values, and you expect some of the records to match.

You need to increase the likelihood that the records will match when they have similar address values. Solution: You increase the relative weights for Address Line 1 of the matching policy.

Does this meet the goal?

- A. Yes
- B. NO

Answer: B

Explanation: Decrease the Min. matching score.

A data matching project consists of a computer-assisted process and an interactive process. The matching project applies the matching rules in the matching policy to the data source to be assessed. This process assesses the likelihood that any two rows are matches in a matching score. Only those records with a probability of a match greater than a value set by the data steward in the matching policy will be considered a match.

References: <https://docs.microsoft.com/en-us/sql/data-quality-services/data-matching>

NEW QUESTION 5

You are designing an indexing strategy for a data warehouse. The data warehouse contains a table named Table1. Data is bulk inserted into Table1.

You plan to create the indexes configured as shown in the following table.

Index name	Indexing specifications
Index1	<ul style="list-style-type: none"> • Index1 contains all the data in Table1. • Queries against Index1 perform aggregation operations against hundreds of millions of rows.
Index2	<ul style="list-style-type: none"> • Index2 returns all the columns in this index. • Index2 contains 80 percent of the columns in Table1. • Index2 is used to assist with queries against other tables by performing point lookups against Table1.

Which type of index should you use to minimize the query times of each index? To answer, drag the appropriate index types to the correct indexes. Each index type 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.

Index Types

Clustered	Clustered columnstore
Hash	Heap
Nonclustered	Nonclustered columnstore

Answer Area

Index1:

Index2:

Answer:

Explanation: Index Types

Clustered	Clustered columnstore
Hash	Heap
Nonclustered	Nonclustered columnstore

Answer Area

Index1:

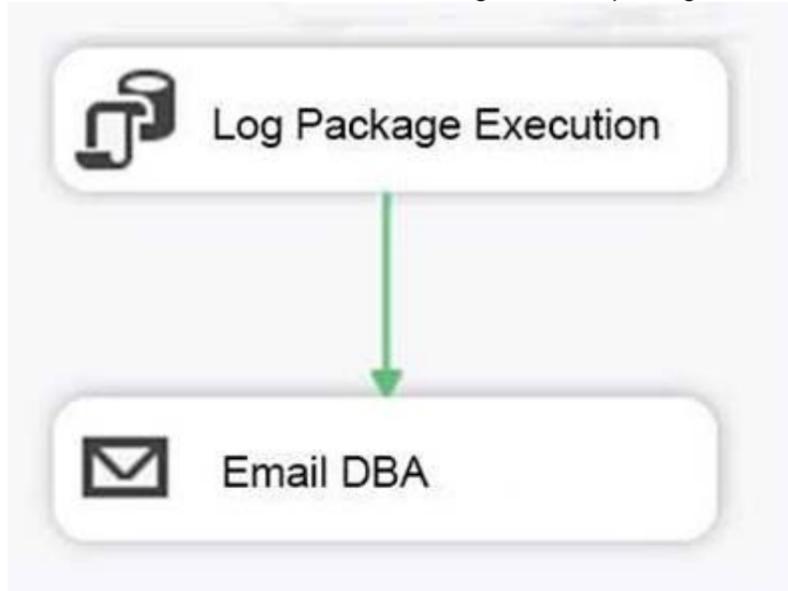
Index2:

NEW QUESTION 6

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are developing a Microsoft SQL Server Integration Services (SSIS) projects. The project consists of several packages that load data warehouse tables. You need to extend the control flow design for each package to use the following control flow while minimizing development efforts and maintenance:



Solution: You add the control flow to a script task. You add an instance of the script task to the storage account in Microsoft Azure. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: A package consists of a control flow and, optionally, one or more data flows. You create the control flow in a package by using the Control Flow tab in SSIS Designer.

References: <https://docs.microsoft.com/en-us/sql/integration-services/control-flow/control-flow>

NEW QUESTION 7

Your company has a Microsoft SQL Server data warehouse instance. The human resources department assigns all employees a unique identifier. You plan to store this identifier in a new table named Employee.

You create a new dimension to store information about employees by running the following Transact-SQL statement:

```
CREATE TABLE [Dimension].[Employee]
(
    [EmployeeID] [int] NOT NULL,
    [EmployeeName] [nvarchar](50) NULL,
    [PreferredName] [nvarchar](50) NULL,
    [IsSalesperson] [bit] NOT NULL,
    [Email] [nvarchar](50) NULL
)
```

You have not added data to the dimension yet. You need to modify the dimension to implement a new column named [EmployeeKey]. The new column must use unique values.

How should you complete the Transact-SQL statements? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

```
ALTER TABLE [Dimension].[Employee]
```

ADD [EmployeeKey] INT IDENTITY(1,1) NULL ADD [EmployeeKey] INT IDENTITY(1,1) NOT NULL ADD [EmployeeID] INT IDENTITY(1,1) NULL ADD [EmployeeID] INT IDENTITY(1,1) NOT NULL
--

```
ALTER TABLE [Dimension].[Employee]
ADD CONSTRAINT PK_Dimension_Employee
```

PRIMARY KEY CLUSTERED ([EmployeeKey]) PRIMARY KEY CLUSTERED ([EmployeeID]) PRIMARY KEY CLUSTERED ([Employee]) PRIMARY KEY CLUSTERED ([PreferredName])
--

Answer:

Explanation:

Answer Area

```
ALTER TABLE [Dimension].[Employee]
```

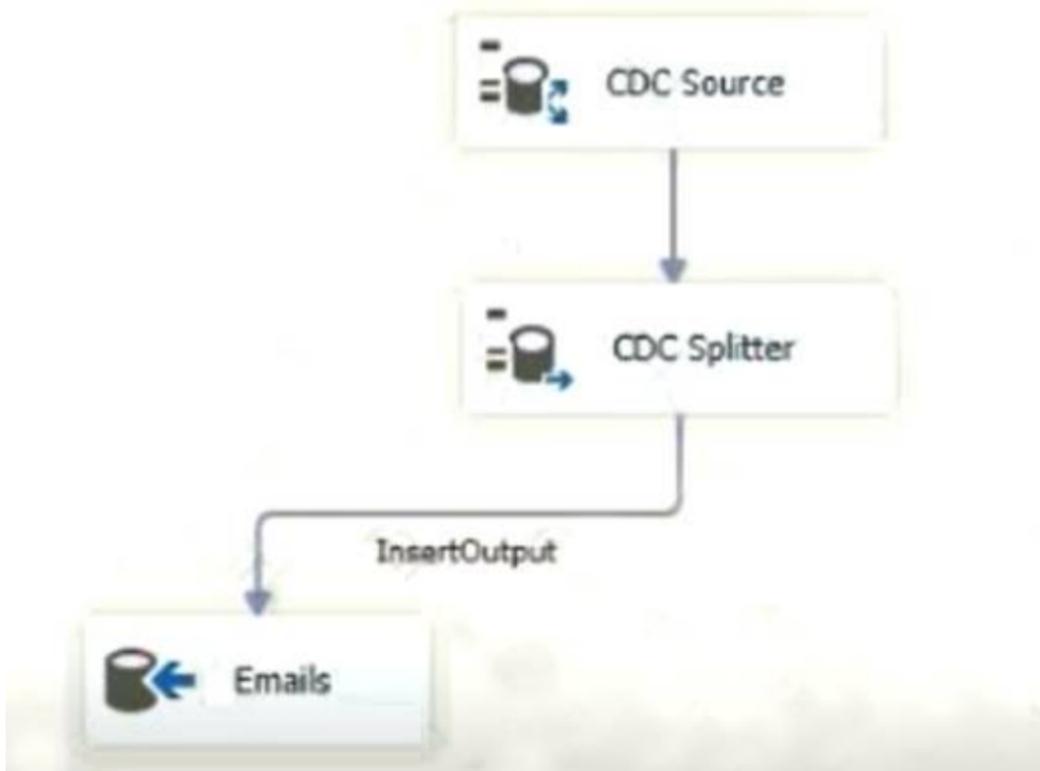
ADD [EmployeeKey] INT IDENTITY(1,1) NULL ADD [EmployeeKey] INT IDENTITY(1,1) NOT NULL ADD [EmployeeID] INT IDENTITY(1,1) NULL ADD [EmployeeID] INT IDENTITY(1,1) NOT NULL

```
ALTER TABLE [Dimension].[Employee]
ADD CONSTRAINT PK_Dimension_Employee
```

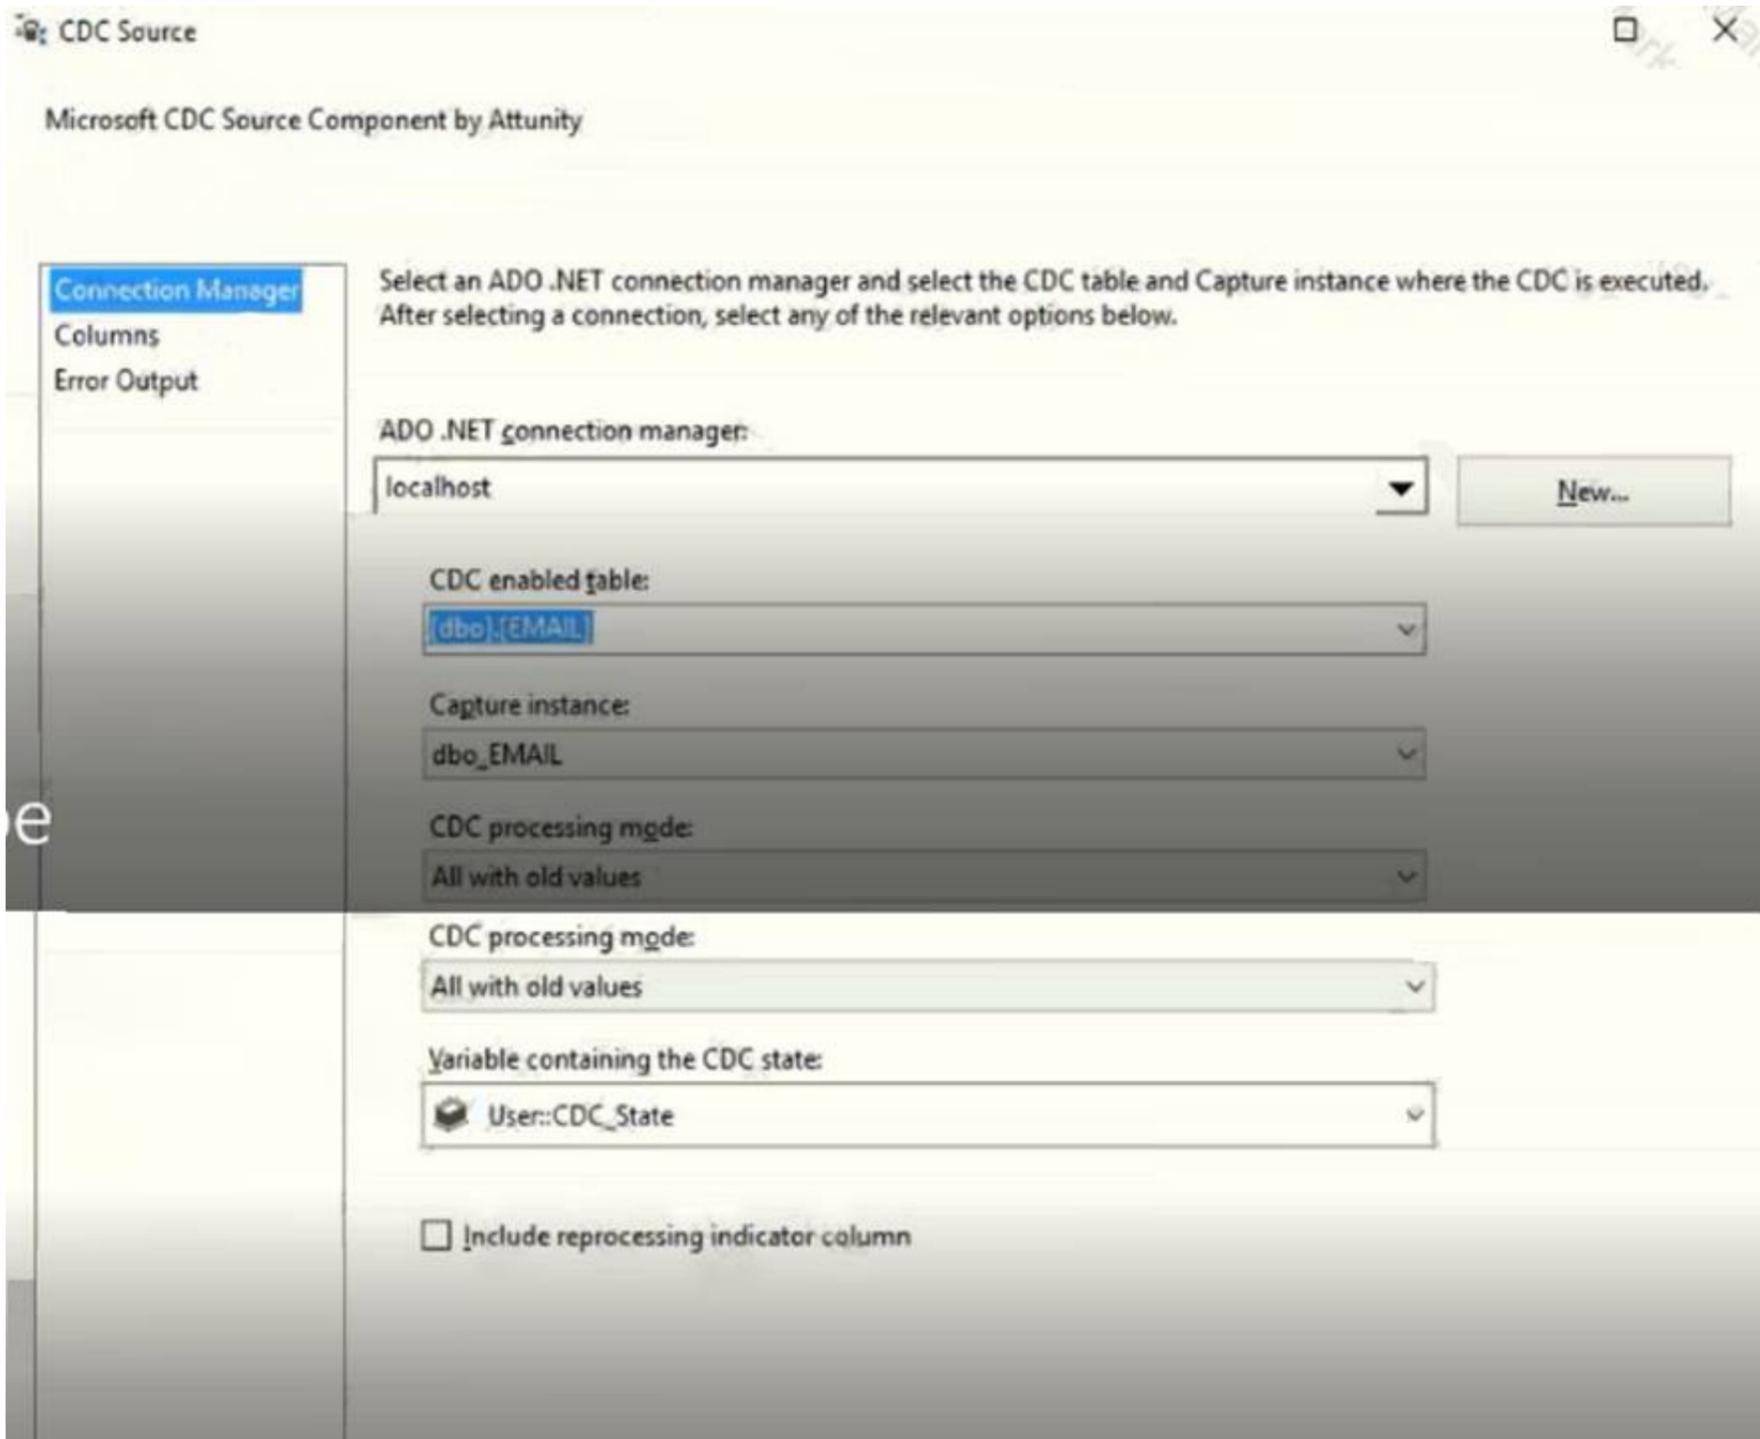
PRIMARY KEY CLUSTERED ([EmployeeKey]) PRIMARY KEY CLUSTERED ([EmployeeID]) PRIMARY KEY CLUSTERED ([Employee]) PRIMARY KEY CLUSTERED ([PreferredName])

NEW QUESTION 8

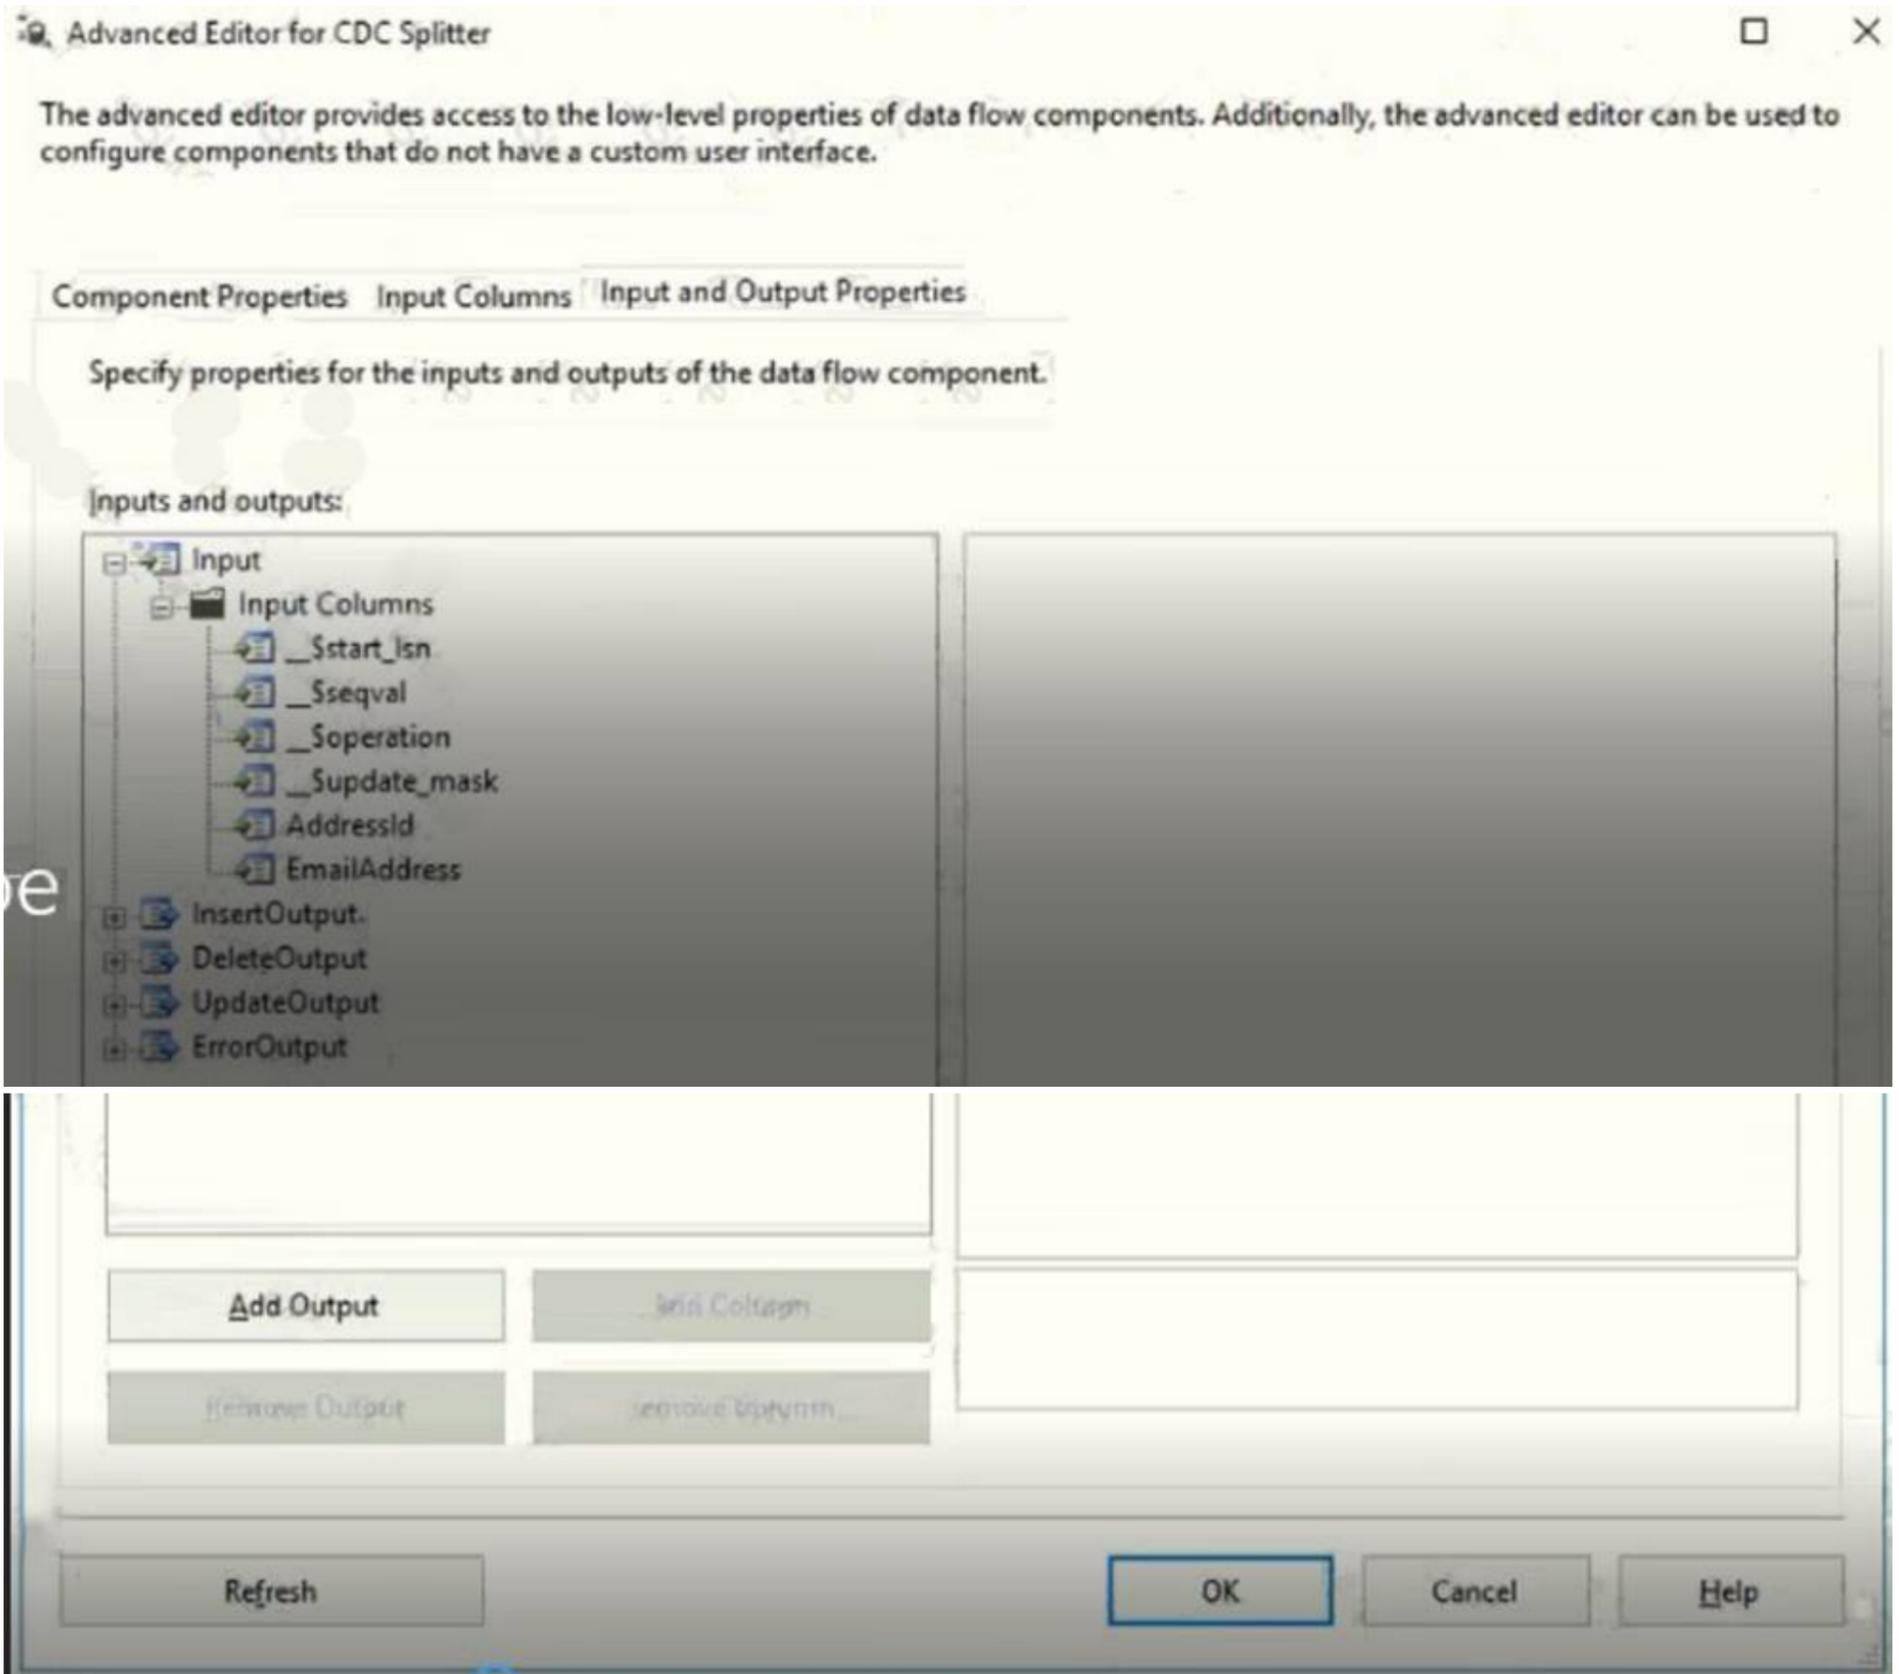
You have a database that contains a table named Email. Change Data Capture (CDC) is enabled for the table. You have a Microsoft SQL Server Integration Services (SSIS) package that contains the Data Flow task shown in the Data Flow exhibit. (Click the Exhibit button.)



You have an existing CDC source as shown in the CDC Source exhibit (Click the Exhibit button)



and a CDC Splitter transform as shown in the CDC Splitter exhibit. (Click the Exhibit button.)



You need to perform an incremental import of customer email addresses. Before importing email addresses, you must move all previous email addresses to another table for later use.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

	Yes	No
You can use the Delete Output data flow and the <code>\$_operation</code> input column for the Splitter task.	<input type="radio"/>	<input type="radio"/>
You can use the Update Output data flow and the <code>\$_operation</code> input column of the Splitter task.	<input type="radio"/>	<input type="radio"/>
You can use the Update Output data flow and the <code>\$_update_mask</code> input column of the Splitter task.	<input type="radio"/>	<input type="radio"/>
You can use the Delete Output data flow and the <code>\$_update_mask</code> input column of the Splitter task.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation: Yes
 Yes Yes No

NEW QUESTION 9

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:

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.

References:

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

NEW QUESTION 10

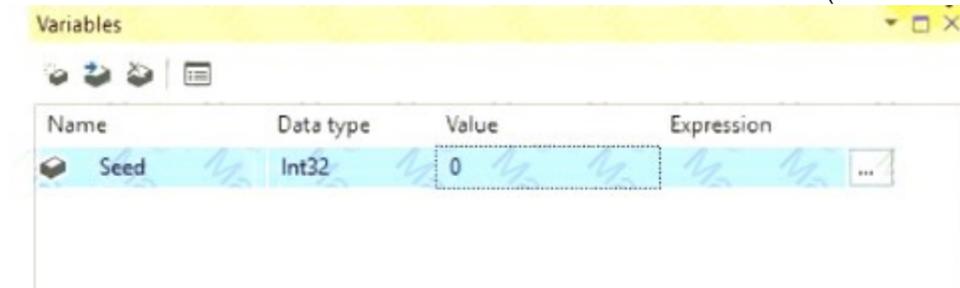
You are testing a Microsoft SQL Server Integration Services (SSIS) package. The package includes the Control Flow task shown in the Control Flow exhibit (Click the Exhibit button)



and the Data Flow task shown in the Data Flow exhibit. (Click the Exhibit button.)



You declare a variable named seed as shown in the Variables exhibit. (Click the Exhibit button.).



The variable is changed by the Script task during execution.

You need to be able to interrogate the value of the seed variable after the Script task completes execution. For each of the following statements, select Yes if the statement is true. Otherwise, select No.

	Yes	No
You can display the variable by adding a data viewer to the data flow.	<input type="radio"/>	<input type="radio"/>
You can display the variable by adding a breakpoint to the OnPostExecute event and using the Locals window.	<input type="radio"/>	<input type="radio"/>
You can display the variable by adding a breakpoint to the OnVariableValueChanged event and using the Watch window.	<input type="radio"/>	<input type="radio"/>
You can display the variable by adding the following code segment to the Script task: <code>MessageBox.Show</code>	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation: No Yes No Yes

NEW QUESTION 10

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. 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
 References:
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-compute-overview>

NEW QUESTION 13

You have a database named DB1. You create a Microsoft SQL Server Integration Services (SSIS) package that incrementally imports data from a table named Customers. The package uses an OLE DB data source for connections to DB1. The package defines the following variables.

Variable name	Data type	Description
LastKey	Int64	LastKey stores the last identifier used in the imported table.
TableName	String	TableName stores the name of the imported table.

To support incremental data loading, you create a table by running the following Transact-SQL segment:

```
CREATE TABLE LastKeyByTable (
    Id int IDENTITY(1,1) PRIMARY KEY,
    TableName sysname UNIQUE,
    LastKey bigint
)
```

You need to create a DML statements that updates the LastKeyByTable table.

How should you complete the Transact-SQL statement? To answer, select the appropriate Transact-SQL segments in the dialog box in the answer area.

Answer Area

```
UPDATE dbo.LastKeyByTable
SET
```

▼

LastKey = ?

LastKey = @A

LastKey = @B

LastKey = @LastKey

```
WHERE
```

▼

TableName = ?

TableName = @A

TableName = @B

TableName = @TableName

Answer:

Explanation:

Answer Area

UPDATE dbo.LastKeyByTable

SET

LastKey = ?	▼
LastKey = @A	
LastKey = @B	
LastKey = @LastKey	

WHERE

TableName = ?	▼
TableName = @A	
TableName = @B	
TableName = @TableName	

NEW QUESTION 18

You are building a server to host a data warehouse.

The planned disk activity for the data warehouse is five percent write activity and 95 percent read activity. You need to recommend a storage solution for the data files of the data warehouse. The solution must meet the following requirements:

*Ensure that the data warehouse is available if two disks fail.

*Minimize hardware costs.

Which RAID configuration should you recommend?

- A. RAID1
- B. RAID 5
- C. RAID 6
- D. RAID 10

Answer: C

Explanation: According to the Storage Networking Industry Association (SNIA), the definition of RAID 6 is: "Any form of RAID that can continue to execute read and write requests to all of a RAID array's virtual disks in the presence of any two concurrent disk failures."

NEW QUESTION 20

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 implementing a Microsoft SQL Server data warehouse with a multi-dimensional data model. You have a fact table that includes sales data for all products. The model includes a dimension named Geography that stores all geographies. You create a dimension that has a foreign key and provides the ability to analyze sales by the following sales channels: Internet or retail store.

You need to update the data model to allow business users to analyze Internet sales by geography without changing the overall structure of the data model. What should you do?

- A. star schema
- B. snowflake schema
- C. conformed dimension
- D. slowly changing dimension (SCD)
- E. fact table
- F. semi-additive measure
- G. non-additive measure
- H. dimension table reference relationship

Answer: D

NEW QUESTION 21

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a Microsoft SQL server that has Data Quality Services (DQS) installed.

You need to review the completeness and the uniqueness of the data stored in the matching policy. Solution: You create a matching rule.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: Use a matching rule, and use completeness and uniqueness data to determine what weight to give a field in the matching process. If there is a high level of uniqueness in a field, using the field in a matching policy can decrease the matching results, so you may want to set the weight for that field to a relatively small value. If you have a low level of uniqueness for a column, but low completeness, you may not want to include a domain for that column.
 References:
<https://docs.microsoft.com/en-us/sql/data-quality-services/create-a-matching-policy?view=sql-server-2017>

NEW QUESTION 25

You have a data warehouse.
 You need to move a table named Fact.ErrorLog to a new filegroup named LowCost.
 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

- Add a file to the LowCost filegroup.
- Rename the Fact.ErrorLog table to Fact.ErrorLogBak.
- Drop the Fact.ErrorLog table.
- Create a new Fact.ErrorLog table on the LowCost filegroup.
- Add a filegroup named LowCost to the database.
- Reorganize the clustered index on the Fact.ErrorLog table in the new filegroup.
- Rebuild the clustered index on the Fact.ErrorLog table in the new filegroup.

Answer Area



Answer:

Explanation: Step 1: Add a filegroup named LowCost to the database. First create a new filegroup.
 Step 2:
 The next stage is to go to the 'Files' page in the same Properties window and add a file to the filegroup (a filegroup always contains one or more files)
 Step 3:
 To move a table to a different filegroup involves moving the table's clustered index to the new filegroup. While this may seem strange at first this is not that surprising when you remember that the leaf level of the clustered index actually contains the table data. Moving the clustered index can be done in a single statement using the DROP_EXISTING clause as follows (using one of the AdventureWorks2008R2 tables as an example) :
`CREATE UNIQUE CLUSTERED INDEX PK_Department_DepartmentID ON HumanResources.Department(DepartmentID) WITH (DROP_EXISTING=ON,ONLINE=ON) ON SECONDARY`
 This recreates the same index but on the SECONDARY filegroup.
 References:
<http://www.sqlmatters.com/Articles/Moving%20a%20Table%20to%20a%20Different%20Filegroup.aspx>

NEW QUESTION 28

You are implementing a Microsoft SQL Server data warehouse with a multi-dimensional data model. Orders are stored in a table named Factorder. The addresses that are associated with all orders are stored in a fact table named FactAddress. A key in the FoctAddress table specifies the type of address for an order. You need to ensure that business users can examine the address data by either of the following:
 • shipping address and billing address
 • shipping address or billing address type Which data model should you use?

- A. star schema
- B. snowflake schema
- C. conformed dimension
- D. slowly changing dimension (SCD)
- E. fact table

- F. semi-additive measure
- G. non-additive measure
- H. dimension table reference relationship

Answer: H

NEW QUESTION 31

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a data warehouse that stores information about products, sales, and orders for a manufacturing company. The instance contains a database that has two tables named SalesOrderHeader and SalesOrderDetail. SalesOrderHeader has 500,000 rows and SalesOrderDetail has 3,000,000 rows.

Users report performance degradation when they run the following stored procedure:

```
CREATE PROCEDURE Sales.GetRecentSales (@date datetime)
AS BEGIN
    IF @date is NULL
        SET @date = DATEADD(MONTH, -3, (SELECT MAX(ORDERDATE) FROM Sales.SalesOrderHeader))
    SELECT * FROM Sales.SalesOrderHeader h, Sales.SalesOrderDetail d
    WHERE h.SalesOrderID = d.SalesOrderID
    AND h.OrderDate > @date
END
```

You need to optimize performance.

Solution: You run the following Transact-SQL statement:

```
CREATE STATISTICS Stat1
ON Sales.SalesOrderHeader (OrderDate)
WITH FULLSCAN
```

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation: UPDATE STATISTICS updates query optimization statistics on a table or indexed view. FULLSCAN computes statistics by scanning all rows in the table or indexed view. FULLSCAN and SAMPLE 100 PERCENT have the same results.

References:

<https://docs.microsoft.com/en-us/sql/t-sql/statements/update-statistics-transact-sql?view=sql-server-2017>

NEW QUESTION 32

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.
- ▶ 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 on 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.

NEW QUESTION 35

You plan to use the dtutil.exe utility with Microsoft SQL Server Integration Services (SSIS) to customize packages. You need to create a new package ID for package1 on Server1. Which dtutil.exe command should you run?

- A. dtutil.exe /FILE c:\repository\package1.dtsx /DestServer Server! /COPY SQL;package1.dtsx
- B. dtutil.exe /I /FILE c:\repository\package1.dtsx
- C. dtutil.exe /SQL package1 /COPY OTS;c:\repository\package1.dtsx
- D. dtutil.exe /SQL package1 /DELETE

Answer: A

NEW QUESTION 37

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.
- ▶ 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:

Explanation: 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-v>

NEW QUESTION 38

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 to 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.
- ▶ 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 data loading for the tables.

Which data loading technology should you use for each table? To answer, select the appropriate options in the answer area.

Table	Technology
Dimension.SalesTerritory	<div style="border: 1px solid black; padding: 5px;"> <p>Change Data Capture (CDC)</p> <p>Change Tracking</p> <p>Temporal table</p> <p>Microsoft SQL Server snapshot replication</p> </div>
Dimension.Customer	<div style="border: 1px solid black; padding: 5px;"> <p>Change Data Capture (CDC)</p> <p>Change Tracking</p> <p>Temporal table</p> <p>Microsoft SQL Server snapshot replication</p> </div>
Dimension.Date	<div style="border: 1px solid black; padding: 5px;"> <p>Change Data Capture (CDC)</p> <p>Change Tracking</p> <p>Temporal table</p> <p>Microsoft SQL Server snapshot replication</p> </div>

Answer:

Explanation: Scenario: The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated
 Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables. Box 1: Change Tracking
 Box 2: Change Tracking Box 3: Temporal Table

Temporal Tables are generally useful in scenarios that require tracking history of data changes.

We recommend you to consider Temporal Tables in the following use cases for major productivity benefits.

* Slowly-Changing Dimensions

Dimensions in data warehousing typically contain relatively static data about entities such as geographical locations, customers, or products.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/tables/temporal-table-usage-scenarios>

NEW QUESTION 43

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a Microsoft Azure SQL Data Warehouse instance. You run the following Transact-SQL statement:

```
SELECT CustomerKey, SUM(SalesAmt) TotalSales
FROM sales.FactOrders
GROUP BY CustomerKey
```

The query fails to return results.

You need to determine why the query fails.

Solution: You run the following Transact-SQL statement:

```
SELECT TOP 1 status, total_elapsed_time, submit_time
FROM sales.FactOrders
WHERE [label] = 'TotalSales'
ORDER BY submit_time
```

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: To use submit_time we must use sys.dm_pdw_exec_requests table. References:
<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec>

NEW QUESTION 45

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have the following line-of-business solutions:

- ▶ 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.

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:

- ▶ 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.

Solution: Perform the following actions: 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.
References: <https://solutioncenter.apexsql.com/enable-use-sql-server-change-data-capture/>

NEW QUESTION 48

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Your company uses Microsoft SQL Server to deploy a data warehouse to an environment that has a SQL Server Analysis Services (SSAS) instance. The data warehouse includes the Fact.Order table as shown in the following table definition. The table has no indexes.

Columns

- Order Key (bigint, not null)
- City Key (int, not null)
- Customer Key (int, not null)
- Stock Item Key (int, not null)
- Order Date Key (date, not null)
- Picked Date Key (date, null)
- Salesperson Key (int, not null)
- Picker Key (int, null)
- Quantity (int, not null)
- Unit Price (decimal(18,2), not null)
- Tax Rate (decimal(18,3), not null)
- Total Excluding Tax (decimal(18,2), not null)
- Tax Amount (decimal(18,2), not null)
- Total Including Tax (decimal(18,2), not null)

```
SELECT AVG([Tax Amount]) AS [Average Tax Amount]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'

SELECT SUM([Total Excluding Tax]) AS [Total Revenue]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'
```

You need to ensure that the queries complete as quickly as possible.
 Solution: You create measure for the Fact.Order table. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: You should use a columnstore index. Columnstore indexes are the standard for storing and querying large data warehousing fact tables. This index uses column-based data storage and query processing to achieve gains up to 10 times the query performance in your data warehouse over traditional row-oriented storage.
 References:
<https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview?view=sql-serv>

NEW QUESTION 50

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You plan to deploy a Microsoft SQL server that will host a data warehouse named DB1. The server will contain four SATA drives configured as a RAID 10 array. You need to minimize write contention on the transaction log when data is being loaded to the database. Solution: You add more data files to DB1. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: There is no performance gain, in terms of log throughput, from multiple log files. SQL Server does not write log records in parallel to multiple log files. Instead you should place the log file on a separate drive. References:
<https://www.red-gate.com/simple-talk/sql/database-administration/optimizing-transaction-log-throughput/> <https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on->

NEW QUESTION 52

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 developing a Microsoft SQL Server Integration Services (SSIS) package. The package design consists of the sources shown in the following diagram:



Each source contains data that is not sorted.
 You need to combine data from all of the sources into a single dataset. Which SSIS Toolbox item should you use?

- A. CDC Control task
- B. CDC Splitter
- C. Union All
- D. XML task
- E. Fuzzy Grouping
- F. Merge
- G. Merge Join

Answer: C

NEW QUESTION 56

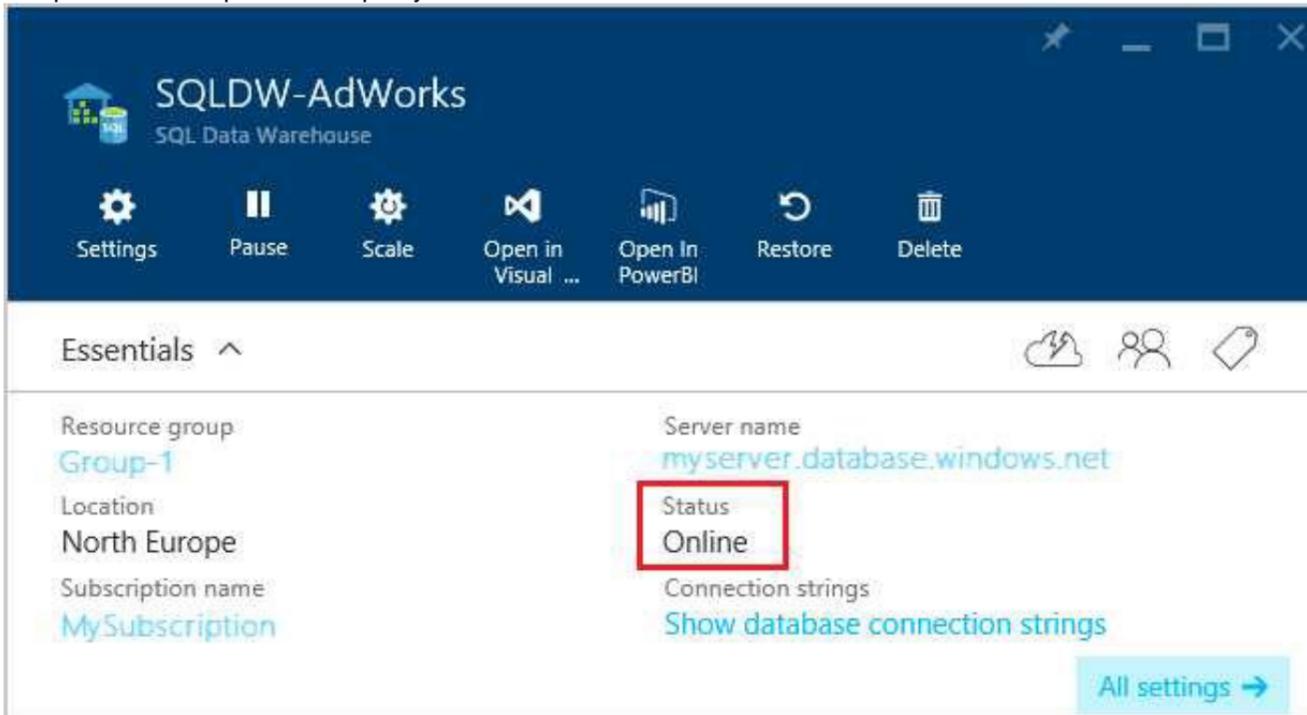
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
 After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
 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

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
 Note: To pause a database:

1. Open the Azure portal and open your database. Notice that the Status is Online.



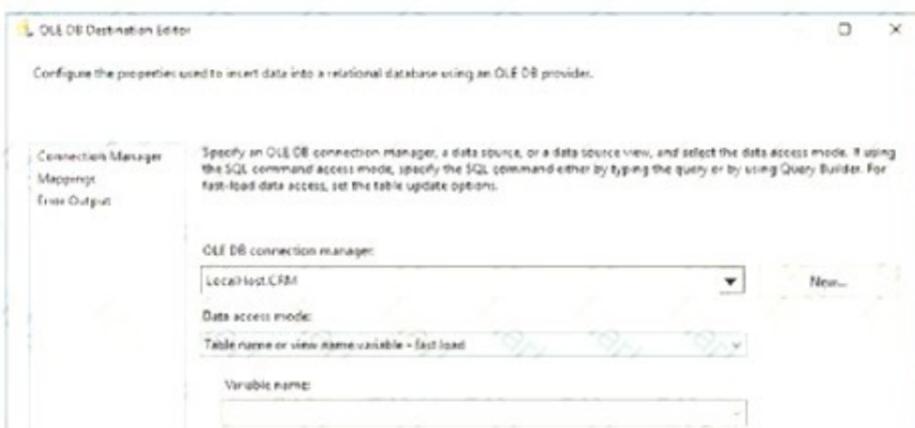
2. To suspend compute and memory resources, click Pause, and then a confirmation message appears. Click yes to confirm or no to cancel.

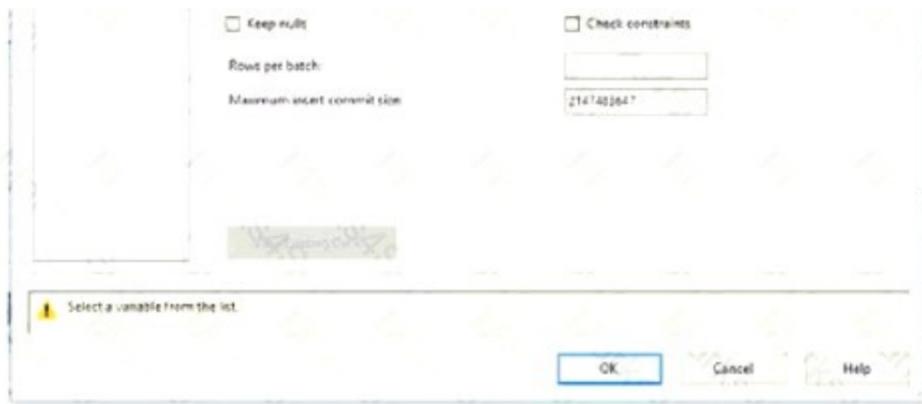
References:

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

NEW QUESTION 61

You need to configure Microsoft SQL Server Integration Services (SSIS) for maximum insert performance. The Integration Services package is configured as shown in the following exhibit.





Use the drop-down menus to select the answer choice that completes each statement based on the information presented in the graphic.
 NOTE: Each correct selection is worth one point.

You need to set Data access mode to [answer choice] to minimize the amount of time it takes to load data to the table.

You need to enable [answer choice] to minimize the amount of time it takes to load data to the table.

Table or view
Table or view – fast load
Table name or view name variable
SQL command

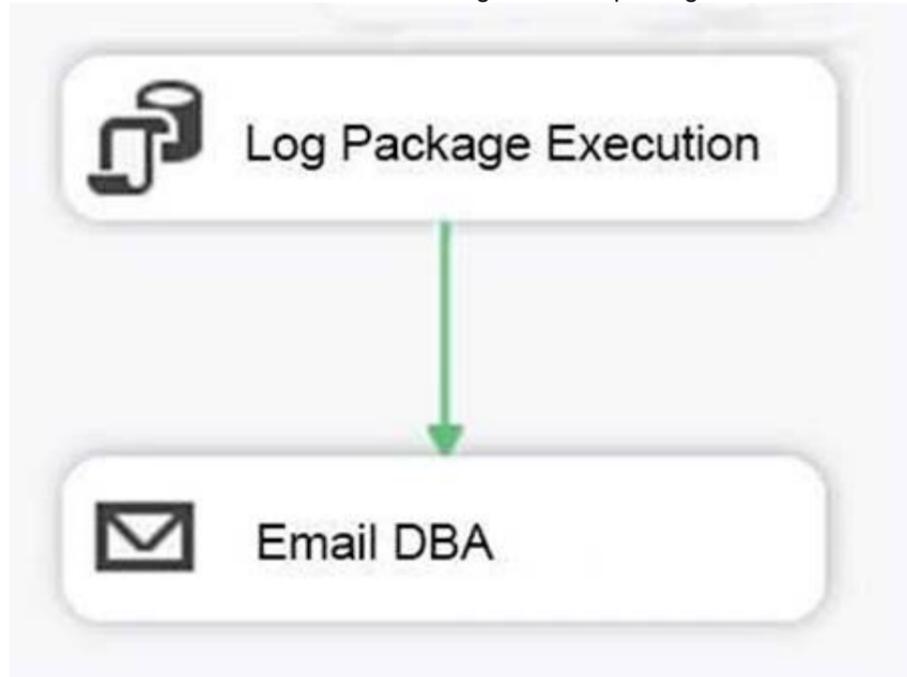
Keep nulls
Keep identity
Check constraints
Table lock

Answer:

Explanation: Table or view Check constraints

NEW QUESTION 66

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You are developing a Microsoft SQL Server Integration Services (SSIS) projects. The project consists of several packages that load data warehouse tables. You need to extend the control flow design for each package to use the following control flow while minimizing development efforts and maintenance:



Solution: You add the control flow to a control flow package part. You add an instance of the control flow package part to each data warehouse load package. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation: A package consists of a control flow and, optionally, one or more data flows. You create the control flow in a package by using the Control Flow tab in SSIS Designer.

References: <https://docs.microsoft.com/en-us/sql/integration-services/control-flow/control-flow>

NEW QUESTION 71

You manage Master Data Services (MDS). You need to create a new entity with the following requirements:

- Maximize the performance of the MDS system.
- Ensure that the Entity change logs are stored.

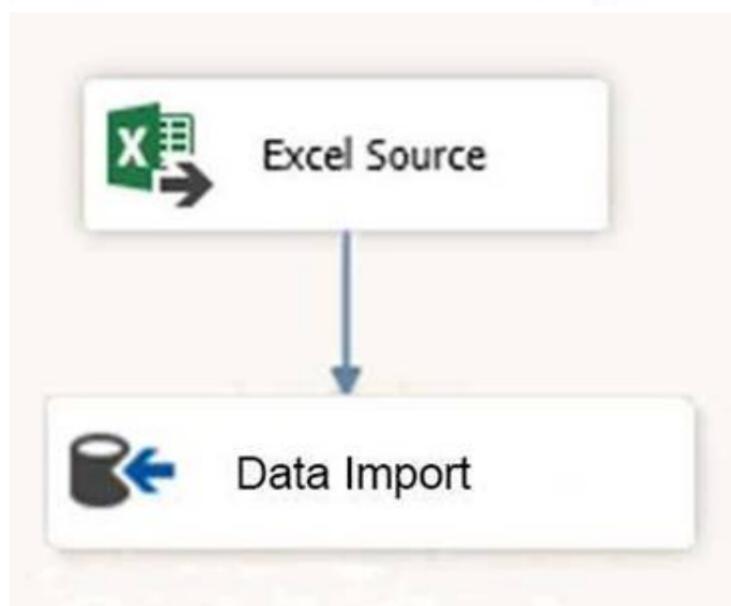
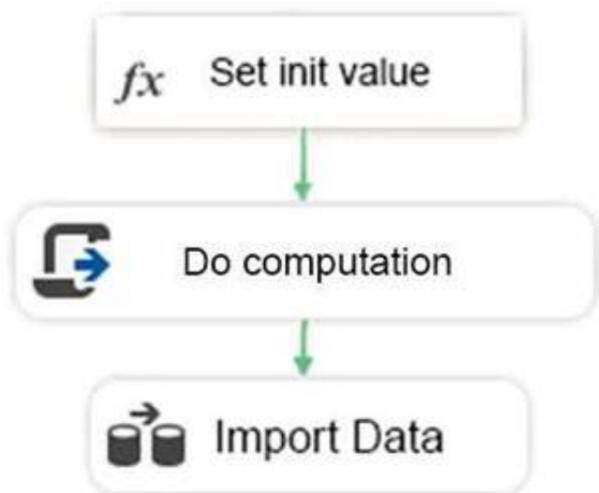
You need to configure the Transaction Log Type setting. Which type should you use?

- A. Full
- B. None
- C. Attribute
- D. Member
- E. Simple

Answer: D

NEW QUESTION 73

You are testing a Microsoft SQL Server Integration Services (SSIS) package. The package includes the Control Flow task shown in the Control Flow exhibit (Click the Exhibit button) and the Data Flow task shown in the Data Flow exhibit. (Click the Exhibit button.)



You declare a variable named Seed as shown in the Variables exhibit. (Click the Exhibit button.) The variable is changed by the Script task during execution.

Variables			
Name	Data type	Value	Expression
Seed	Int32	0	...

You need to be able to interrogate the value of the Seed variable after the Script task completes execution. For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer Area

	Yes	No
You can display the variable by adding a data viewer to the data flow.	<input type="radio"/>	<input type="radio"/>
You can display the variable by adding a breakpoint to the OnPostExecute event and using the Locals window.	<input type="radio"/>	<input type="radio"/>
You can display the variable by adding a breakpoint to the OnVariableValueChanged event and using the Watch window.	<input type="radio"/>	<input type="radio"/>
You can display the variable by adding the following code segment to the Script task: <code>MessageBox.Show</code>	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation: References:
<https://docs.microsoft.com/en-us/sql/integration-services/variables-window>

NEW QUESTION 77

You are the administrator of a Microsoft SQL Server Master Data Services (MDS) model. The model was developed to provide consistent and validated snapshots of master data to the ETL processes by using subscription views. A new model version has been created. You need to ensure that the ETL processes retrieve the latest snapshot of master data. What should you do?

- A. Add a version flag to the new version, and create new subscription views that use this version flag.
- B. Create new subscription views for the new version.
- C. Update the subscription views to use the new version.
- D. Update the subscription views to use the last committed version.

Answer: A

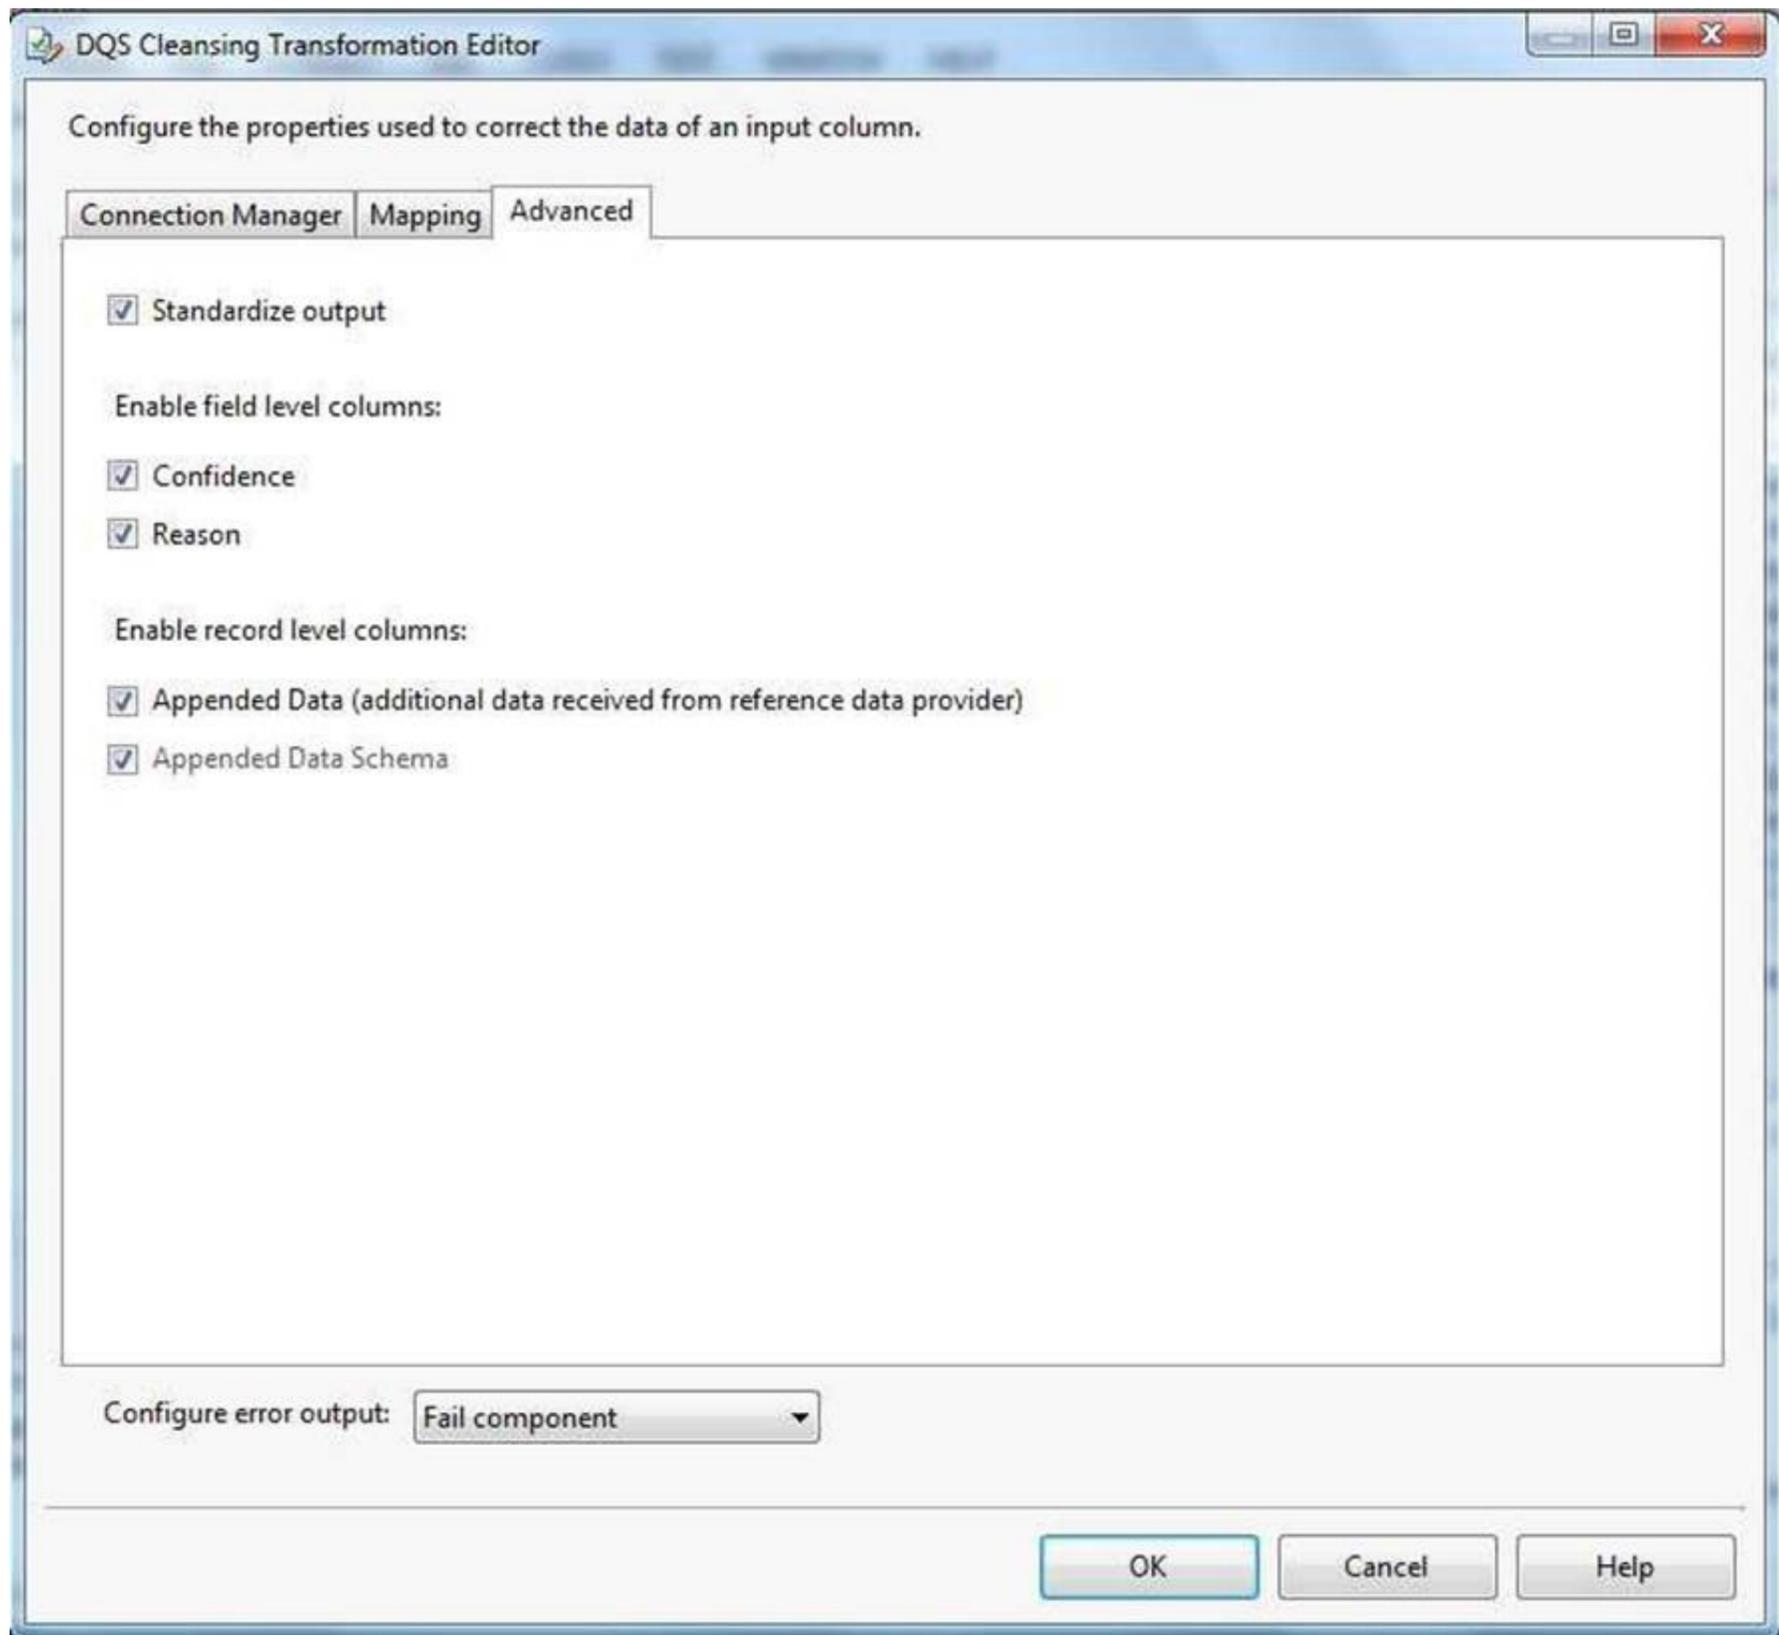
Explanation: When a version is ready for users or for a subscribing system, you can set a flag to identify the version. You can move this flag from version to version as needed. Flags help users and subscribing systems identify which version of a model to use.
 References: <https://docs.microsoft.com/en-us/sql/master-data-services/versions-master-data-services>

NEW QUESTION 79

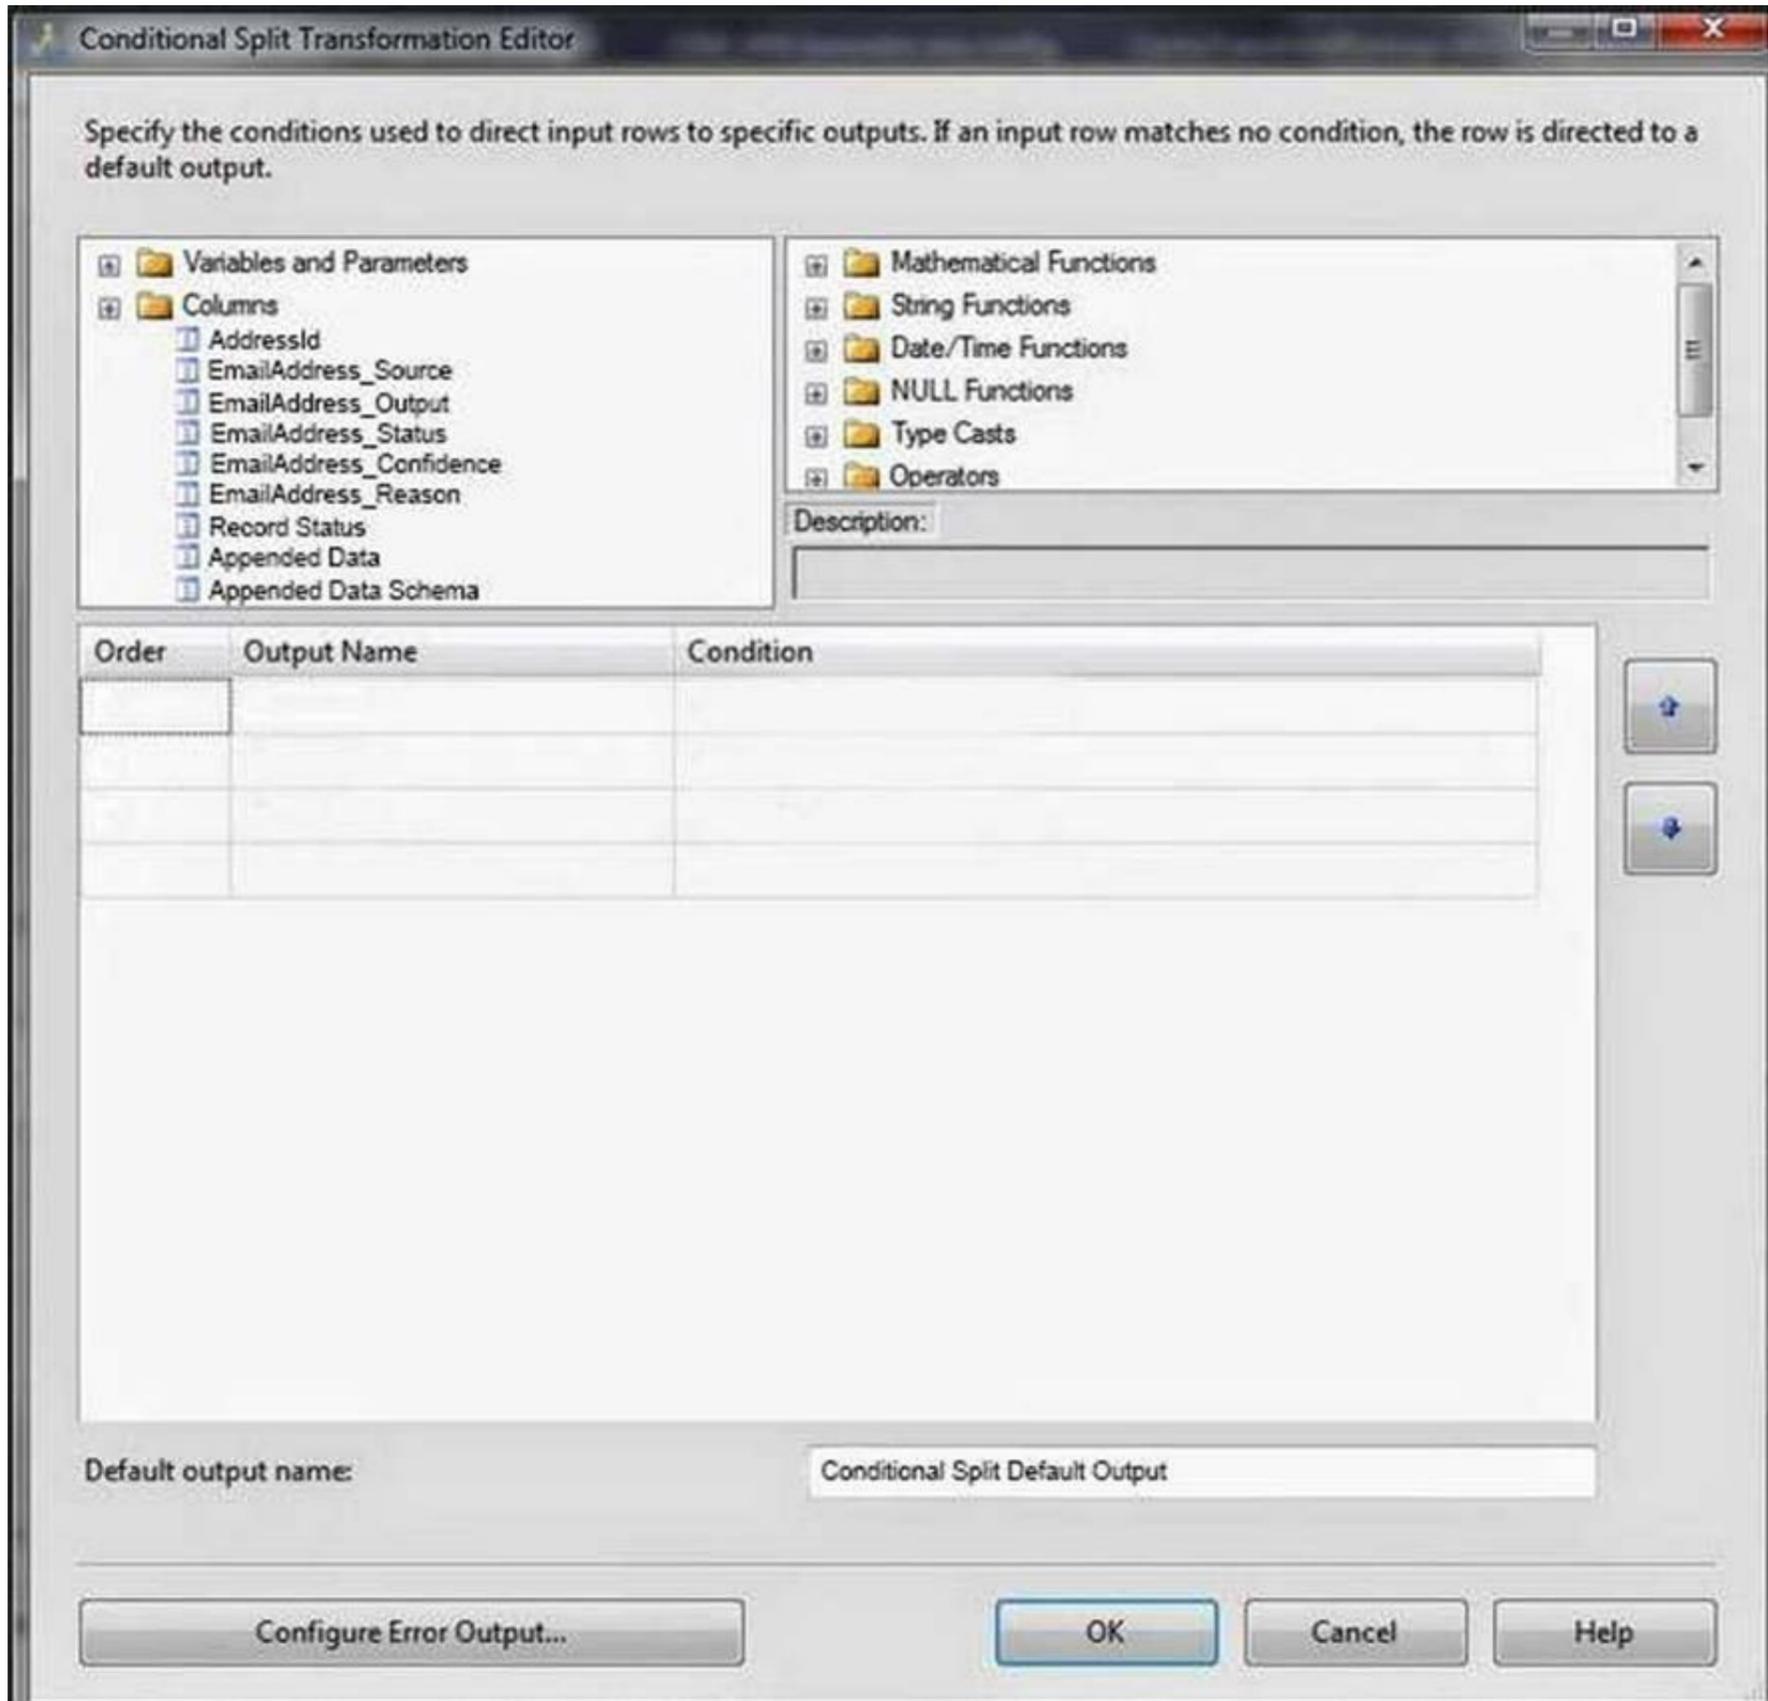
You have a Microsoft SQL Server Integration Services (SSIS) package that contains a Data Flow task as shown in the Data Flow exhibit. (Click the Exhibit button.)



You install Data Quality Services (DQS) on the same server that hosts SSIS and deploy a knowledge base to manage customer email addresses. You add a DQS Cleansing transform to the Data Flow as shown in the Cleansing exhibit. (Click the Exhibit button.)



You create a Conditional Split transform as shown in the Splitter exhibit. (Click the Exhibit button.)



You need to split the output of the DQS Cleansing task to obtain only Correct values from the EmailAddress column. For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer Area

Yes

No

You can use the EmailAddress_Output column to split the output.

You can use the EmailAddress_Status column to split the output.

You can use the EmailAddress_Reason column to split the output.

Answer:

Explanation: The DQS Cleansing component takes input records, sends them to a DQS server, and gets them back corrected. The component can output not only the corrected data, but also additional columns that may be useful for you. For example - the status columns. There is one status column for each mapped field, and another one that aggregated the status for the whole record. This record status column can be very useful in some scenarios, especially when records are further processed in different ways depending on their status. In such cases, it is recommended to use a Conditional Split component below the DQS Cleansing component, and configure it to split the records to groups based on the record status (or based on other columns such as specific field status).
 References: <https://blogs.msdn.microsoft.com/dqs/2011/07/18/using-the-ssis-dqs-cleansing-component/>

NEW QUESTION 81

You need to load data from a CSV file to a table.

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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: Each correct selection is worth one point.

Transact-SQL segments	Answer Area
<div style="display: flex; gap: 10px;"> <div style="border: 1px solid black; padding: 2px 5px;">BULK</div> <div style="border: 1px solid black; padding: 2px 5px;">INSERT</div> </div>	<div style="border: 1px solid black; width: 50px; height: 20px; display: inline-block;"></div> <div style="border: 1px solid black; width: 50px; height: 20px; display: inline-block;"></div> Sales.Invoices
<div style="display: flex; gap: 10px;"> <div style="border: 1px solid black; padding: 2px 5px;">FROM</div> <div style="border: 1px solid black; padding: 2px 5px;">WITH</div> </div>	<div style="border: 1px solid black; width: 50px; height: 20px; display: inline-block;"></div> '\\share\data\file1.csv'
<div style="border: 1px solid black; padding: 2px 5px;">MERGE</div>	<div style="border: 1px solid black; width: 50px; height: 20px; display: inline-block;"></div> (FORMAT = 'CSV')

Answer:

Explanation: The Merge transformation combines two sorted datasets into a single dataset. The rows from each dataset are inserted into the output based on values in their key columns.

By including the Merge transformation in a data flow, you can merge data from two data sources, such as tables and files.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/merge-transformation?view>

NEW QUESTION 86

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

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.

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

NEW QUESTION 90

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 the series.

Start of repeated scenario

Contoso. Ltd. has a Microsoft SQL Server environment that includes SQL Server Integration Services (SSIS), a data warehouse, and SQL Server Analysis Services (SSAS) Tabular and multidimensional models.

The data warehouse stores data related to your company sales, financial transactions and financial budgets All data for the data warehouse originates from the company's business financial system.

The data warehouse includes the following tables:

Table	Notes
dbo.load_City	
dbo.stage_City	
dbo.dim_City	
fact.Sale	
fact.Transaction	This table contains more than 20,000,000 rows. There are currently no indexes on the table. The table has a column named [sale key]. Most queries that target fact.Transaction return recent data based on this column and a column named Description.

The company plans to use Microsoft Azure to store older records from the data warehouse. You must modify the database to enable the Stretch Database capability.

Users report that they are becoming confused about which city table to use for various queries. You plan to create a new schema named Dimension and change the name of the dbo.du_city table to Dimension.city. Data loss is not permissible, and you must not leave traces of the old table in the data warehouse. Pal to create a measure that calculates the profit margin based on the existing measures.

You must improve performance for queries against the fact.Transaction table. You must implement appropriate indexes and enable the Stretch Database capability.

End of repeated scenario

You need to resolve the problems reported about the dia city table.

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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.

Transact-SQL segments

- EXEC sp_rename 'dbo.dim_City', 'City'
- ALTER SCHEMA Dimension TRANSFER dbo.City
- DROP TABLE dbo.dim_City
GO
CREATE TABLE Dimension.City(...)
- SELECT *
INTO Dimension.City
FROM dbo.dim_City
- ALTER TABLE dbo.dim_City
ADD Dimension.City VARCHAR(20) NULL

Answer area

```
CREATE SCHEMA Dimension
GO
Transact-SQL segment
Transact-SQL segment
```

Answer:

Explanation:

Transact-SQL segments

```
EXEC sp_rename 'dbo.dim_City', 'City'

ALTER SCHEMA Dimension TRANSFER dbo.City

DROP TABLE dbo.dim_City
GO
CREATE TABLE Dimension.City( ... )

SELECT *
INTO Dimension.City
FROM dbo.dim_City

ALTER TABLE dbo.dim_City
ADD Dimension.City VARCHAR(20) NULL
```

Answer area

```
CREATE SCHEMA Dimension
GO
ALTER TABLE dbo.dim_City
ADD Dimension.City VARCHAR(20) NULL

DROP TABLE dbo.dim_City
GO
CREATE TABLE Dimension.City( ... )
```

NEW QUESTION 94

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You have a Microsoft SQL server that has Data Quality Services (DQS) installed. You need to review the completeness and the uniqueness of the data stored in the matching policy. Solution: You modify the weight of the domain in the matching rule.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation: Use a matching rule, and use completeness and uniqueness data to determine what weight to give a field in the matching process.

If there is a high level of uniqueness in a field, using the field in a matching policy can decrease the matching results, so you may want to set the weight for that field to a relatively small value. If you have a low level of uniqueness for a column, but low completeness, you may not want to include a domain for that column.

References:

<https://docs.microsoft.com/en-us/sql/data-quality-services/create-a-matching-policy?view=sql-server-2017>

NEW QUESTION 97

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

You are the administrator of a Microsoft SQL Server Master Data Services (MDS) instance. The instance contains a model named Geography and a model named customer. The Geography model contains an entity named countryRegion.

You need to ensure that the countryRegion entity members are available in the customer model. Solution: Configure an entity sync relationship to replicate the CountryRegion entity.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 100

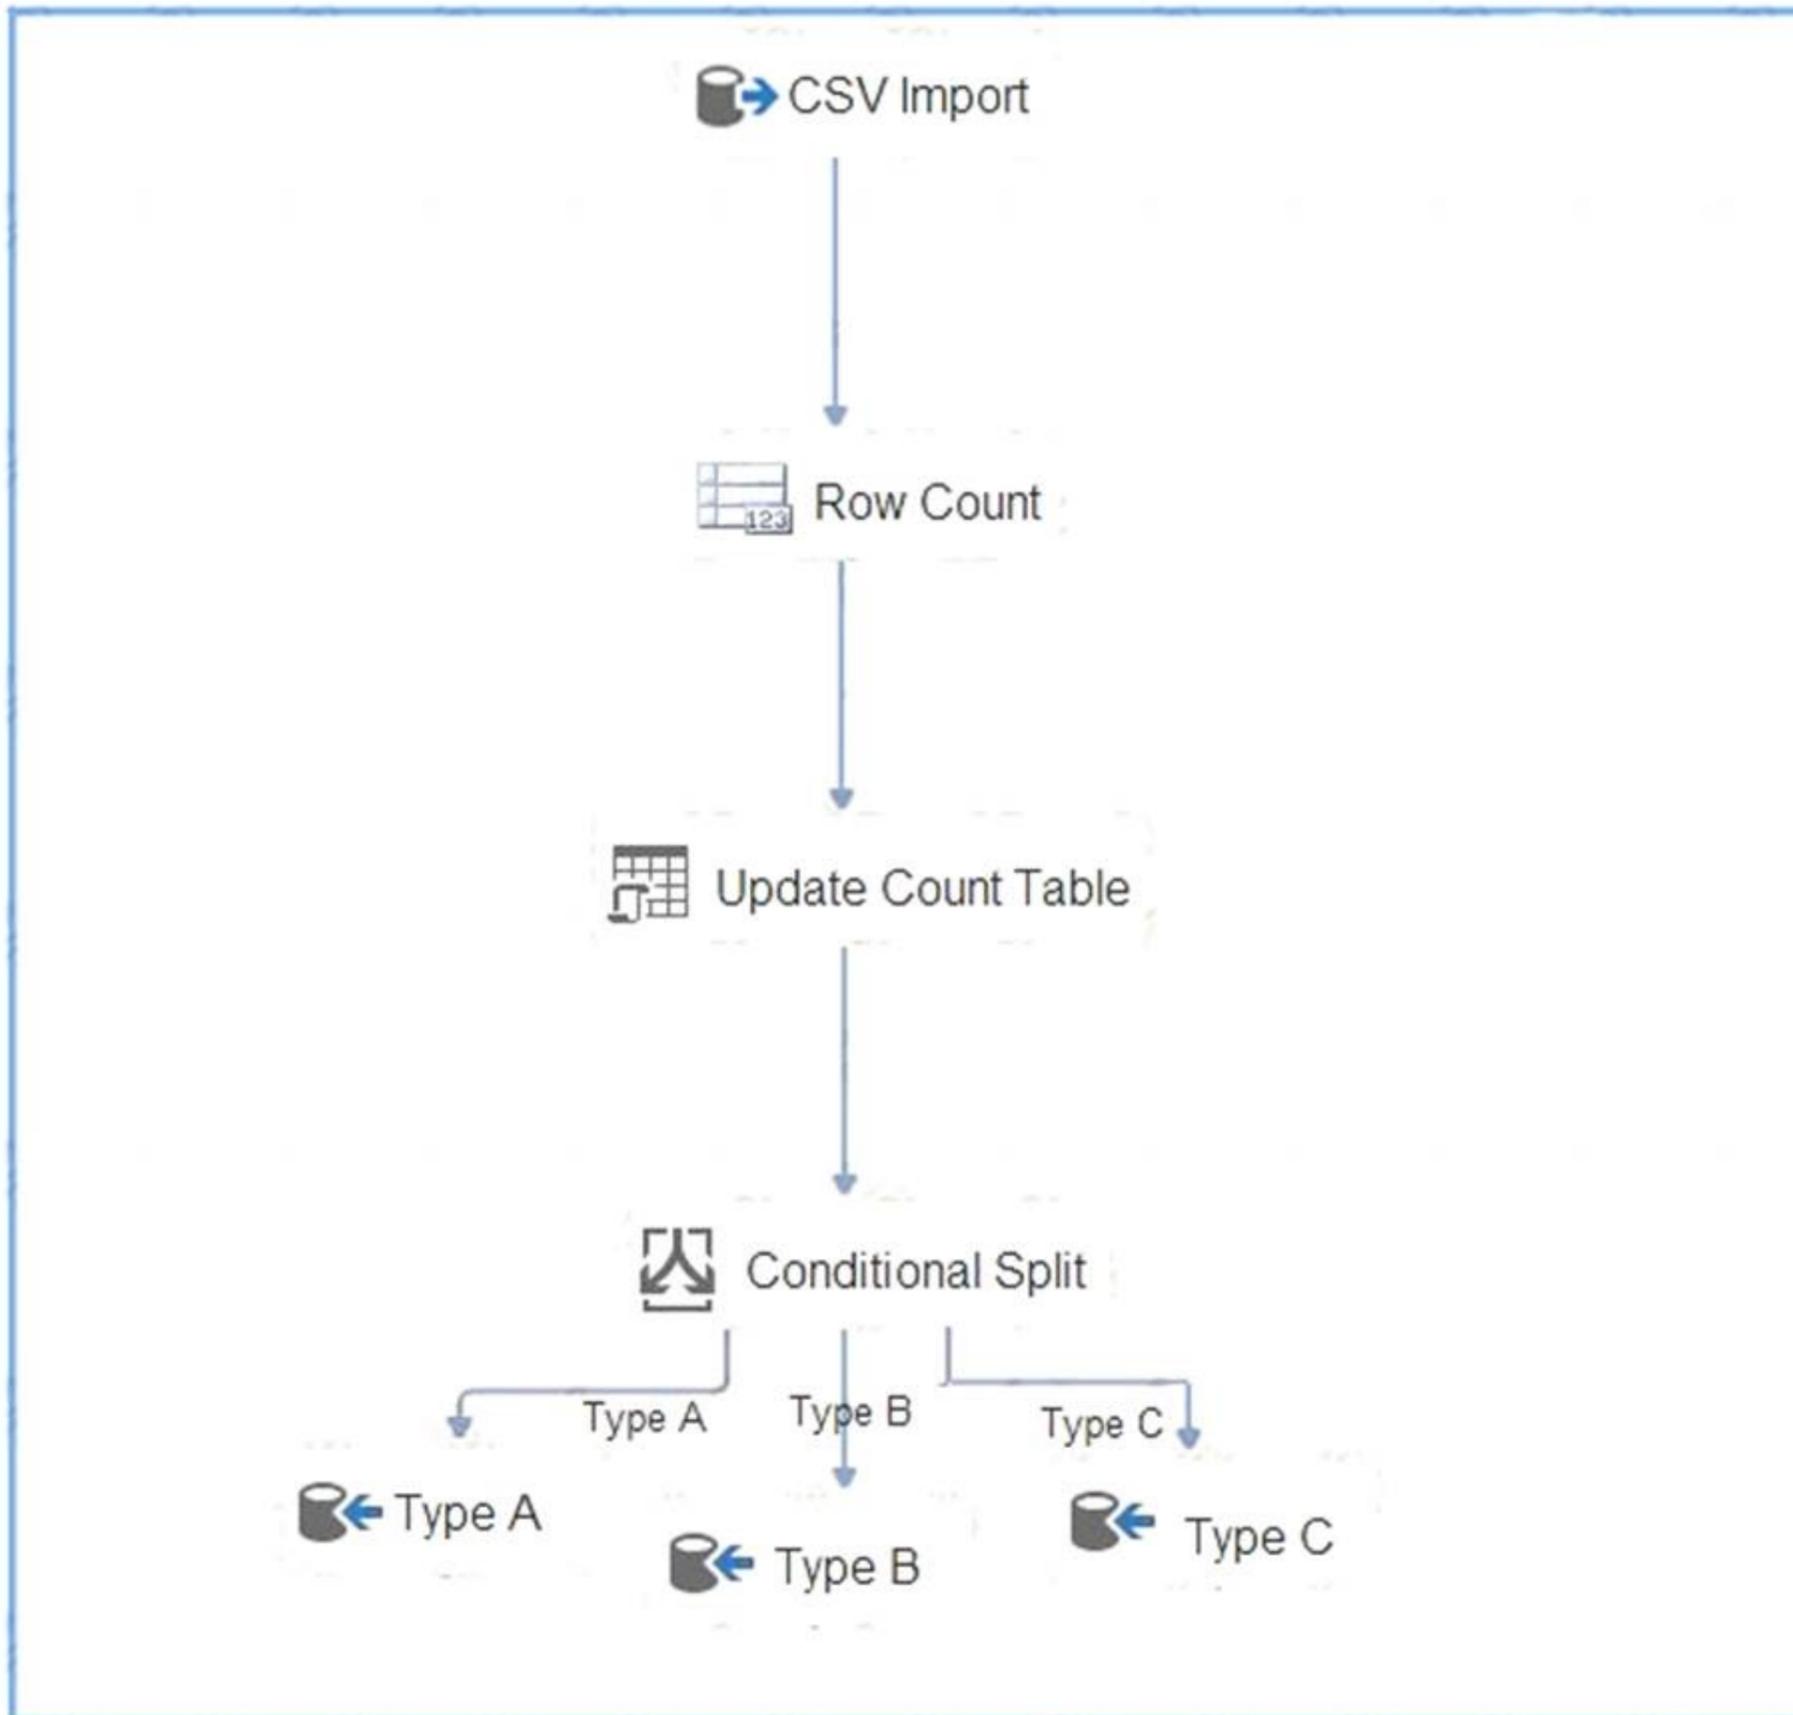
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Each night you receive a comma separated values (CSV) file that contains different types of rows. Each row type has a different structure. Each row in the CSV file is unique. The first column in every row is named Type. This column identifies the data type.

For each data type, you need to load data from the CSV file to a target table. A separate table must contain the number of rows loaded for each data type.

Solution: You create a SQL Server Integration Services (SSIS) package as shown in the exhibit. (Click the Exhibit tab.)



Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: The conditional split must be before the count.

NEW QUESTION 103

You have a Microsoft SQL Server Integration Services (SSIS) package that loads data into a data warehouse each night from a transactional system. The package also loads data from a set of Comma-Separated Values (CSV) files that are provided by your company’s finance department.

The SSIS package processes each CSV file in a folder. The package reads the file name for the current file into a variable and uses that value to write a log entry to a database table.

You need to debug the package and determine the value of the variable before each file is processed.

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.

Actions

Click the **Start** toolbar button to commence debugging the package.

When a breakpoint is reached, view the value of the variable by using the Variables window.

Open the Control Flow editor for the package.

When a breakpoint is reached, view the value of the variable by using the Locals window.

Set a breakpoint on the For Loop container.

Set a breakpoint on the Sequence container.

Open the Data Flow editor for the package.

Set a breakpoint on the Foreach Loop container.

Answer Area



Answer:

Explanation: You debug control flows.

The Foreach Loop container is used for looping through a group of files. Put the breakpoint on it.

The Locals window displays information about the local expressions in the current scope of the Transact-SQL debugger.

References: <https://docs.microsoft.com/en-us/sql/integration-services/troubleshooting/debugging-control-flow>

<http://blog.pragmaticworks.com/looping-through-a-result-set-with-the-foreach-loop>

NEW QUESTION 108

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 have a database named DB1.

You need to track auditing data for four tables in DB1 by using change data capture. Which stored procedure should you execute first?

- A. catalog.deploy_project
- B. catalog.restore_project
- C. catalog.stop_operation
- D. sys.sp_cdc_add_job
- E. sys.sp_cdc_change_job
- F. sys.sp_cdc_disable_db

Answer: D

Explanation: Because the cleanup and capture jobs are created by default, the sys.sp_cdc_add_job stored procedure is necessary only when a job has been explicitly dropped and must be recreated.

Note: sys.sp_cdc_add_job creates a change data capture cleanup or capture job in the current database. A cleanup job is created using the default values when the first table in the database is enabled for change data capture. A capture job is created using the default values when the first table in the database is enabled for change data capture and no transactional publications exist for the database. When a transactional publication exists, the transactional log reader is used to drive the capture mechanism, and a separate capture job is neither required nor allowed.

Note: sys.sp_cdc_change_job

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/track-changes/track-data-changes-sqlserver>

NEW QUESTION 112

You manage an inventory system that has a table named Products. The Products table has several hundred columns. You generate a report that relates two columns named ProductReference and ProductName from the Products table. The result is sorted by a column named QuantityInStock from largest to smallest. You need to create an index that the report can use. How should you complete the Transact-SQL statement? To answer, select the appropriate Transact-SQL segments in the answer area.

Answer Area

CREATE INDEX lx_product

CLUSTERED
 NONCLUSTERED

ON dbo.Products

(ProductReference)
 (QuantityInStock)
 (ProductName)

INCLUDE

(Products)
 (ProductReference)
 (ProductName, ProductReference)

Answer:

Explanation:

Answer Area

CREATE INDEX lx_product

CLUSTERED
 NONCLUSTERED

ON dbo.Products

(ProductReference)
 (QuantityInStock)
 (ProductName)

INCLUDE

(Products)
 (ProductReference)
 (ProductName, ProductReference)

NEW QUESTION 114

You are developing a data warehouse. You run the following Transact-SQL statement:

```
USE AdventureWorks
GO
CREATE TABLE Production.TransactionHistoryArchive(
TransactionID INT IDENTITY (1, 1) NOT NULL,
CONSTRAINT PK_TransactionHistoryArchive_TransactionID PRIMARY KEY CLUSTERED (TransactionID)
)
```

Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic. NOTE: Each correct selection is worth one point.

What is the name of the table created?

AdventureWorks
Production
TransactionHistoryArchive

What is the name of the primary key?

Identity
Production
TransactionID

Answer:

Explanation:

What is the name of the table created?

AdventureWorks
Production
TransactionHistoryArchive

What is the name of the primary key?

Identity
Production
TransactionID

NEW QUESTION 115

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 have a database named DB1 that has change data capture enabled.

A Microsoft SQL Server Integration Services (SSIS) job runs once weekly. The job loads changes from DB1 to a data warehouse by querying the change data capture tables.

Users report that an application that uses DB1 is suddenly unresponsive.

You discover that the Integration Services job causes severe blocking issues in the application. You need to ensure that the users can run the application as quickly as possible. Your SQL Server login is a member of only the ssis.admin database role.

Which stored procedure should you execute?

- A. catalog.deploy_project
- B. catalog.restore_project
- C. catalog.stop.operation
- D. sys.sp.cdc.addjob
- E. sys.sp.cdc.changejob
- F. sys.sp_cdc_disable_db
- G. sys.sp_cdc_enable_db
- H. sys.sp_cdc.stopJob

Answer: E

Explanation: sys.sp_cdc_change_job modifies the configuration of a change data capture cleanup or capture job in the current database.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-change-job-trans>

NEW QUESTION 120

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Your company uses Microsoft SQL Server to deploy a data warehouse to an environment that has a SQL Server Analysis Services (SSAS) instance. The data warehouse includes the Fact.Order table as shown in the following table definition. The table has no indexes.

- Columns
- Order Key (bigint, not null)
 - City Key (int, not null)
 - Customer Key (int, not null)
 - Stock Item Key (int, not null)
 - Order Date Key (date, not null)
 - Picked Date Key (date, null)
 - Salesperson Key (int, not null)
 - Picker Key (int, null)
 - Quantity (int, not null)
 - Unit Price (decimal(18,2), not null)
 - Tax Rate (decimal(18,3), not null)
 - Total Excluding Tax (decimal(18,2), not null)
 - Tax Amount (decimal(18,2), not null)
 - Total Including Tax (decimal(18,2), not null)

You must minimize the amount of space that indexes for the Fact.Order table consume. You run the following queries frequently. Both queries must be able to use a columnstore index:

```
SELECT AVG([Tax Amount]) AS [Average Tax Amount]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'

SELECT SUM([Total Excluding Tax]) AS [Total Revenue]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'
```

You need to ensure that the queries complete as quickly as possible.

Solution You create two nonclustered indexes. The first includes the [Order Date Key] and [Tax Amount] columns. The second will include the [Order Date Key] and [Total Excluding Tax] columns.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 125

A database has tables named Table1, Table2, and Table3.

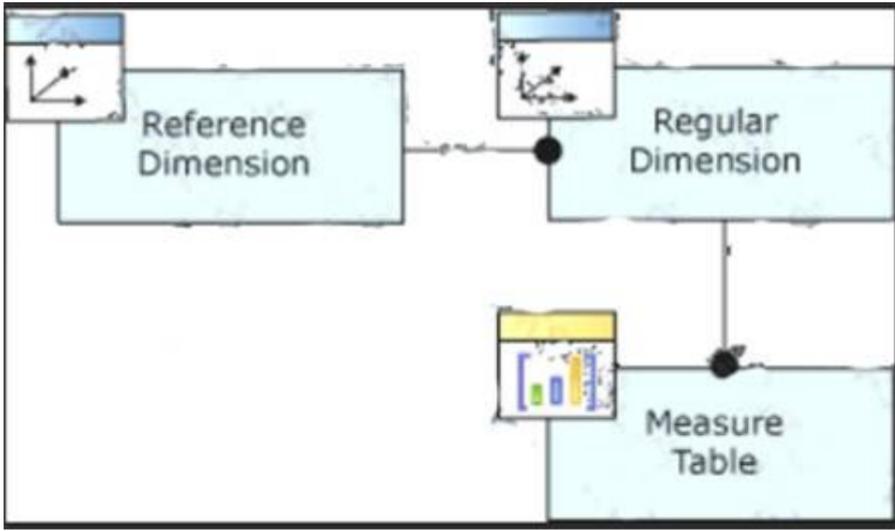
- ▶ Table1 has a foreign key relationship with Table2.
- ▶ Table2 has a foreign key relationship with Table3.
- ▶ Table1 does not have a direct relationship with Table3.

You need to recommend an appropriate dimension usage relationship. What should you recommend?

- A. many-to-one relationship
- B. referenced relationship
- C. regular dimension relationship
- D. fact relationship

Answer: B

Explanation: A reference dimension relationship between a cube dimension and a measure group exists when the key column for the dimension is joined indirectly to the fact table through a key in another dimension table, as shown in the following illustration.



NEW QUESTION 130

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it As a result, these questions will not appear in the review screen. You have an on-premises Microsoft SQL Server instance and a Microsoft Azure SQL Data Warehouse instance. You move data from the on-premises database to the data warehouse once each day by using a SQL Server Integration Services (SSIS) package. You observe that the package no longer completes within the allotted time. You need to determine which tasks are taking a long time to complete. Solution: You enable package logging within SSIS. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 135

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 the series. Start of repeated scenario Contoso. Ltd. has a Microsoft SQL Server environment that includes SQL Server Integration Services (SSIS), a data warehouse, and SQL Server Analysis Services (SSAS) Tabular and multidimensional models. The data warehouse stores data related to your company sales, financial transactions and financial budgets. All data for the data warehouse originates from the company's business financial system. The data warehouse includes the following tables:

Table	Notes
dbo.load_City	
dbo.stage_City	
dbo.dim_City	
fact.Sale	
fact.Transaction	This table contains more than 20,000,000 rows. There are currently no indexes on the table. The table has a column named [sale key]. Most queries that target fact.Transaction return recent data based on this column and a column named Description.

The company plans to use Microsoft Azure to store older records from the data warehouse. You must modify the database to enable the Stretch Database capability. Users report that they are becoming confused about which city table to use for various queries. You plan to create a new schema named Dimension and change the name of the dbo.du_city table to Dimension.city. Data loss is not permissible, and you must not leave traces of the old table in the data warehouse. Pal to create a measure that calculates the profit margin based on the existing measures. You must implement a partitioning scheme few the fact. Transaction table to move older data to less expensive storage. Each partition will store data for a single calendar year, as shown in the exhibit (Click the Exhibit button.) You must align the partitions.

	Transaction Key	Date Key	Customer Key	Bill To Customer Key	Supplier Key	Transaction Type Key	Payment Method Key	WWI Invoice ID
1	7	2013-01-01	375	202	0	1	0	7
2	11	2013-01-01	387	202	0	1	0	11
3	12	2013-01-01	330	202	0	1	0	12
4	13	2013-01-01	274	202	0	1	0	13
5	16	2013-01-01	215	202	0	1	0	16
6	25	2013-01-01	298	202	0	1	0	25
7	26	2013-01-01	285	202	0	1	0	26
8	30	2013-01-01	368	202	0	1	0	30
9	35	2013-01-01	232	202	0	1	0	35
10	39	2013-01-01	346	202	0	1	0	39
11	41	2013-01-01	216	202	0	1	0	41
12	63	2013-01-02	224	202	0	1	0	42
13	64	2013-01-02	264	202	0	1	0	43
14	65	2013-01-02	268	202	0	1	0	44
15	70	2013-01-02	376	202	0	1	0	49
16	74	2013-01-02	387	202	0	1	0	53
17	75	2013-01-02	330	202	0	1	0	54
16	74	2013-01-02	387	202	0	1	0	53
17	75	2013-01-02	330	202	0	1	0	54
18	76	2013-01-02	274	202	0	1	0	55
19	78	2013-01-02	215	202	0	1	0	57
20	85	2013-01-02	298	202	0	1	0	64
21	86	2013-01-02	285	202	0	1	0	65
22	90	2013-01-02	368	202	0	1	0	69
23	94	2013-01-02	232	202	0	1	0	73

You must improve performance for queries against the fact.Transaction table. You must implement appropriate indexes and enable the Stretch Database capability.

End of repeated scenario

You need to resolve the problems reported about the dia city table.

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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.

Answer area

```
CREATE CLUSTERED COLUMNSTORE INDEX idx_fact_sale ON fact.Sale
CREATE NONCLUSTERED COLUMNSTORE INDEX idx_fact_sale ON fact.Sale
ALTER INDEX idx_fact_sale ON fact.Sale DISABLE
```

```
WITH (DROP_EXISTING = ON)
DROP INDEX idx_fact_sale ON fact.Sale
ALTER INDEX idx_fact_sale ON fact.Sale REBUILD
CREATE CLUSTERED COLUMNSTORE INDEX idx_fact_sale_cs ON fact.Sale
```

Answer:

Explanation:

Answer area

```
CREATE CLUSTERED COLUMNSTORE INDEX idx_fact_sale ON fact.Sale  
CREATE NONCLUSTERED COLUMNSTORE INDEX idx_fact_sale ON fact.Sale  
ALTER INDEX idx_fact_sale ON fact.Sale DISABLE
```

```
WITH (DROP_EXISTING = ON)  
DROP INDEX idx_fact_sale ON fact.Sale  
ALTER INDEX idx_fact_sale ON fact.Sale REBUILD  
CREATE CLUSTERED COLUMNSTORE INDEX idx_fact_sale_cs ON fact.Sale
```

NEW QUESTION 136

You have a database named DB1 that contains millions of rows. You plan to perform a weekly audit of the changes to the rows. You need to ensure that you can view which rows were modified and the hour that the modification occurred. What should you do?

- A. Enable Policy-Based Management
- B. Configure Stretch Database.
- C. Configure an SSIS database.
- D. Enable change data capture.

Answer: D

Explanation: SQL Server 2017 provides two features that track changes to data in a database: change data capture and change tracking. Change data capture provides historical change information for a user table by capturing both the fact that DML changes were made and the actual data that was changed. Changes are captured by using an asynchronous process that reads the transaction log and has a low impact on the system.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/track-changes/track-data-changes-sql-server>

NEW QUESTION 138

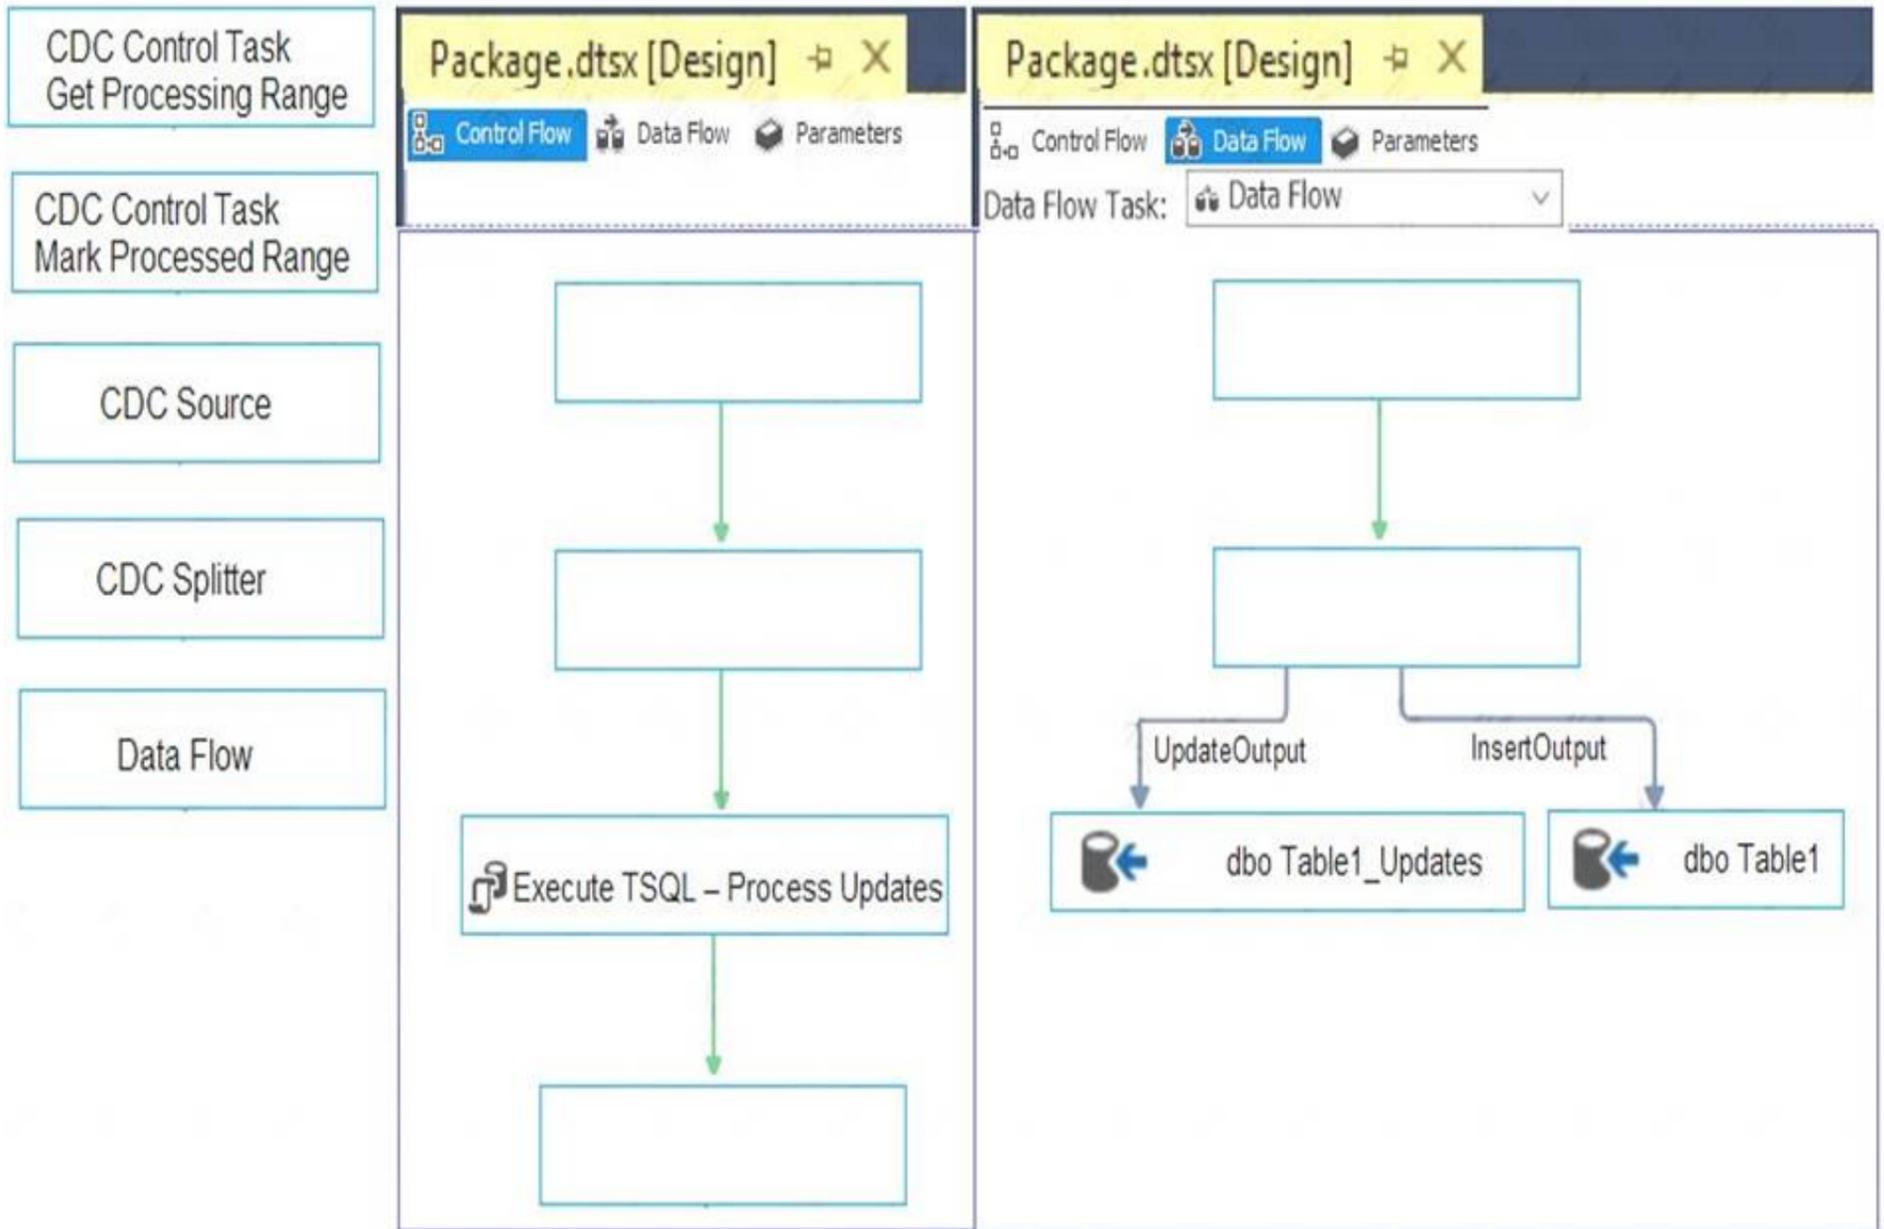
You are developing a Microsoft SQL Server Integration Services (SSIS) package to incrementally load new and changed records from a data source. The SSIS package must load new records into Table1 and updated records into Table1_Updates. After loading records, the package must call a Transact-SQL statement to process updated rows according to existing business logic.

You need to complete the design of the SSIS package.

Which tasks should you use? To answer, drag the appropriate SSIS objects to the correct targets. Each SSIS object 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: Each correct selection is worth one point.

Answer Area



Answer:

Explanation: Step 1: CDC Control Task Get Processing Range Step 2: Mark Processed Range

Step 3: Data Flow

The Data Flow task encapsulates the data flow engine that moves data between sources and destinations, and lets the user transform, clean, and modify data as it is moved. Addition of a Data Flow task to a package control flow makes it possible for the package to extract, transform, and load data.

Step 4: CDC Source

The CDC source reads a range of change data from SQL Server 2017 change tables and delivers the changes downstream to other SSIS component.

Step 5: CDC Splitter

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.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/control-flow/cdc-control-task> <https://docs.microsoft.com/en-us/sql/integration-services/control-flow/data-flow-task> <https://docs.microsoft.com/en-us/sql/integration-services/data-flow/cdc-splitter?view=sql-server-2017>

NEW QUESTION 142

You create a Master Data Services (MDS) model that manages the master data for a Product dimension. The Product dimension has the following properties: All the members of the Product dimension have a product type, a product subtype, and a unique product name.

Each product has a single product type and a single product subtype. The product type has a one-to-many relationship to the product subtype.

You need to ensure that the relationship between the product name, the product type, and the product subtype is maintained when products are added to or updates in the database.

What should you add to the model?

- A. a subscription view
- B. a derived hierarchy
- C. a recursive hierarchy
- D. an explicit hierarchy

Answer: B

Explanation: A Master Data Services derived hierarchy is derived from the domain-based attribute relationships that already exist between entities in a model. You can create a derived hierarchy to highlight any of the existing domain-based attribute relationships in the model.

NEW QUESTION 144

You have a Microsoft SQL Server Data Warehouse instance that uses SQL Server Analysis Services (SSAS). The instance has a cube containing data from an on-

premises SQL Server instance. A measure named Measure1 is configured to calculate the average of a column. You plan to change Measure1 to a full additive measure and create a new measure named Measure2 that evaluates data based on the first populated row. You need to configure the measures. What should you do? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

Measure	Action
Measure1	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px;">▼</div> <p>Turn off semi-additive behavior.</p> <p>Enable the First Child semi-additive function.</p> <p>Enable the FirstNonEmpty semi-additive function.</p> <p>Enable the LastNoneEmpty semi-additive function.</p> <p>Enable the Count semi-additive function.</p> <p>Enable the None semi-additive function.</p> </div>
Measure2	<div style="border: 1px solid gray; padding: 5px;"> <div style="background-color: #f0f0f0; padding: 2px;">▼</div> <p>Turn off semi-additive behavior.</p> <p>Enable the First Child semi-additive function.</p> <p>Enable the FirstNonEmpty semi-additive function.</p> <p>Enable the LastNoneEmpty semi-additive function.</p> <p>Enable the Count semi-additive function.</p> <p>Enable the None semi-additive function.</p> </div>

Answer:

Explanation: Box 1:

The default setting is SUM (fully additive). Box 2:

FirstNonEmpty: The member value is evaluated as the value of its first child along the time dimension that contains data.

References:

<https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models/define-semiadditive-behavior>

NEW QUESTION 148

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 designing a data warehouse and the load process for the data warehouse.

You have a source system that contains two tables named Table1 and Table2. All the rows in each table have a corresponding row in the other table.

The primary key for Table1 is named Key1. The primary key for Table2 is named Key2.

You need to combine both tables into a single table named Table3 in the data warehouse. The solution must ensure that all the nonkey columns in Table1 and Table2 exist in Table3. Which component should you use to load the data to the data warehouse?

- A. the Slowly Changing Dimension transformation
- B. the Conditional Split transformation
- C. the Merge transformation
- D. the Data Conversion transformation
- E. an Execute SQL task
- F. the Aggregate transformation
- G. the Lookup transformation

Answer: G

Explanation: The Lookup transformation performs lookups by joining data in input columns with columns in a reference dataset. You use the lookup to access additional information in a related table that is based on values in common columns.

You can configure the Lookup transformation in the following ways: Specify joins between the input and the reference dataset.

Add columns from the reference dataset to the Lookup transformation output. Etc.

NEW QUESTION 149

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 implementing a Microsoft SQL Server data warehouse with a multi-dimensional data model. When testing a pilot version of the data warehouse, business users observe that the number of products in

stock is inaccurate. The number of products in stock always increases and represents the total number of products that have ever been in stock.

You need to correct the existing model and ensure that it reflects the number of in-stock products. You must not change the overall structure of the data model.

What should you do?

- A. star schema

- B. snowflake schema
- C. conformed dimension
- D. slowly changing dimension (SCD)
- E. fact table
- F. semi-additive measure
- G. non-additive measure
- H. dimension table reference relationship

Answer: H

NEW QUESTION 153

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You plan to deploy a Microsoft SQL server that will host a data warehouse named DB1. The server will contain four SATA drives configured as a RAID 10 array. You need to minimize write contention on the transaction log when data is being loaded to the database. Solution: You configure the server to automatically delete the transaction logs nightly. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation: You should place the log file on a separate drive. References: <https://www.red-gate.com/simple-talk/sql/database-administration/optimizing-transaction-log-throughput/> <https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on->

NEW QUESTION 158

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

NEW QUESTION 160

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have a data warehouse that stores information about products, sales, and orders for a manufacturing company. The instance contains a database that has two tables named SalesOrderHeader and SalesOrderDetail. SalesOrderHeader has 500,000 rows and SalesOrderDetail has 3,000,000 rows. Users report performance degradation when they run the following stored procedure:

```
CREATE PROCEDURE Sales.GetRecentSales (@date datetime)
AS BEGIN
    IF @date is NULL
        SET @date = DATEADD(MONTH, -3, (SELECT MAX(ORDERDATE) FROM Sales.SalesOrderHeader))
    SELECT * FROM Sales.SalesOrderHeader h, Sales.SalesOrderDetail d
    WHERE h.SalesOrderID = d.SalesOrderID
    AND h.OrderDate > @date
END
```

You need to optimize performance. Solution: You run the following Transact-SQL statement:

```
CREATE STATISTICS Stat1
On Sales.SalesOrderHeader (OrderDate)
WITH SAMPLE 100 ROWS
```

Does the solution meet the goal?

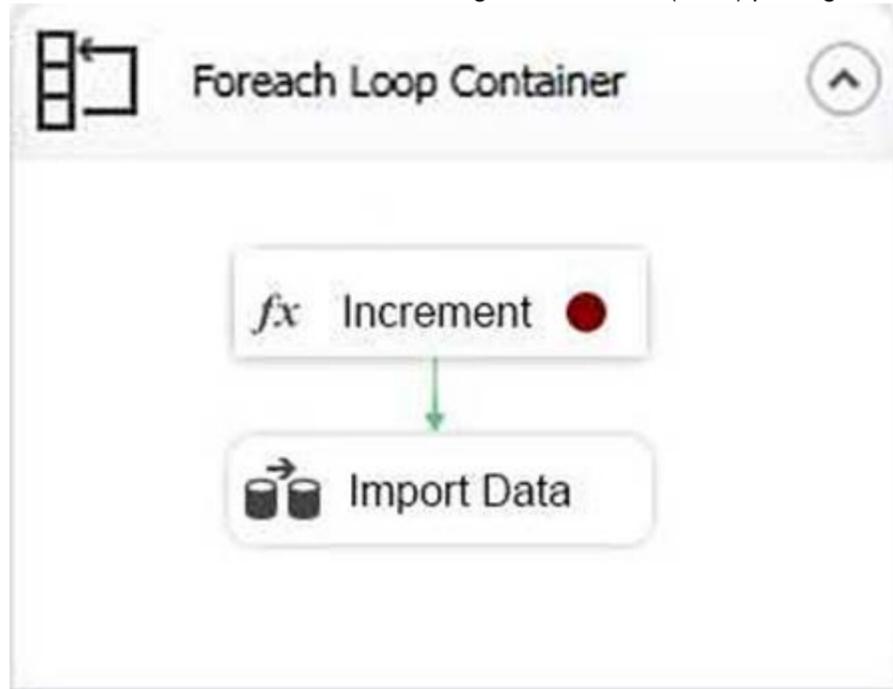
- A. Yes
- B. No

Answer: B

Explanation: 100 out of 500,000 rows is a too small sample size.
 References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-statistics>

NEW QUESTION 165

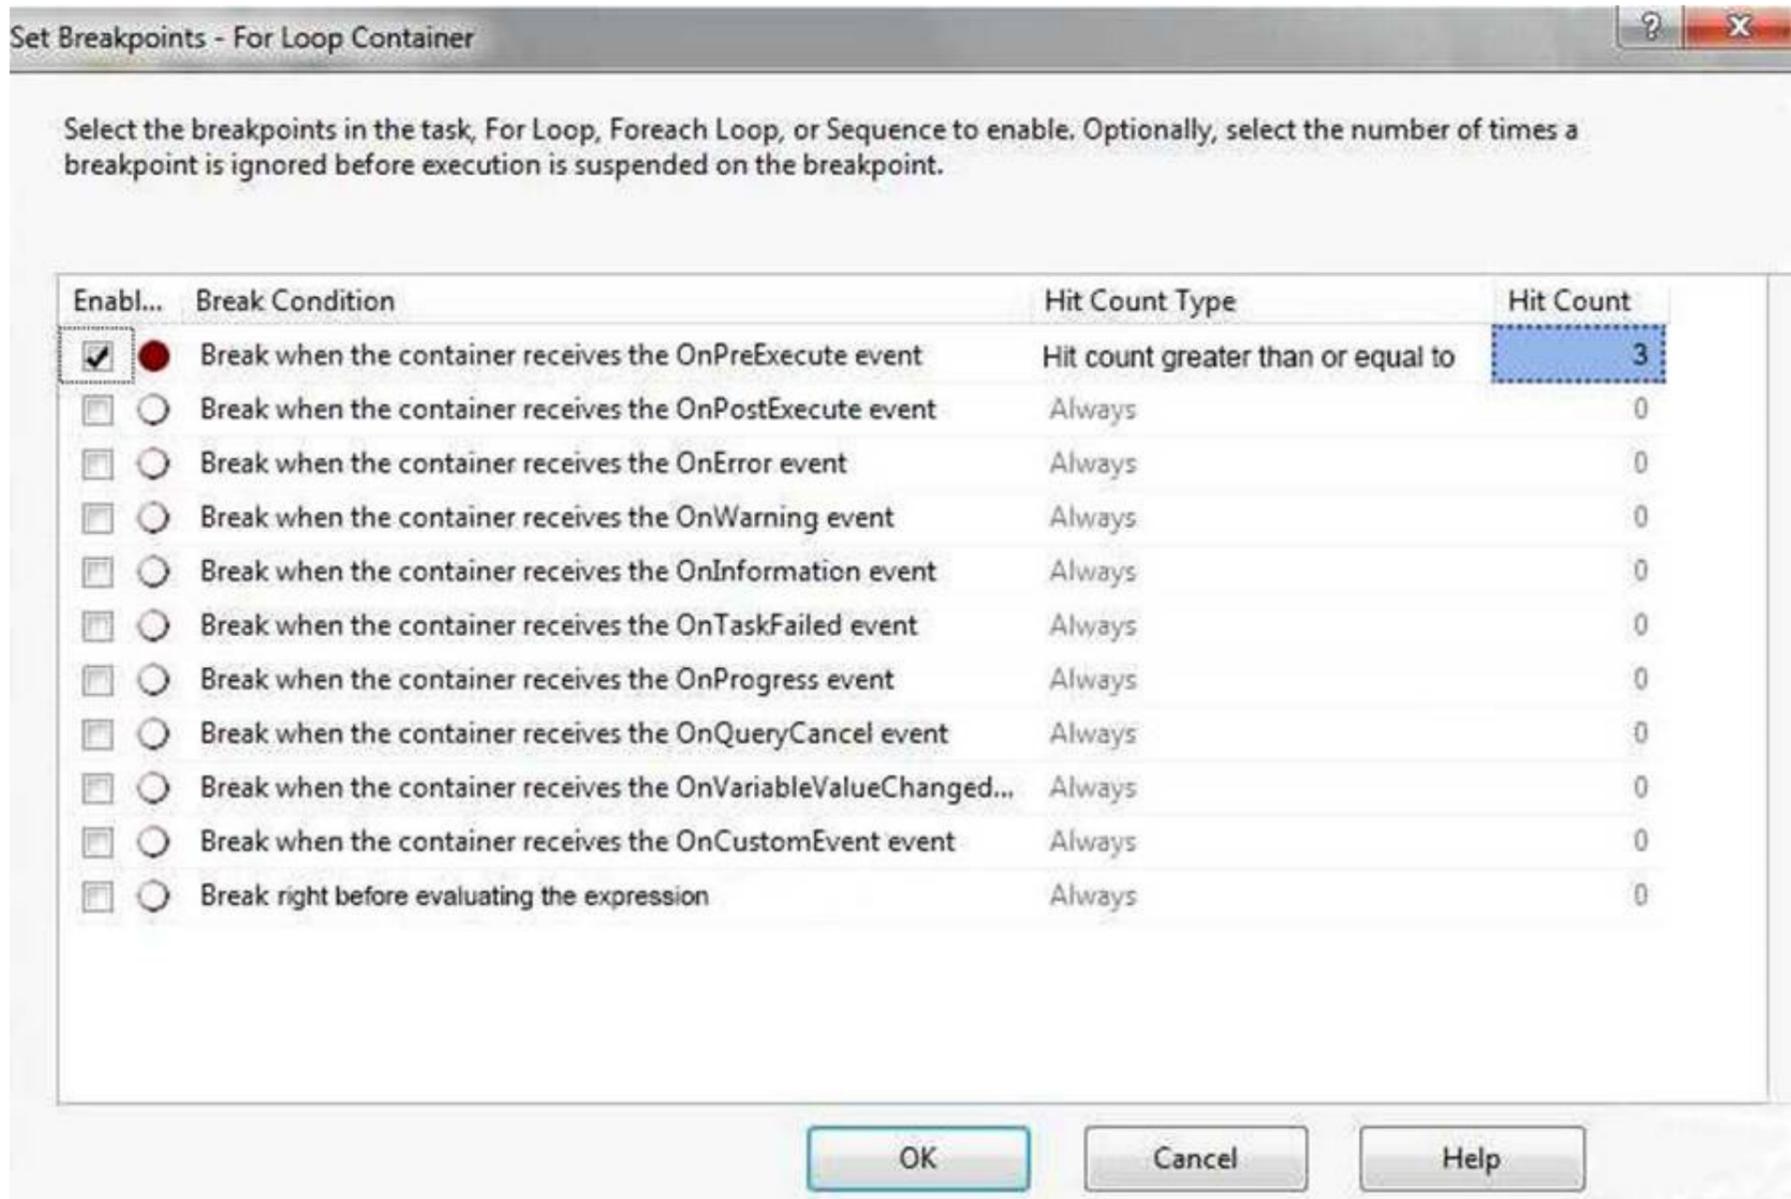
You have the Microsoft SQL Server Integration Services (SSIS) package shown in the Control flow exhibit. (Click the Exhibit button.)



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.

The 'Expression Builder' dialog box is open. The title bar says 'Expression Builder'. The main text says 'Specify the expression for the property: ConnectionString.'. On the left, there is a 'Variables' folder. On the right, there are several function categories: 'Mathematical Functions', 'String Functions', 'Date/Time Functions', 'NULL Functions', 'Type Casts', and 'Operators'. The 'Expression' field contains the text '@[User::Loop] = @[User::Loop] + 10'. Below it is an 'Evaluated value' field which is empty. At the bottom, there are three buttons: 'Evaluate Expression', 'OK', and 'Cancel'.

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.)



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.	<input type="radio"/>	<input type="radio"/>
The loop variable resets to 0 when the breakpoint is reached.	<input type="radio"/>	<input type="radio"/>
When the code stops at a breakpoint, you can change the value of the loop variable.	<input type="radio"/>	<input type="radio"/>

Answer:

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.

NEW QUESTION 167

You have a server that has Data Quality Services (DQS) installed.

You create a matching policy that contains one matching rule.

You need to configure the Similarity of Similar percentage that defines a match. Which similarity percentage will always generate a similarity score of 0?

- A. 55
- B. 80
- C. 70
- D. 75

Answer: A

Explanation: The minimum similarity between the values of a field is 60%. If the calculated matching score for a field of two records is less than 60, the similarity score is automatically set to 0.

References:

<https://docs.microsoft.com/en-us/sql/data-quality-services/create-a-matching-policy?view=sql-server-2017>

NEW QUESTION 172

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>

NEW QUESTION 176

You are developing a Microsoft SQL Server Data Warehouse. You use SQL Server Integration Services (SSIS) packages to import files from a Microsoft Azure blob storage to the data warehouse.

You plan to use multiple SQL Server instances and SSIS Scale Out to complete the workload faster. You must configure three SQL Server instances to run the SSIS package.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Install The SSIS Scale Out Worker feature on two server
- B. Install the Scale Out Master role feature on one server.
- C. Deploy the SSIS project to the SSIS catalog only on the SQL Server which has the Scale Out Master role installed.
- D. Install the SSIS Scale Out Worker feature on all three server
- E. Install the Scale Out Master role on one server.
- F. Deploy the SSIS project to the SSIS catalog on all three SQL Servers in the SSIS Scale Out environment.

Answer: AD

NEW QUESTION 180

You have a data warehouse named DW1.

InDvfe you plan to create a table named Table1 that will be partitioned by hour. Table1 will contain the last three hours of data.

You plan to implement a sliding window process for inserting data into Table1.

You need to recommend the minimum number of partitions that must be included in Table1 to support the planned implementation. The solution must minimize the number of transaction log records created during the insert process.

How many partitions should you recommend?

- A. 3
- B. 5
- C. 9
- D. 24

Answer: B

NEW QUESTION 184

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 loading data from an OLTP database to a data warehouse. The database contains a table named Sales.

Sales contains details of records that have a type of refund and records that have a type of sales. The data warehouse design contains a table for sales data and a table for refund data.

Which component should you use to load the data to the warehouse?

- A. the Slowly Changing Dimension transformation
- B. the Conditional Split transformation
- C. the Merge transformation
- D. the Data Conversion transformation
- E. an Execute SQL task
- F. the Aggregate transformation
- G. the Lookup transformation

Answer: B

Explanation: The Conditional Split transformation can route data rows to different outputs depending on the content of the data. The implementation of the Conditional Split transformation is similar to a CASE decision structure in a programming language. The transformation evaluates expressions, and based on the results, directs the data row to the specified output. This transformation also provides a default output, so that if a row matches no expression it is directed to the default output.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/conditionalsplit-Transformation>

NEW QUESTION 187

You have a fact table in a data warehouse that stores financial data. The table contains eight column configured as shown in the following table.

DateID	Stock-ID	Open-ingPrice	Closing-Price	Quanti-tyTraded	Bro-kerID	Num-berOfTra-des	Market-ID
20170301	22	30.20	34.23	100	10	1	1
20170301	31	10.05	12.23	110	10	2	2
20170302	22	30.89	34.76	899	5	1	1

You need to identify a column that can be aggregated across all dimensions. Which column should you identify?

- A. OpeningPrice
- B. StockID
- C. NumberOfTrades
- D. MarketID

Answer: C

Explanation: Aggregates are sometimes referred to as pre-calculated summary data, since aggregations are usually precomputed, partially summarized data, that are stored in new aggregated tables.

References: [https://en.wikipedia.org/wiki/Aggregate_\(data_warehouse\)](https://en.wikipedia.org/wiki/Aggregate_(data_warehouse))

NEW QUESTION 190

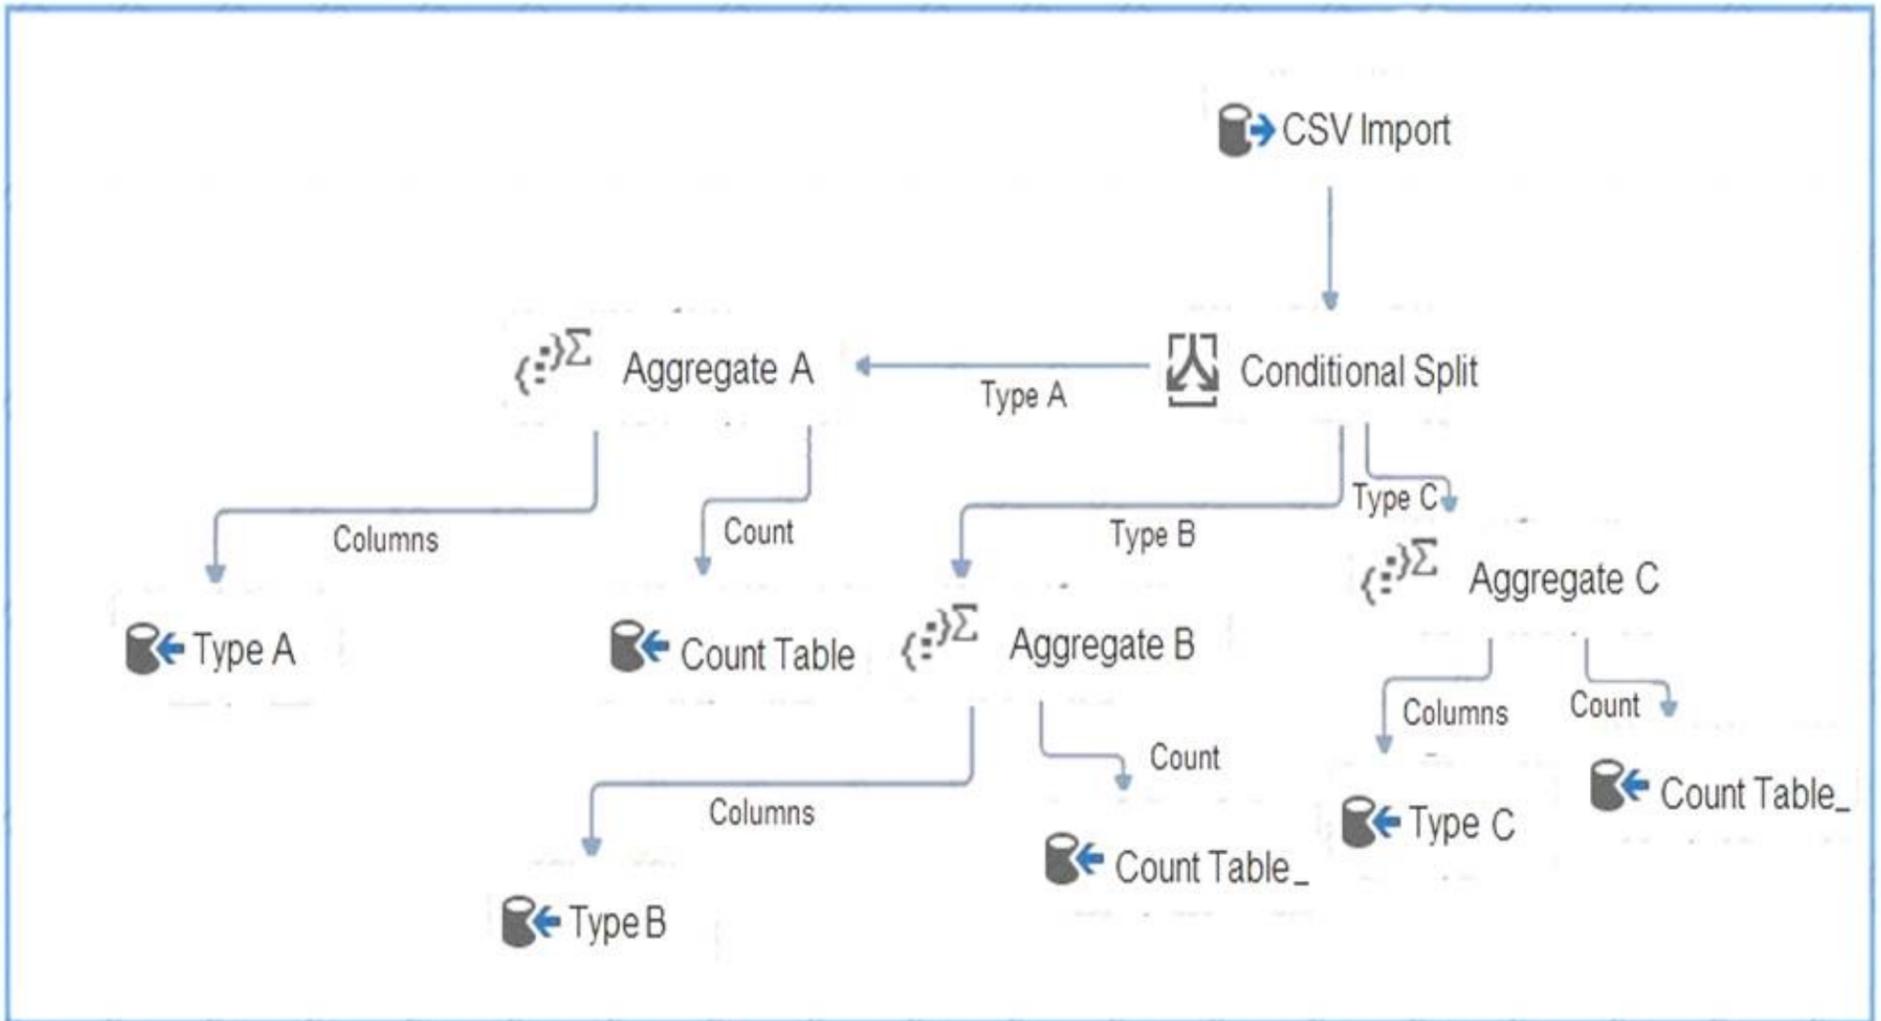
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Each night you receive a comma separated values (CSV) file that contains different types of rows. Each row type has a different structure. Each row in the CSV file is unique. The first column in every row is named Type. This column identifies the data type.

For each data type, you need to load data from the CSV file to a target table. A separate table must contain the number of rows loaded for each data type.

Solution: You create a SQL Server Integration Services (SSIS) package as shown in the exhibit. (Click the Exhibit tab.)



Does the solution meet the goal?

- A. Yes
- B. NO

Answer: A

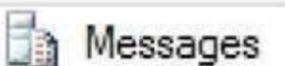
Explanation: The conditional split is correctly placed before the count.

NEW QUESTION 193

You manage a data warehouse in a Microsoft SQL Server instance. Company employee information is imported from the human resources system to a table named Employee in the data warehouse instance. The Employee table was created by running the query shown in the Employee Schema exhibit. (Click the Exhibit button.)

```
CREATE TABLE dbo.DimEmployee
(
    EmployeeID int IDENTITY (1,1) PRIMARY KEY,
    EmployeeSSN int NULL UNIQUE,
    EmployeeName nvarchar(100) NOT NULL
)
```

The personal identification number is stored in a column named EmployeeSSN. All values in the EmployeeSSN column must be unique. When importing employee data, you receive the error message shown in the SQL Error exhibit. (Click the Exhibit button.)



Msg 2627, Level 14, State 1, Line 13

Violation of UNIQUE Key constraint 'UQ_DimEmplo_8549FE539cf2eca'. Cannot insert duplicate key object 'dbo.DimEmployee'. The duplicate key value is (<NULL>).

The statement has been terminated.

You determine that the Transact-SQL statement shown in the Data Load exhibit in the cause of the error. (Click the Exhibit button.)

```
INSERT dbo.DimEmployee (EmployeeSSN, EmployeeName)
SELECT NULL, EmployeeName
FROM HR.dbo.Employee
```

You remove the constraint on the EmployeeSSN column. You need to ensure that values in the EmployeeSSN column are unique. 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
Creating a clustered unique index on the EmployeeSSN column solves the issue.	<input type="radio"/>	<input type="radio"/>
Creating a filtered unique index on the EmployeeSSN column solves the issue.	<input type="radio"/>	<input type="radio"/>

Answer:

Explanation: With the ANSI standards SQL:92, SQL:1999 and SQL:2003, an UNIQUE constraint must disallow duplicate non-NULL values but accept multiple NULL values. In the Microsoft world of SQL Server however, a single NULL is allowed but multiple NULLs are not. From SQL Server 2008, you can define a unique filtered index based on a predicate that excludes NULLs. References: <https://stackoverflow.com/questions/767657/how-do-i-create-a-unique-constraint-that-also-allows-nulls>

NEW QUESTION 195

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 developing a Microsoft SQL Server Integration Services (SSIS) package. The package design consists of two differently structured sources in a single data flow. The Sales source retrieves sales transactions from a SQL Server database, and the Product source retrieves product details from an XML file. You need to combine the two data flow sources into a single output dataset. Which SSIS Toolbox item should you use?

- A. CDC Control task
- B. CDC Splitter
- C. Union All
- D. XML task
- E. Fuzzy Grouping
- F. Merge
- G. Merge Join

Answer: G

Explanation: The Merge Join transformation provides an output that is generated by joining two sorted datasets using a FULL, LEFT, or INNER join. For example, you can use a LEFT join to join a table that includes product information with a table that lists the country/region in which a product was manufactured. The result is a table that lists all products and their country/region of origin. References: <https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/merge-join-transformation>

NEW QUESTION 198

You are developing a Microsoft SQL Server Integration Services (SSIS) package. You create a data flow that has the following characteristics:

- The package moves data from the table [source].Table1 to DW.Table1.
- All rows from [source].Table1 must be captured in DW.Table1 for error.Table1.
- The table error.Table1 must accept rows that fail upon insertion into DW.Table1 due to violation of nullability or data type errors such as an invalid date, or invalid characters in a number.
- The behavior for the Error Output on the "OLE DB Destination" object is Redirect.
- The data types for all columns in [source].Table1 are VARCHAR. Null values are allowed.
- The Data access mode for both OLE DB destinations is set to Table or view - fast load.

The table definitions are as follows:

```
CREATE TABLE [source].Table1
```

```
(
  ID INT NULL,
  CreateDate VARCHAR(100) NULL,
  Date1 DATETIME2(7) NULL,
  Number1 VARCHAR(100) NULL
)
```

```
CREATE TABLE error.Table1
```

```
(
  ID INT NULL,
  CreateDate VARCHAR(100) NULL,
  Date1 DATETIME2(7) NULL,
  Number1 VARCHAR(100) NULL,
  ErrorDescription VARCHAR(255) NULL
)
```

Use the drop-down menus to select the answer choice that answers each question.

The ErrorDescription column is not yet populated in error.Table1. You must capture the error description for any rows redirected to the "Error OLE DB Destination". What should you do next?

- In "OLE DB Destination Error", map the ErrorCode field to ErrorDescription.
- Create an INSERT trigger on [Error].[Table1] to populate the ErrorDescription from ErrorCode.
- Add a Derived Column transformation before "OLE DB Destination". Use ErrorCode to populate ErrorDescription.
- Add a Script Component transformation before "OLE DB Destination Error". Capture the ErrorDescription with VB or C# code.

You execute the package. You note that all rows are redirected to OLE DB Destination Error, including both rows with bad data and rows with valid data. What is the next step?

- Uncheck the Check Constraints option in OLE DB Destination.
- Change the Data access mode for OLE DB Destination to Table or View.
- Uncheck the options Table Lock and Check Constraints for OLE DB Destination.
- Change the ValidateExternalMetadata setting for the OLE DB Destination Error object to False.
- Add a Conditional Split transformation before OLE DB Destination. Create outputs based on ErrorCode.

Answer:

Explanation:

The ErrorDescription column is not yet populated in error.Table1. You must capture the error description for any rows redirected to the "Error OLE DB Destination". What should you do next?

- In "OLE DB Destination Error", map the ErrorCode field to ErrorDescription.
- Create an INSERT trigger on [Error].[Table1] to populate the ErrorDescription from ErrorCode.
- Add a Derived Column transformation before "OLE DB Destination". Use ErrorCode to populate ErrorDescription.
- Add a Script Component transformation before "OLE DB Destination Error". Capture the ErrorDescription with VB or C# code.

You execute the package. You note that all rows are redirected to OLE DB Destination Error, including both rows with bad data and rows with valid data. What is the next step?

- Uncheck the Check Constraints option in OLE DB Destination.
- Change the Data access mode for OLE DB Destination to Table or View.
- Uncheck the options Table Lock and Check Constraints for OLE DB Destination.
- Change the ValidateExternalMetadata setting for the OLE DB Destination Error object to False.
- Add a Conditional Split transformation before OLE DB Destination. Create outputs based on ErrorCode.

NEW QUESTION 202

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.

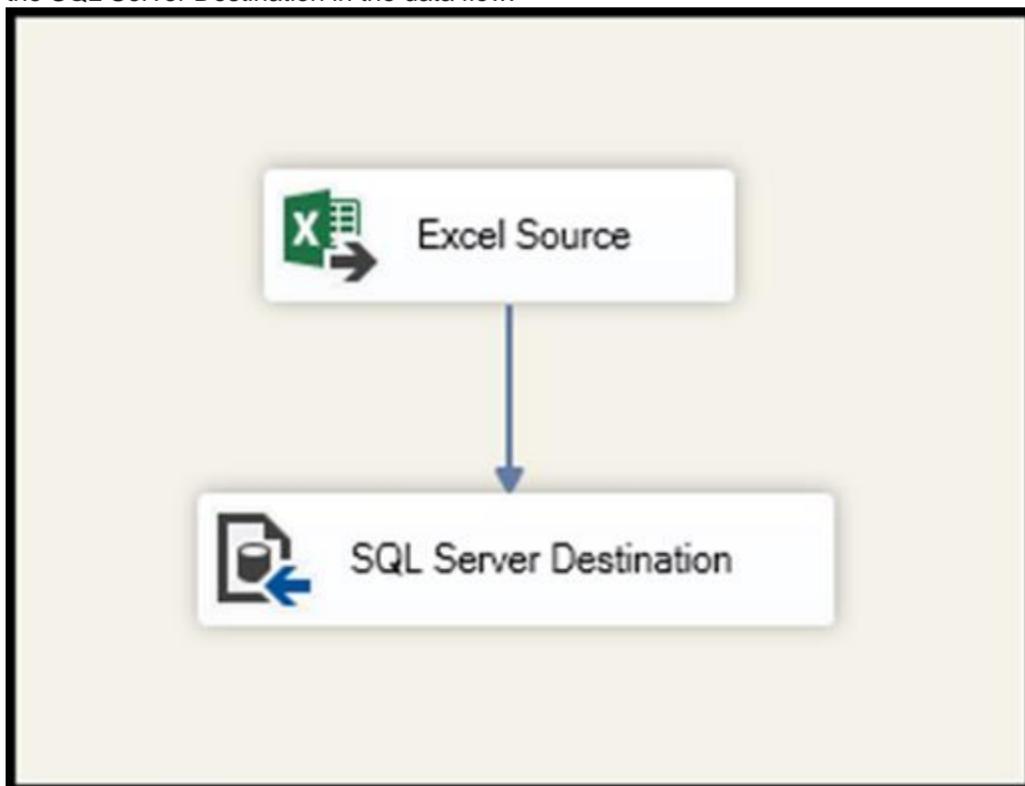
Database	Description
DB1	This database supports the online store.
DB2	This is the data warehouse for the company. DB2 contains a table named OnlineOrder that is partitioned in hourly increments. The LOCK_ESCALATION option is set to AUTO . The data flow contains 24 OLE DB destinations, one for each partition.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

Answer: A

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>

NEW QUESTION 207

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. 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.
 References: <https://www.timitchell.net/post/2016/01/18/getting-started-with-change-tracking-in-sql-server/>

NEW QUESTION 211

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 the series.

Start of repeated scenario

Contoso. Ltd. has a Microsoft SQL Server environment that includes SQL Server Integration Services (SSIS), a data warehouse, and SQL Server Analysis Services (SSAS) Tabular and multi-dimensional models.

The data warehouse stores data related to your company sales, financial transactions and financial budgets. All data for the data warehouse originates from the company's business financial system.

The data warehouse includes the following tables:

Table	Notes
dbo.load_City	
dbo.stage_City	
dbo.dim_City	
fact.Sale	
fact.Transaction	This table contains more than 20,000,000 rows. There are currently no indexes on the table. The table has a column named [sale key]. Most queries that target fact.Transaction return recent data based on this column and a column named Description.

The company plans to use Microsoft Azure to store older records from the data warehouse. You must modify the database to enable the Stretch Database capability.

Users report that they are becoming confused about which city table to use for various queries. You plan to create a new schema named Dimension and change the name of the dbo.dia_city table to Dimension.city. Data loss is not permissible, and you must not leave traces of the old table in the data warehouse.

You must implement a partitioning scheme for the fact.Transaction table to move older data to less expensive storage. Each partition will store data for a single calendar year, as shown in the exhibit (Click the Exhibit button.) You must align the partitions.

You must improve performance for queries against the fact.Transaction table. You must implement appropriate indexes and enable the Stretch Database capability.

End of repeated scenario

You need to configure the fact. Transaction table.

Which three Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

Results Messages

	Transaction Key	Date Key	Customer Key	Bill To Customer Key	Supplier Key	Transaction Type Key	Payment Method Key	WWI Invoice ID
1	7	2013-01-01	375	202	0	1	0	7
2	11	2013-01-01	387	202	0	1	0	11
3	12	2013-01-01	330	202	0	1	0	12
4	13	2013-01-01	274	202	0	1	0	13
5	16	2013-01-01	215	202	0	1	0	16
6	25	2013-01-01	298	202	0	1	0	25
7	26	2013-01-01	285	202	0	1	0	26
8	30	2013-01-01	368	202	0	1	0	30
9	35	2013-01-01	232	202	0	1	0	35
10	39	2013-01-01	346	202	0	1	0	39
11	41	2013-01-01	216	202	0	1	0	41
12	63	2013-01-02	224	202	0	1	0	42
13	64	2013-01-02	264	202	0	1	0	43
14	65	2013-01-02	268	202	0	1	0	44
15	70	2013-01-02	375	202	0	1	0	49
16	74	2013-01-02	387	202	0	1	0	53
17	75	2013-01-02	330	202	0	1	0	54
16	74	2013-01-02	387	202	0	1	0	53
17	75	2013-01-02	330	202	0	1	0	54
18	76	2013-01-02	274	202	0	1	0	55
19	78	2013-01-02	215	202	0	1	0	57
20	85	2013-01-02	298	202	0	1	0	64
21	86	2013-01-02	285	202	0	1	0	65
22	90	2013-01-02	368	202	0	1	0	69
23	94	2013-01-02	232	202	0	1	0	73

Transact-SQL segments

```
ALTER DATABASE Contoso SET REMOTE_DATA_ARCHIVE
= ON (
SERVER =
'MyStretchDatabaseServer.database.windows.net',
CREDENTIAL = TestAzure
)
GO
```

```
CREATE FUNCTION dbo.fn_stretch_by_date(@date
DATETIME2)
RETURNS TABLE
WITH SCHEMABINDING
AS
RETURN SELECT 1 AS is_eligible WHERE @date <
CONVERT(datetime2, '1/1/2015', 101)
GO
```

```
ALTER TABLE fact.Transaction
SET(REMOTE_DATA_ARCHIVE = ON (
FILTER_PREDICATE = dbo.fn_stretch_by_date
([Date Key]), MIGRATION_STATE = OUTBOUND
))
GO
```

```
ALTER DATABASE master SET REMOTE_DATA_ARCHIVE
= ON (
SERVER =
'MyStretchDatabaseServer.database.windows.net',
CREDENTIAL = TestAzure
)
GO
```

Answer area



Answer:

Explanation:

```

Transact-SQL statements
ALTER DATABASE Contoso SET REMOTE_DATA_ARCHIVE
= ON (
SERVER =
'MyStretchDatabaseServer.database.windows.net',
CREDENTIAL = TestAzure
)
GO

CREATE FUNCTION dbo.fn_stretch_by_date(@date
DATETIME2)
RETURNS TABLE
WITH SCHEMABINDING
AS
RETURN SELECT 1 AS is_eligible WHERE @date <
CONVERT(datetime2, '1/1/2015', 101)
GO

ALTER TABLE fact.Transaction
SET(REMOTE_DATA_ARCHIVE = ON (
FILTER_PREDICATE = dbo.fn_stretch_by_date
([Date Key]), MIGRATION_STATE = OUTBOUND
))
GO

ALTER DATABASE master SET REMOTE_DATA_ARCHIVE
= ON (
SERVER =
'MyStretchDatabaseServer.database.windows.net',
CREDENTIAL = TestAzure
)
GO
    
```

Answer area

```

CREATE FUNCTION dbo.fn_stretch_by_date(@date
DATETIME2)
RETURNS TABLE
WITH SCHEMABINDING
AS
RETURN SELECT 1 AS is_eligible WHERE @date <
CONVERT(datetime2, '1/1/2015', 101)
GO
    
```

```

ALTER DATABASE master SET REMOTE_DATA_ARCHIVE
= ON (
SERVER =
'MyStretchDatabaseServer.database.windows.net',
CREDENTIAL = TestAzure
)
GO
    
```

```

ALTER TABLE fact.Transaction
SET(REMOTE_DATA_ARCHIVE = ON (
FILTER_PREDICATE = dbo.fn_stretch_by_date
([Date Key]), MIGRATION_STATE = OUTBOUND
))
GO
    
```

Navigation arrows: > < < > < >

NEW QUESTION 215

You administer a Microsoft SQL Server Master Data Services (MDS) model. All model entity members have passed validation. The current model version should be committed to form a record of master data that can be audited and create a new version to allow the ongoing management of the master data. You lock the current version. You need to manage the model versions. 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

- Commit the current version.
- Set the new version status to **Open**.
- Unlock the current version.
- Unlock the new version.
- Set the current version status to **Open**.
- Create a copy of the current version.
- Validate the current version.

Answer Area

Navigation arrows: < > < >

Answer:

Explanation: Box 1: Validate the current version.

In Master Data Services, validate a version to apply business rules to all members in the model version. You can validate a version after it has been locked.

Box 2: Commit the current version.

In Master Data Services, commit a version of a model to prevent changes to the model's members and their attributes. Committed versions cannot be unlocked.

Prerequisites:

Box 3: Create a copy of the current version.

In Master Data Services, copy a version of the model to create a new version of it. Note:

References:

NEW QUESTION 218

You deploy a Microsoft Server database that contains a staging table named EmailAddress_Import. Each night, a bulk process will import customer information from an external database, cleanse the data, and then insert it into the EmailAddress table. Both tables contain a column named EmailAddressValue that stores the email address.

You need to implement the logic to meet the following requirements:

- ▶ Email addresses that are present in the EmailAddress_Import table but not in the EmailAddress table must be inserted into the EmailAddress table.
- ▶ Email addresses that are not in the EmailAddress_Import but are present in the EmailAddress table must be deleted from the EmailAddress table.

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment 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.

Transact-SQL segments

- EmailAddress
- EmailAddress_Import
- NOT MATCHED BY SOURCE
- NOT MATCHED BY TARGET
- MATCHED

Answer area

```

MERGE Transact-SQL segment AS B
USING Transact-SQL segment AS A
ON A.EmailAddressValue = B.EmailAddressValue
WHEN Transact-SQL segment
THEN INSERT (EmailAddressValue) VALUES (A.EmailAddressValue)
WHEN Transact-SQL segment
THEN DELETE
    
```

Answer:

Explanation: Box 1: EmailAddress

The EmailAddress table is the target. Box 2: EmailAddress_import

The EmailAddress_import table is the source. Box 3: NOT MATCHED BY TARGET

Box 4: NOT MATCHED BY SOURCE

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/merge-transact-sql>

NEW QUESTION 219

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>

NEW QUESTION 223

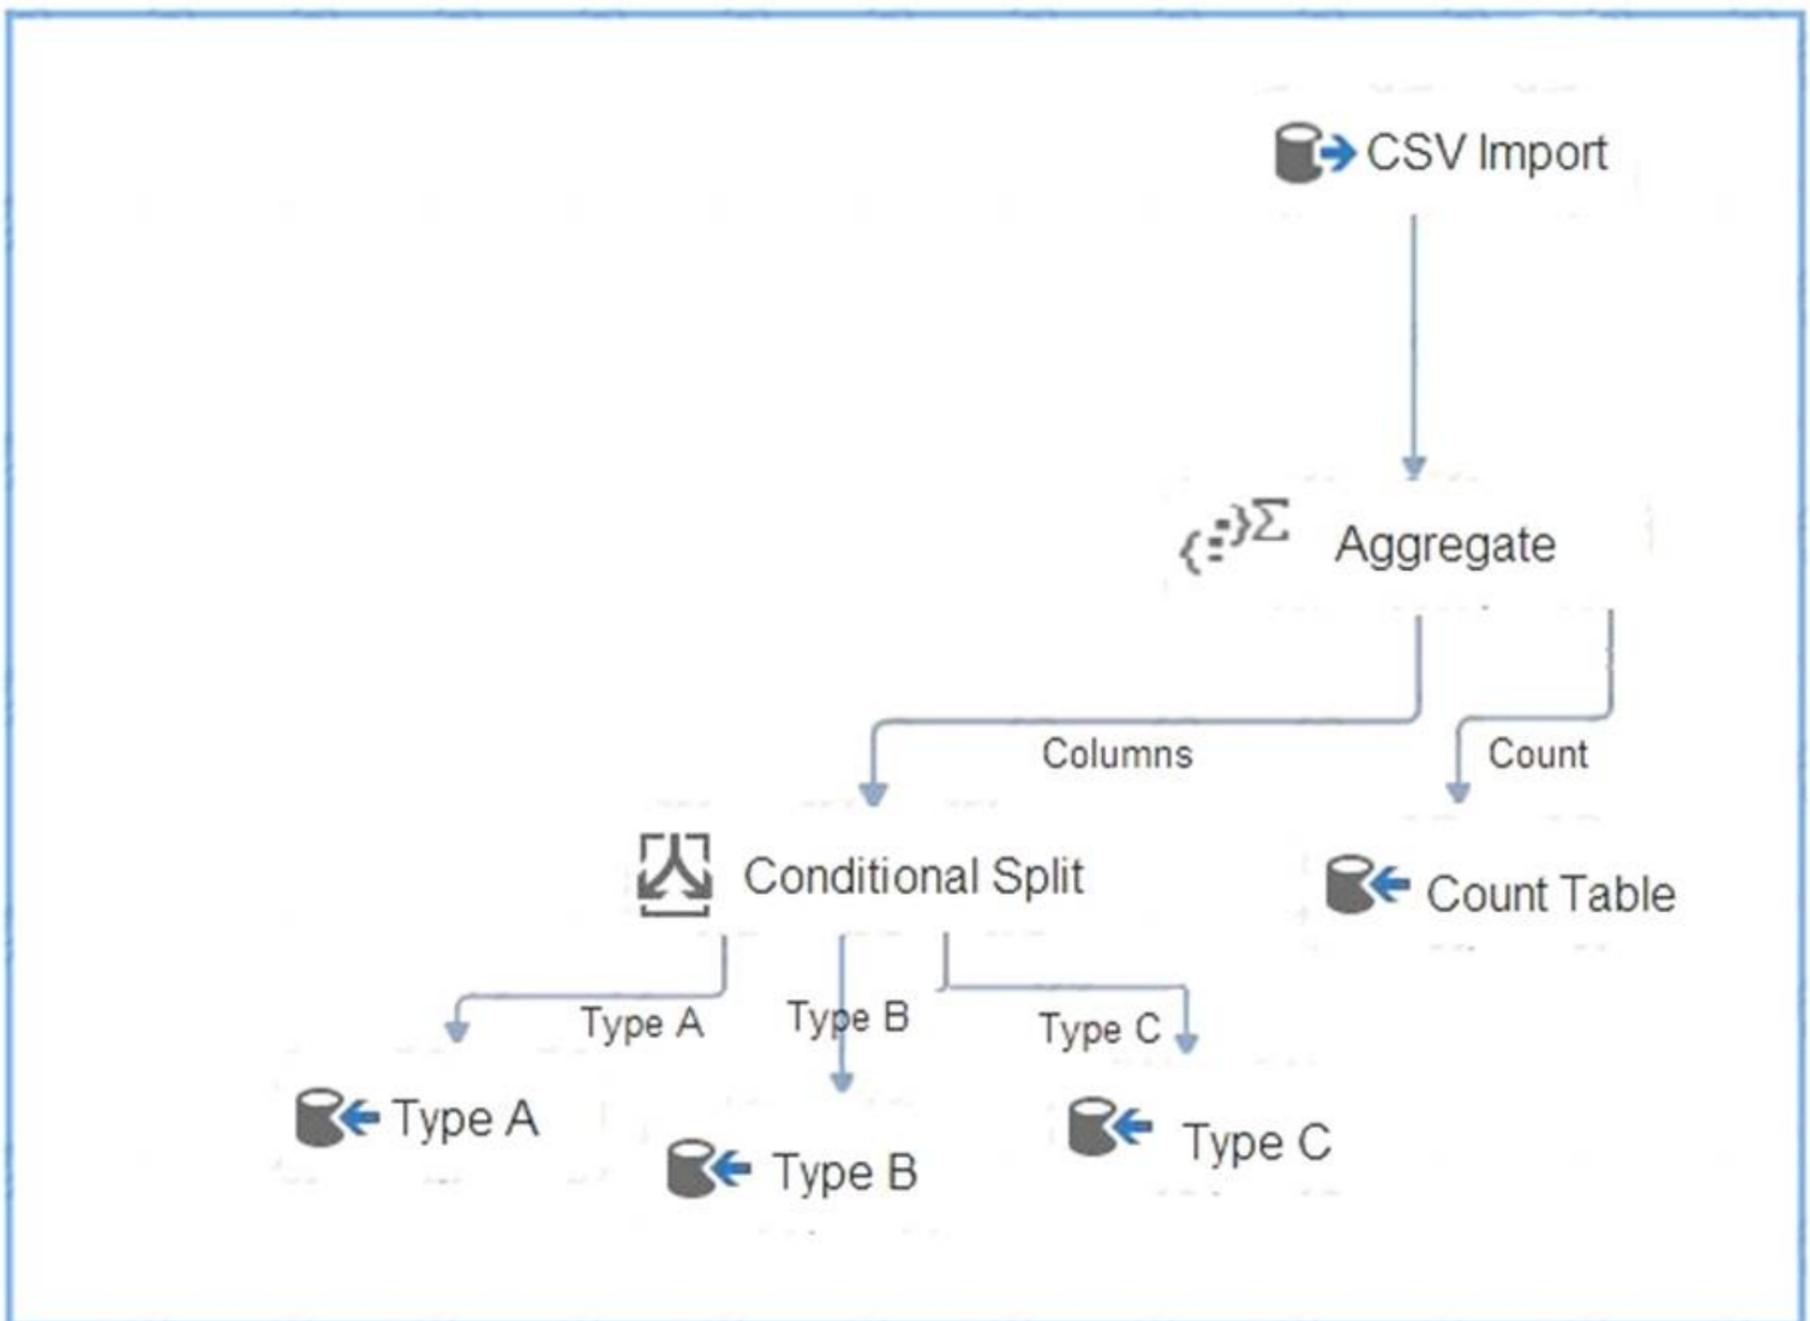
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

Each night you receive a comma separated values (CSV) file that contains different types of rows. Each row type has a different structure. Each row in the CSV file is unique. The first column in every row is named Type. This column identifies the data type.

For each data type, you need to load data from the CSV file to a target table. A separate table must contain the number of rows loaded for each data type.

Solution: You create a SQL Server Integration Services (SSIS) package as shown in the exhibit. (Click the Exhibit tab.)



Does the solution meet the goal?

- A. Yes
- B. NO

Answer: B

Explanation: The conditional split must be before the count.

NEW QUESTION 227

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 have a database named DB1 that has change data capture enabled.

A Microsoft SQL Server Integration Services (SSIS) job runs once weekly. The job loads changes from DB1 to a data warehouse by querying the change data capture tables. You discover that the job loads changes from the previous three days only. You need re ensure that the job loads changes from the previous week. Which stored procedure should you execute?

- A. catalog.deploy_project
- B. catalog.restore_project
- C. catalog.stop.operation
- D. sys.sp_cdc.addJob
- E. sys.sp.cdc.changejob
- F. sys.sp_cdc_disable_db
- G. sys.sp_cdc_enable_db
- H. sys.sp_cdc.stopJob

Answer: A

Explanation: catalog.deploy_project deploys a project to a folder in the Integration Services catalog or updates an existing project that has been deployed previously.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/system-stored-procedures/catalog-deploy-project-ssisd>

NEW QUESTION 228

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

Answer: H

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>

NEW QUESTION 233

You are developing a Microsoft SQL Server Master Data Services (MDS) solution.

The model contains an entity named Product. The Product entity has three user-defined attributes named Category, Subcategory, and Price, respectively.

You need to ensure that combinations of values stored in the Category and Subcategory attributes are unique. What should you do?

- A. Create an attribute group that consists of the Category and Subcategory attribute
- B. Publish a business rule for the attribute group.
- C. Publish a business rule that will be used by the Product entity.
- D. Create a derived hierarchy based on the Category and Subcategory attribute
- E. Use the Category attribute as the top level for the hierarchy.
- F. Set the value of the Attribute Type property for the Category and Subcategory attributes to Domainbased.

Answer: B

Explanation: In Master Data Services, business rule actions are the consequence of business rule condition evaluations. If a condition is true, the action is initiated.

The Validation action "must be unique": The selected attribute must be unique independently or in combination with defined attributes.

NEW QUESTION 238

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 the series.

Start of repeated scenario

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 to daily. The FactOrder 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.

- 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 FactOrder 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.

End of repeated scenario

You need to optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.

Which technology should you use for each table?

To answer, select the appropriate technologies in the answer area.

Answer area

Table	Technology
Dimension.SalesTerritory	<input type="text"/>
Dimension.Customer	<input type="text"/>
Dimension.Date	<input type="text"/>

Table	Technology
Dimension.SalesTerritory	<input type="text"/> <ul style="list-style-type: none"> Change Data Capture (CDC) Change Tracking Temporal table Microsoft SQL Server snapshot replication
Dimension.Customer	<input type="text"/> <ul style="list-style-type: none"> Change Data Capture (CDC) Change Tracking Temporal table Microsoft SQL Server snapshot replication
Dimension.Date	<input type="text"/> <ul style="list-style-type: none"> Change Data Capture (CDC) Change Tracking Temporal table Microsoft SQL Server snapshot replication

Answer:

Explanation: Box 1: Temporal table Box 2: Temporal table

Compared to CDC, Temporal tables are more efficient in storing historical data as it ignores insert actions. Box 3: Change Data Capture (CDC)

By using change data capture, you can track changes that have occurred over time to your table. This kind of functionality is useful for applications, like a data warehouse load process that need to identify changes, so they can correctly apply updates to track historical changes over time.

CDC is good for maintaining slowly changing dimensions.

Scenario: Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.

The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. References:

<https://www.mssqltips.com/sqlservertip/5212/sql-server-temporal-tables-vs-change-data-capture-vs-change-trac> <https://docs.microsoft.com/en-us/sql/relational-databases/tables/temporal-table-usage-scenarios?view=sql-server>

NEW QUESTION 241

You need to recommend a storage solution for a data warehouse that minimizes load times. The solution must provide availability if a hard disk fails.

Which RAID configuration should you recommend for each type of database file? To answer, drag the appropriate RAID configurations to the correct database file types. Each RAID configuration 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: Each correct selection is worth one point.

RAID Configurations

RAID 0	RAID 5
RAID 6	RAID 10

Answer Area

Data file:	RAID configuration
Transaction log:	RAID configuration

Answer:

Explanation: Box 1: RAID 5

RAID 5 is similar to that of RAID 0 provided that the number of disks is the same. However, due to the fact that it is useless to read the parity data, the read speed is just (N-1) times faster but not N times as in RAID 0.

Box 2: RAID 10

Always place log files on RAID 1+0 (or RAID 1) disks. This provides better protection from hardware failure, and better write performance.

Note: In general RAID 1+0 will provide better throughput for write-intensive applications. The amount of performance gained will vary based on the HW vendor's RAID implementations. Most common alternative to RAID 1+0 is RAID 5. Generally, RAID 1+0 provides better write performance than any other RAID level providing data protection, including RAID 5.

NEW QUESTION 245

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-767 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-767 Product From:

<https://www.2passeasy.com/dumps/70-767/>

Money Back Guarantee

70-767 Practice Exam Features:

- * 70-767 Questions and Answers Updated Frequently
- * 70-767 Practice Questions Verified by Expert Senior Certified Staff
- * 70-767 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 70-767 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year