# Microsoft

## Exam Questions 70-767

Implementing a SQL Data Warehouse (beta)

**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 hapen to existing queries that are running?

| |
|---|
| The data warehouse instance pauses when existing queries have completed. No new queries are permitted. |
| The existing queries will be immediately terminated. |
| The existing queries will be paused untill the data warehouse instance is resumed. |

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

| |
|---|
| You will stop being charged for compute resources but will continue to be charged for storage. |
| You will continue to be charged for both compute resources and storage. |
| You are no longer charged for storage but continue to pay for the assigned data warehouse instance units. |

A. Mastered
B. Not Mastered

**Answer:** A

**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**
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 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 statements:

```
SELECT CustomerKey, SUM(SalesAmt) TotalSales
FROM sales.FactOrders
GROUP BY CustomerKey
OPTION (LABEL = 'TotalSales')

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

Does the solution meet the goal?

A. Yes

B. No

**Answer:** B

**Explanation:**
We must use Label, not QueryID in the WHERE clause. References:
https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec

**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.)

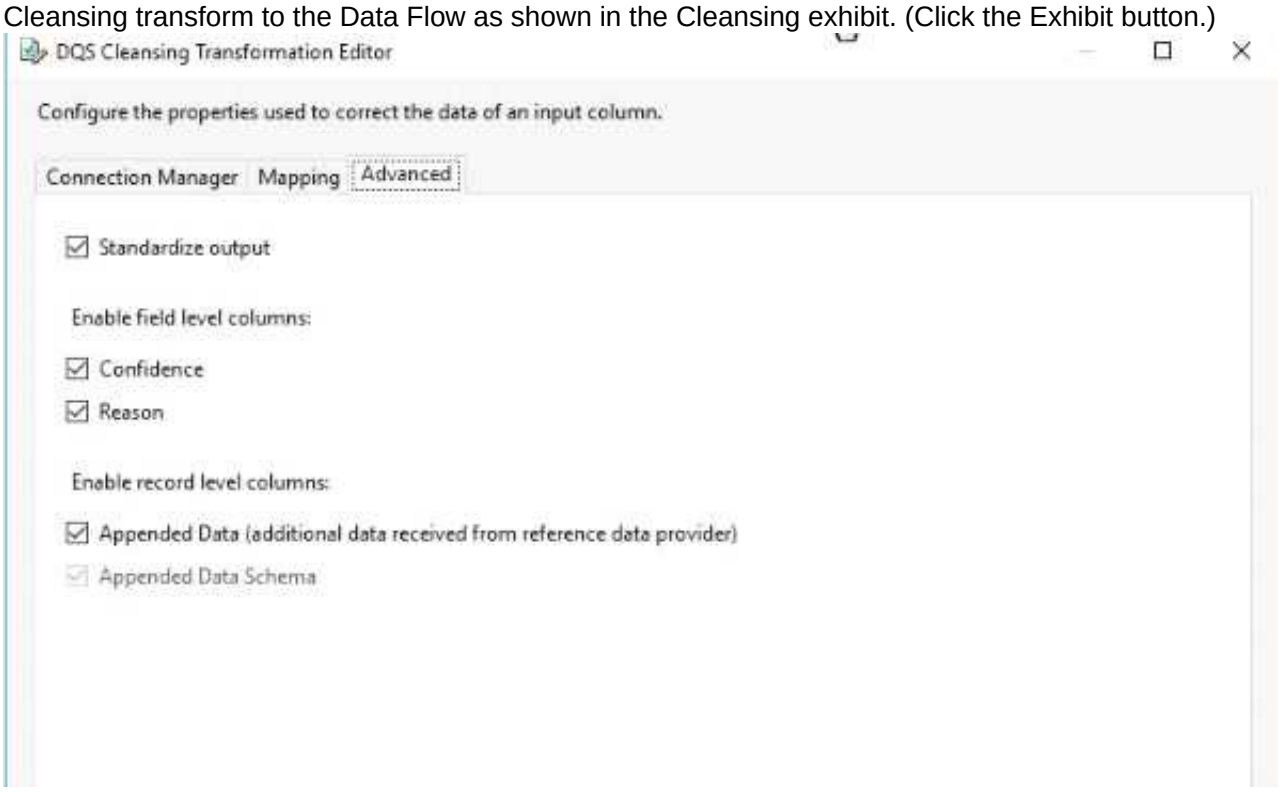


You create a Conditional Split transform as shown in the Splitter exhibit. (Click the Exhibit button.)

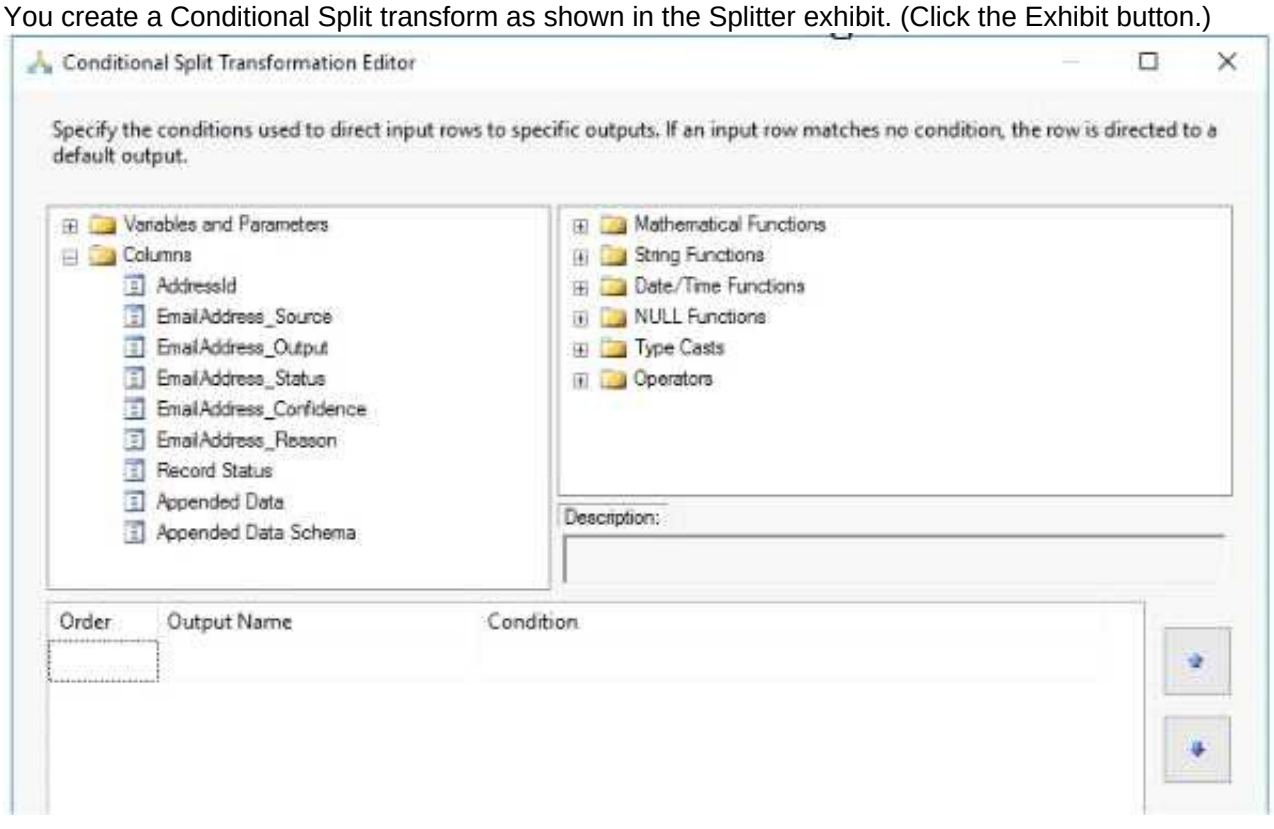


You need to split the output of the DQ5 Cleansing task to obtain only Correct values from the EmailAddress column. For each of the following statements, select Yes if the statement is true. Otherwise, select No.

**Answer area**

| | Yes | No |
|---|---|---|
| You can use the EmailAddress_Output column to split the output. | ○ | ○ |
| You can use the EmailAddress_Status column to split the output. | ○ | ○ |

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

| | Yes | No |
|---|---|---|
| You can use the EmailAddress_Output column to split the output. | ○ | ● |
| You can use the EmailAddress_Status column to split the output. | ● | ○ |

**NEW QUESTION 4**
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 | • Index1 contains all the data in Table1.<br>• Queries against Index1 perform aggregation operations against hundreds of millions of rows. |
| Index2 | • Index2 returns all the columns in this index.<br>• Index2 contains 80 percent of the columns in Table1.<br>• 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: Index type
Index2: Index type

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

**Index Types**

Clustered | Clustered columnstore | Hash | Heap | Nonclustered | Nonclustered columnstore

**Answer Area**

Index1: Clustered columnstore
Index2: Nonclustered columnstore

**NEW QUESTION 5**
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.

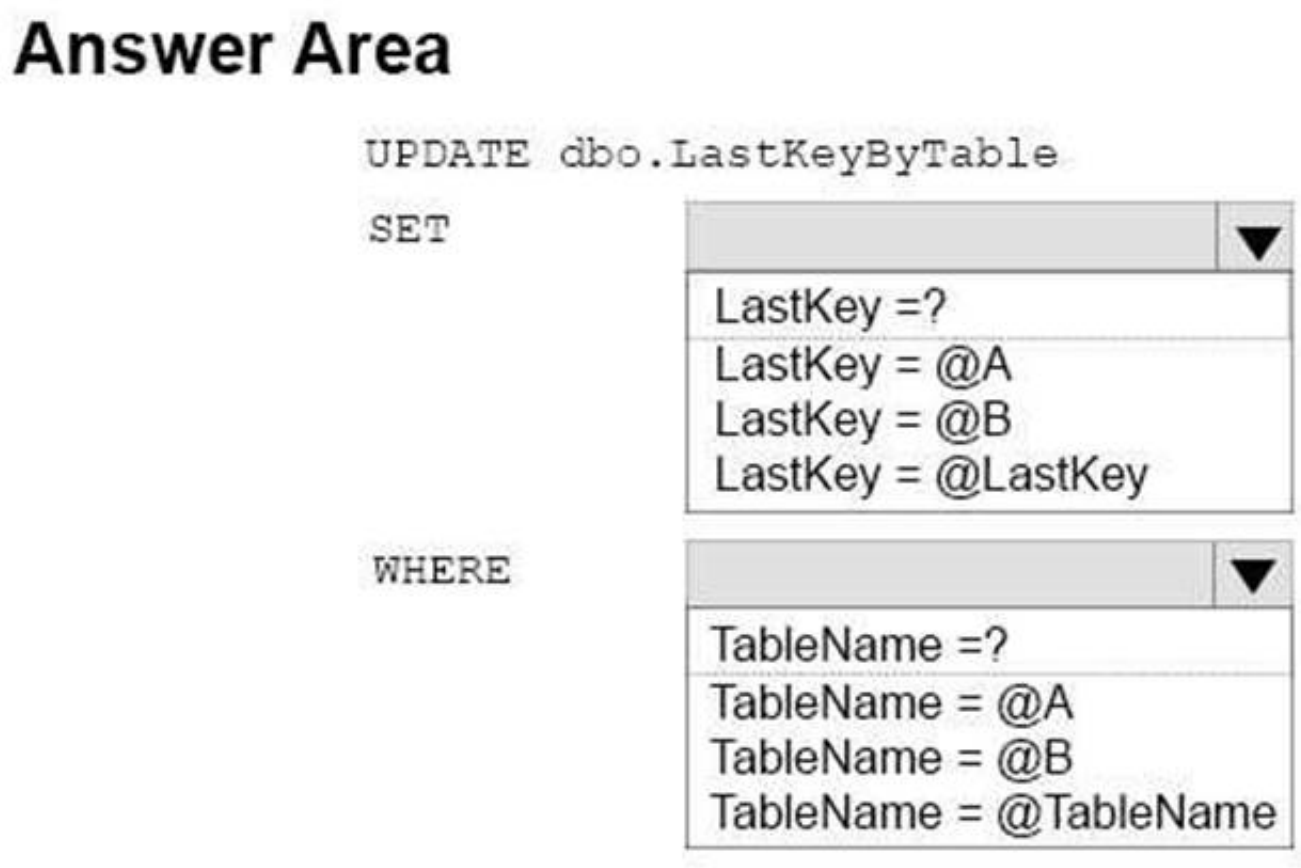


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

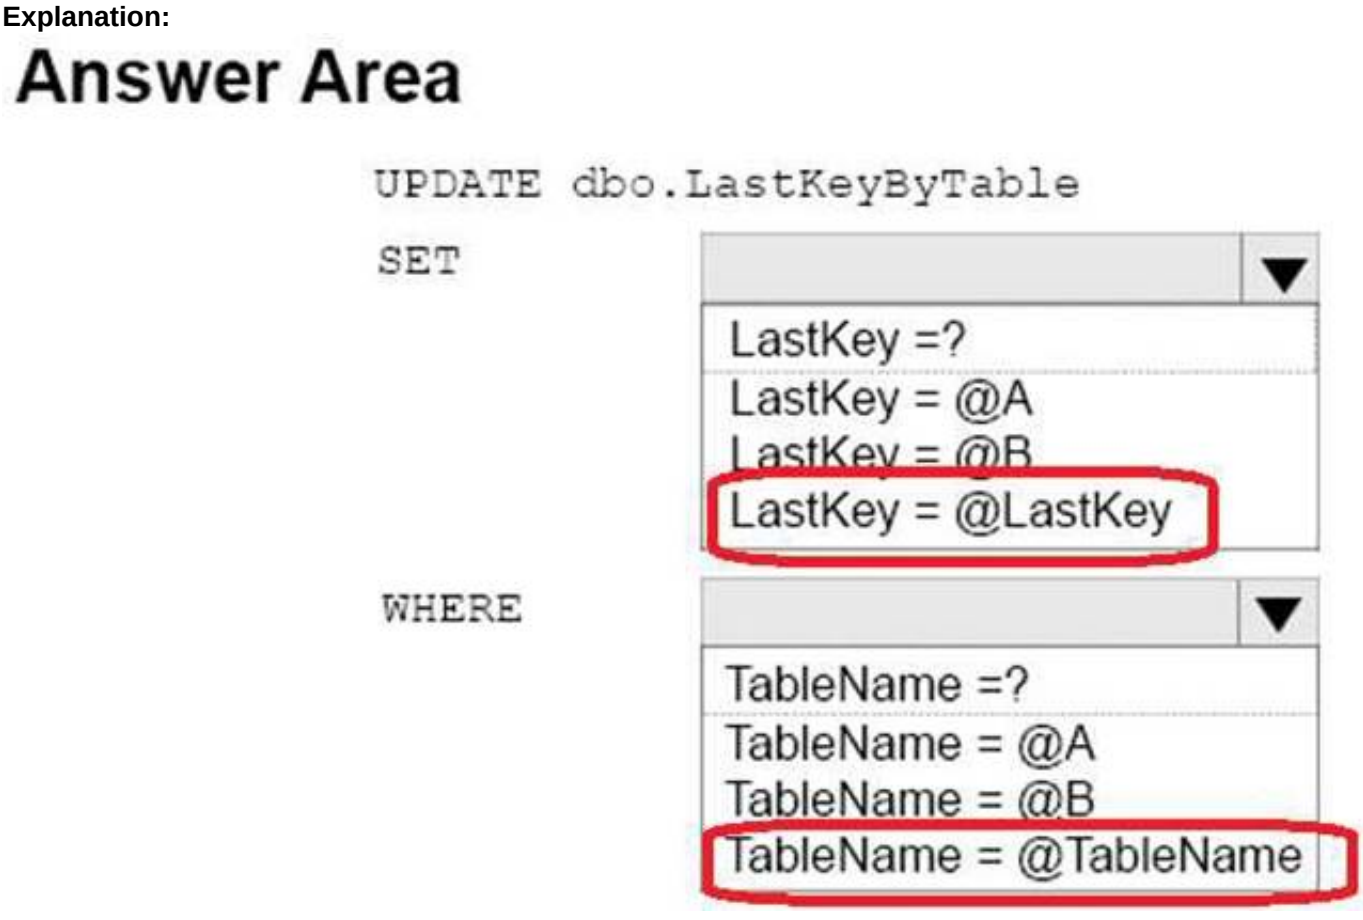


**NEW QUESTION 6**

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 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: In the Customer model, add a domain-based attribute to reference the CountryRegion entity in the Geography model.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** A

**NEW QUESTION 7**
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 in Master Data Services (MDS) as shown in the following exhibit.

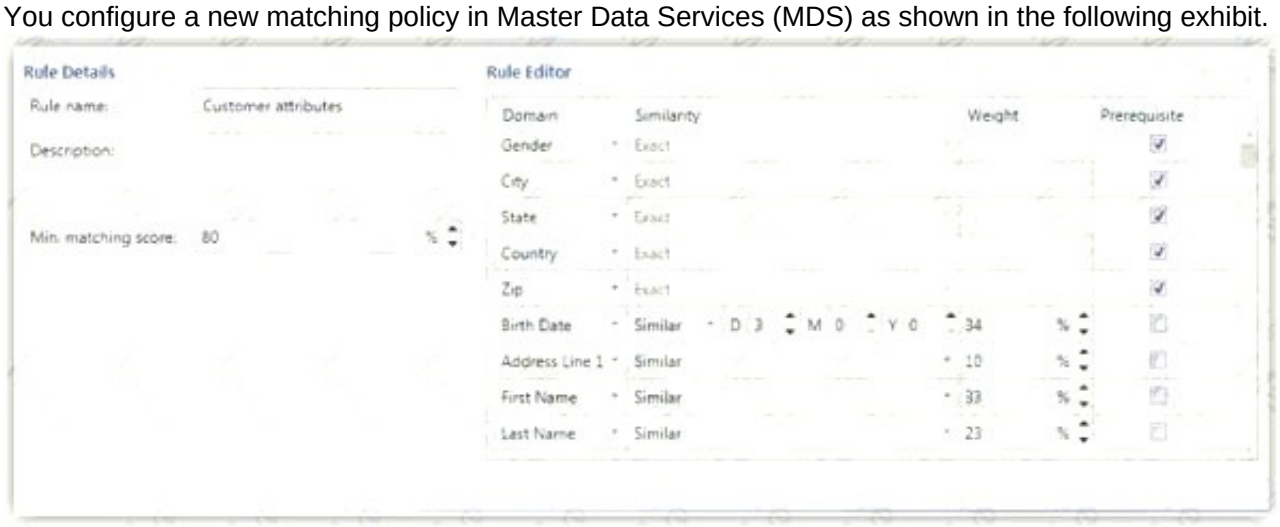


You review the Matching Results of the policy and find that the number of new values matches the new values.
You verify that the data contains multiple records that have similar address values, and you expect some of the records to match. You need to increase the likelihood that the records will match when they have similar address values.
Solution: You decrease the relative weights for Address Line 1 of the matching policy. Does this meet the goal?

A. Yes
B. No

**Answer:** A

**NEW QUESTION 8**
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 0 PERCENT
```

Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**Explanation:**
Microsoft recommend against specifying 0 PERCENT or 0 ROWS in a CREATE STATISTICS..WITH SAMPLE statement. When 0 PERCENT or ROWS is specified, the statistics object is created but does not contain statistics data.
References: https://docs.microsoft.com/en-us/sql/t-sql/statements/create-statistics-transact-sql

**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:

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

You are not permitted to make changes to the client applications. You need to optimize data loading for the Dimension.Customer 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.
NOTE: You will not need all of the Transact-SQL segments.

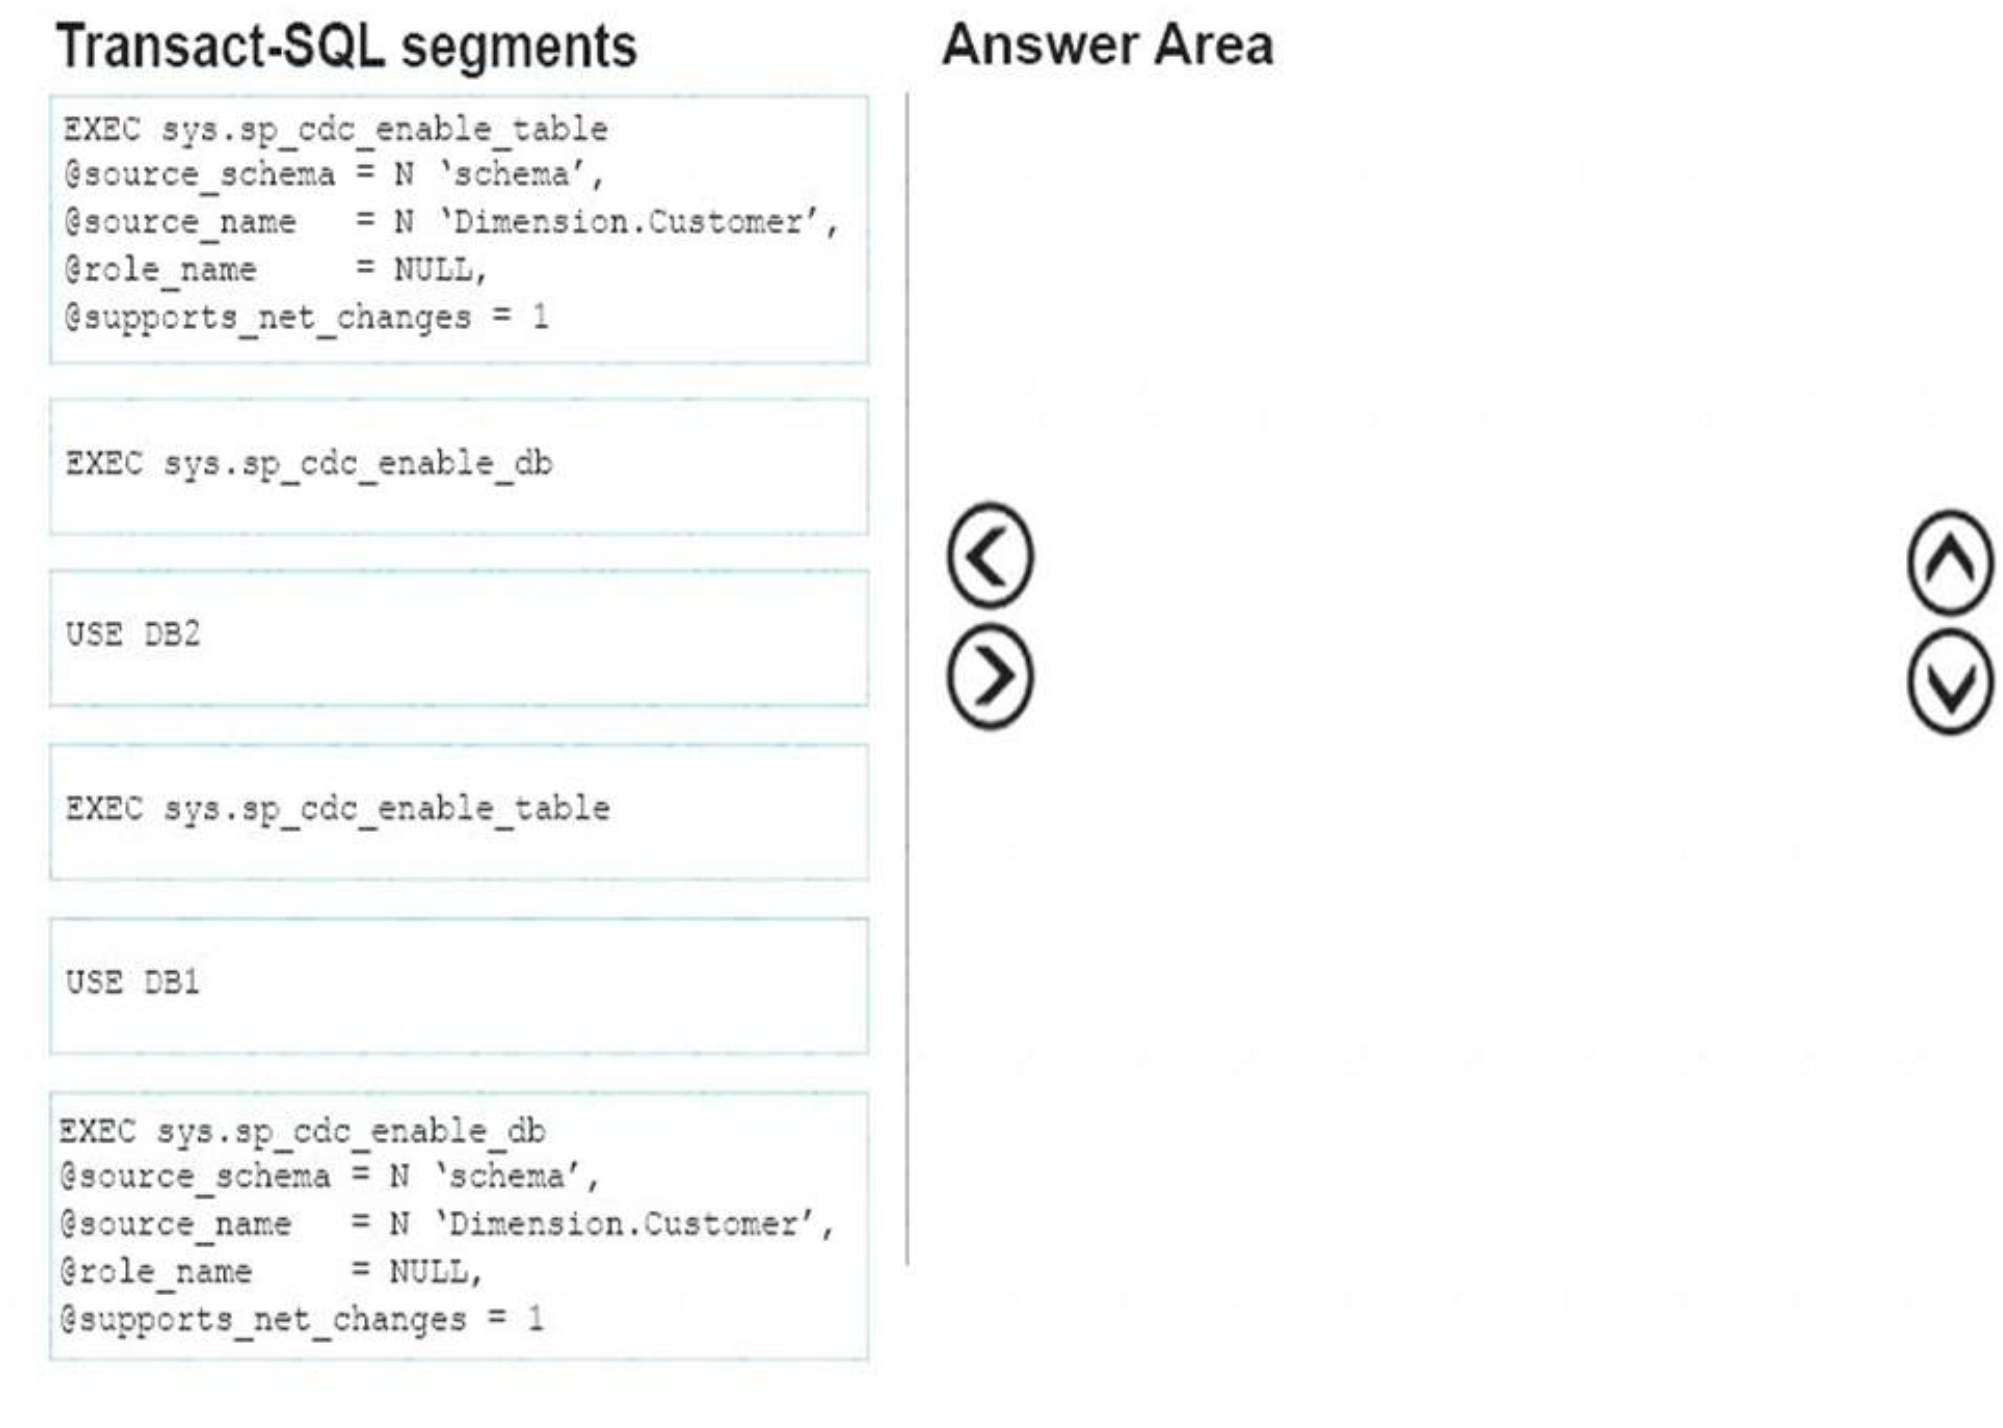


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Step 1: USE DB1
From Scenario: 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.
Step 2: EXEC sys.sp_cdc_enable_db
Before you can enable a table for change data capture, the database must be enabled. To enable the database, use the sys.sp_cdc_enable_db stored procedure. sys.sp_cdc_enable_db has no parameters. Step 3: EXEC sys.sp_cdc_enable_table
@source schema = N 'schema' etc.
Sys.sp_cdc_enable_table enables change data capture for the specified source table in the current database. Partial syntax:
sys.sp_cdc_enable_table
[ @source_schema = ] 'source_schema',
[ @source_name = ] 'source_name' , [,[ @capture_instance = ] 'capture_instance' ] [,[ @supports_net_changes = ] supports_net_changes ]
Etc.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-enable-table-tran
https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-enable-db-transa

**NEW QUESTION 10**
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 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 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 13**
You need to build a knowledge base in Data Quality Services (DQS).
You need to ensure that the data is validated by using a third-party data source before DQS processes the data. 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.

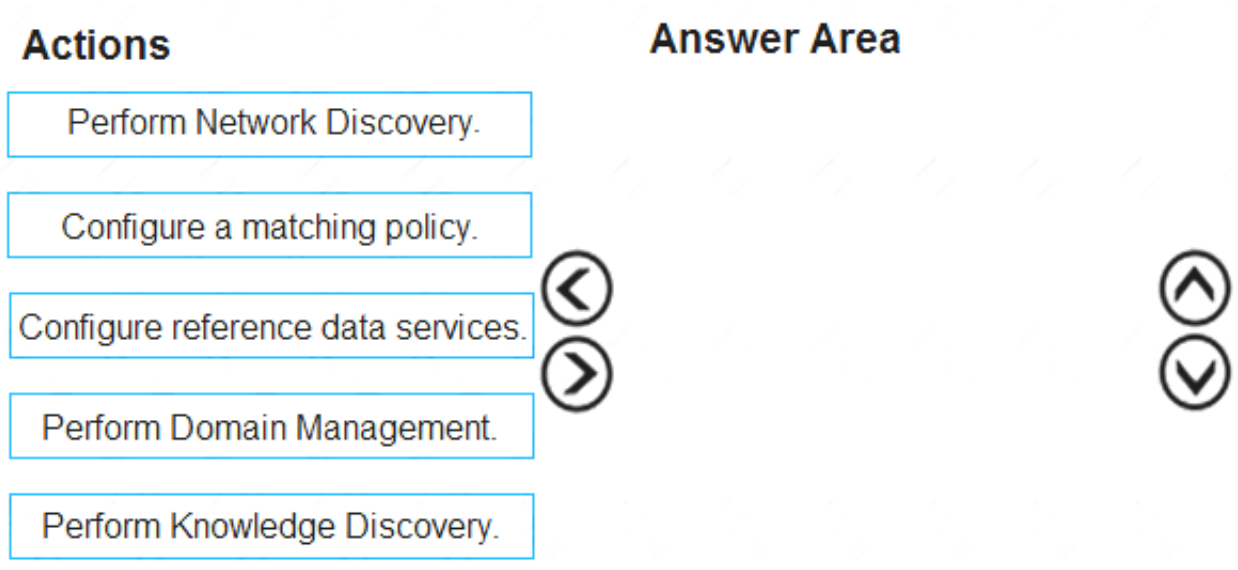


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Building a DQS knowledge base involves the following processes and components: Step 1: Perform Knowledge Discovery
A computer-assisted process that builds knowledge into a knowledge base by processing a data sample Step 2: Perform Domain Management
An interactive process that enables the data steward to verify and modify the knowledge that is in knowledge base domains, each of which is associated with a data field. This can include setting field-wide properties, creating rules, changing specific values, using reference data services, or setting up term-based or cross-field relationships.
Step 3: Configure reference Data Services
A process of domain management that enables you to validate your data against data maintained and guaranteed by a reference data provider.
Step 4: Configure a Matching Policy
A policy that defines how DQS processes records to identify potential duplicates and non-matches, built into the knowledge base in a computer-assisted and interactive process.
References: https://docs.microsoft.com/en-us/sql/data-quality-services/dqs-knowledge-bases-and-domains

**NEW QUESTION 18**
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.

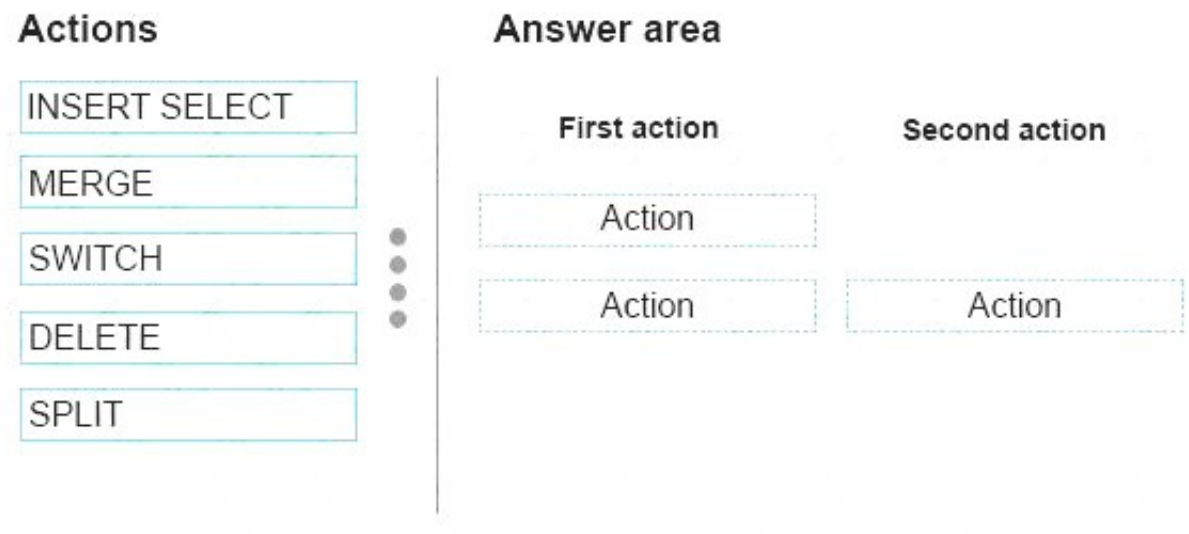

A. Mastered
B. Not Mastered

**Answer:** A

**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 22**
You are designing a warehouse named DW1.
A table named Table1 is partitioned by using the following partitioning scheme and function.

```
AS RANGE LEFT FOR VALUES ('20150101', '20160101', '20170101', '20180101', '20190101'
'20200101');
GO
CREATE PARTITION SCHEME schema1
AS PARTITION function1
ALL TO ([primary]);
GO
```

```
CREATE TABLE table1
(MyId BIGINT IDENTITY (1,1),
OrderDate datetime,
DueDate datetime,
AccountNumber nvarchar(15)
…
PRIMARY KEY (MyId, OrderDate))
ON schema1 (OrderDate)
GO
```

Reports are generated from the data in Table1.
You need to ensure that queries to DW1 return results as quickly as possible. Which column should appear in the WHERE statement clause of the query?

A. AccountNumber
B. MyId
C. DueDate
D. OrderDate

**Answer:** D

**NEW QUESTION 27**
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 configuring a Microsoft SQL server named dw1for a new data warehouse. The server contains eight drives and eight processor cores. Each drive uses a separate physical disk.
You need to configure storage for the tempdb database. The solution must minimize the amount of time it takes to process daily ETLjobs.
Solution: You configure eight files for the tempdb database. You place the files on a drive that will NOT store the user database files.
Does this meet the goal?

A. Yes
B. No

**Answer:** B

**NEW QUESTION 32**
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.
You are importing data from databases at retail stores into a central data warehouse. All stores use the same database schema.
The query being executed against the retail stores is shown below:

```
SELECT *
FROM dbo.Sales
WHERE SalesDate >= CAST(date, GETDATE() -1)
ORDER BY ID
```

The data source property named IsSorted is set to True. The output of the transform must be sorted.
You need to add a component to the data flow. 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 33**
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:

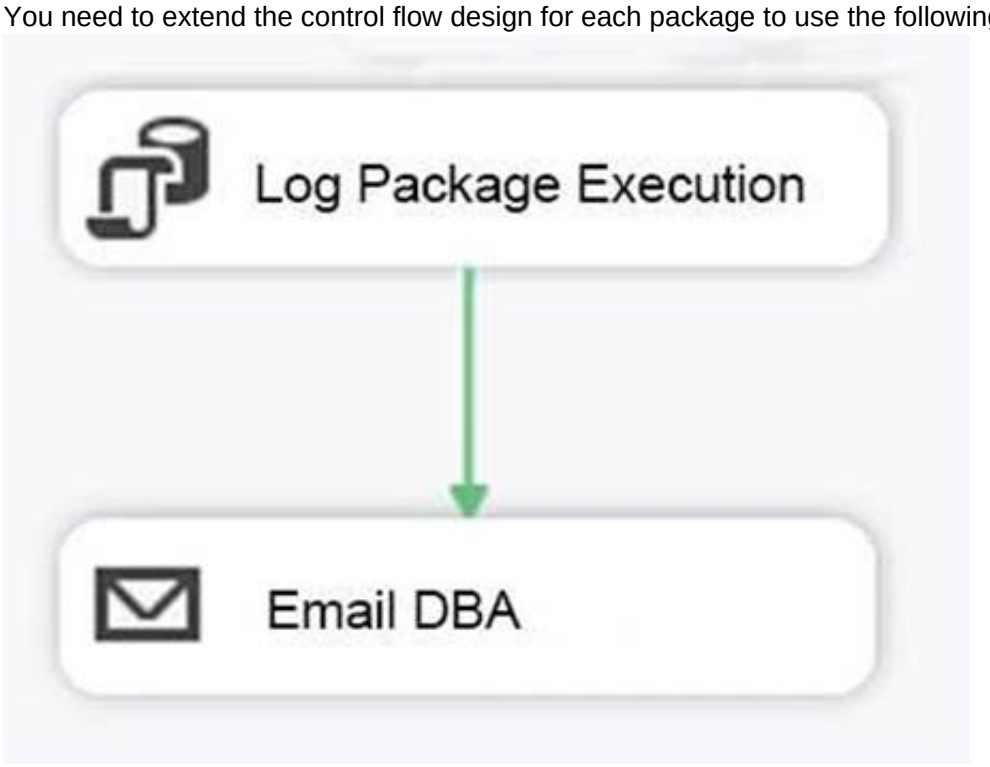


Solution: You add the control flow to a control flow package part. You add an instance of the control flow package part to each data warehouse load package.
Does the solution meet the goal?

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 38**
You manage Master Data Services (MDS). You plan to create entities and attributes and load them with the data. You also plan to match data before loading it into Data Quality Services (DQS).
You need to recommend a solution to perform the actions.
What should you recommend?

A. MDS Add-in for Microsoft Excel
B. MDS Configuration Manager
C. Data Quality Matching
D. MDS repository

**Answer:** A

**Explanation:**
In the Master Data Services Add-in for Excel, matching functionality is provided by Data Quality Services (DQS). This functionality must be enabled to be used.
To enable Data Quality Services integration
Open Master Data Services Configuration Manager.
In the left pane, click Web Configuration.
On the Web Configuration page, select the website and web application.
In the Enable DQS Integration section, click Enable integration with Data Quality Services.
On the confirmation dialog box, click OK.
References:
https://docs.microsoft.com/en-us/sql/master-data-services/install-windows/enable-data-quality-services-integrati

**NEW QUESTION 40**
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 43**
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.)

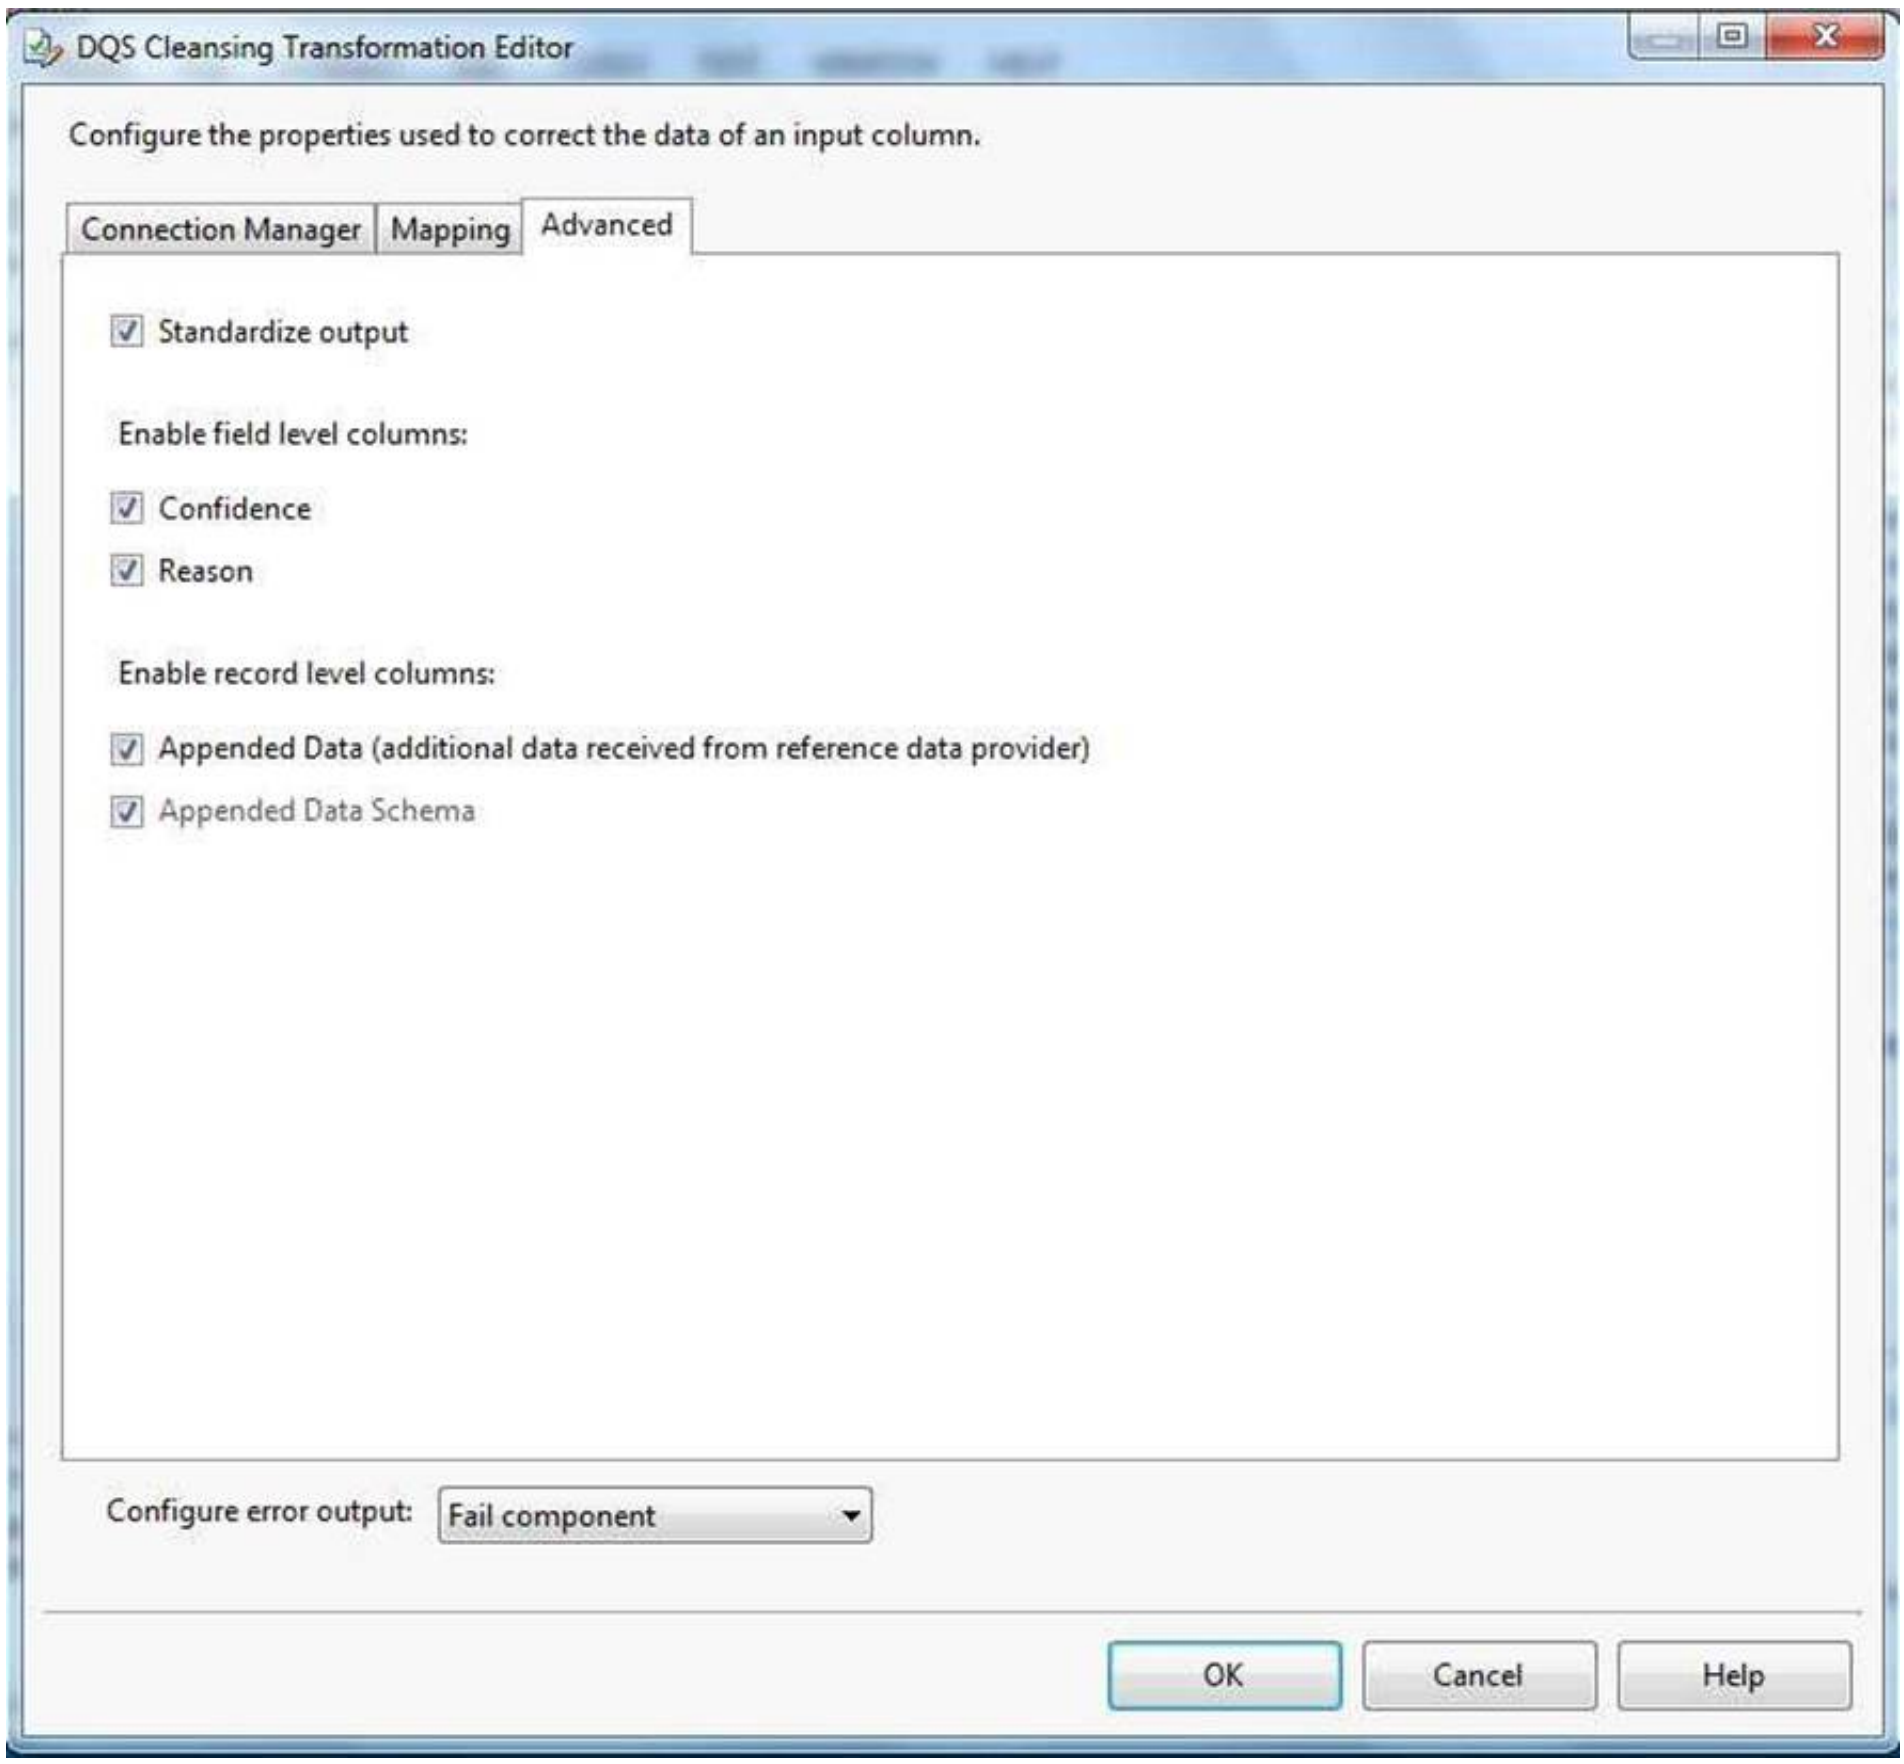


You create a Conditional Split transform as shown in the Splitter exhibit. (Click the Exhibit button.)

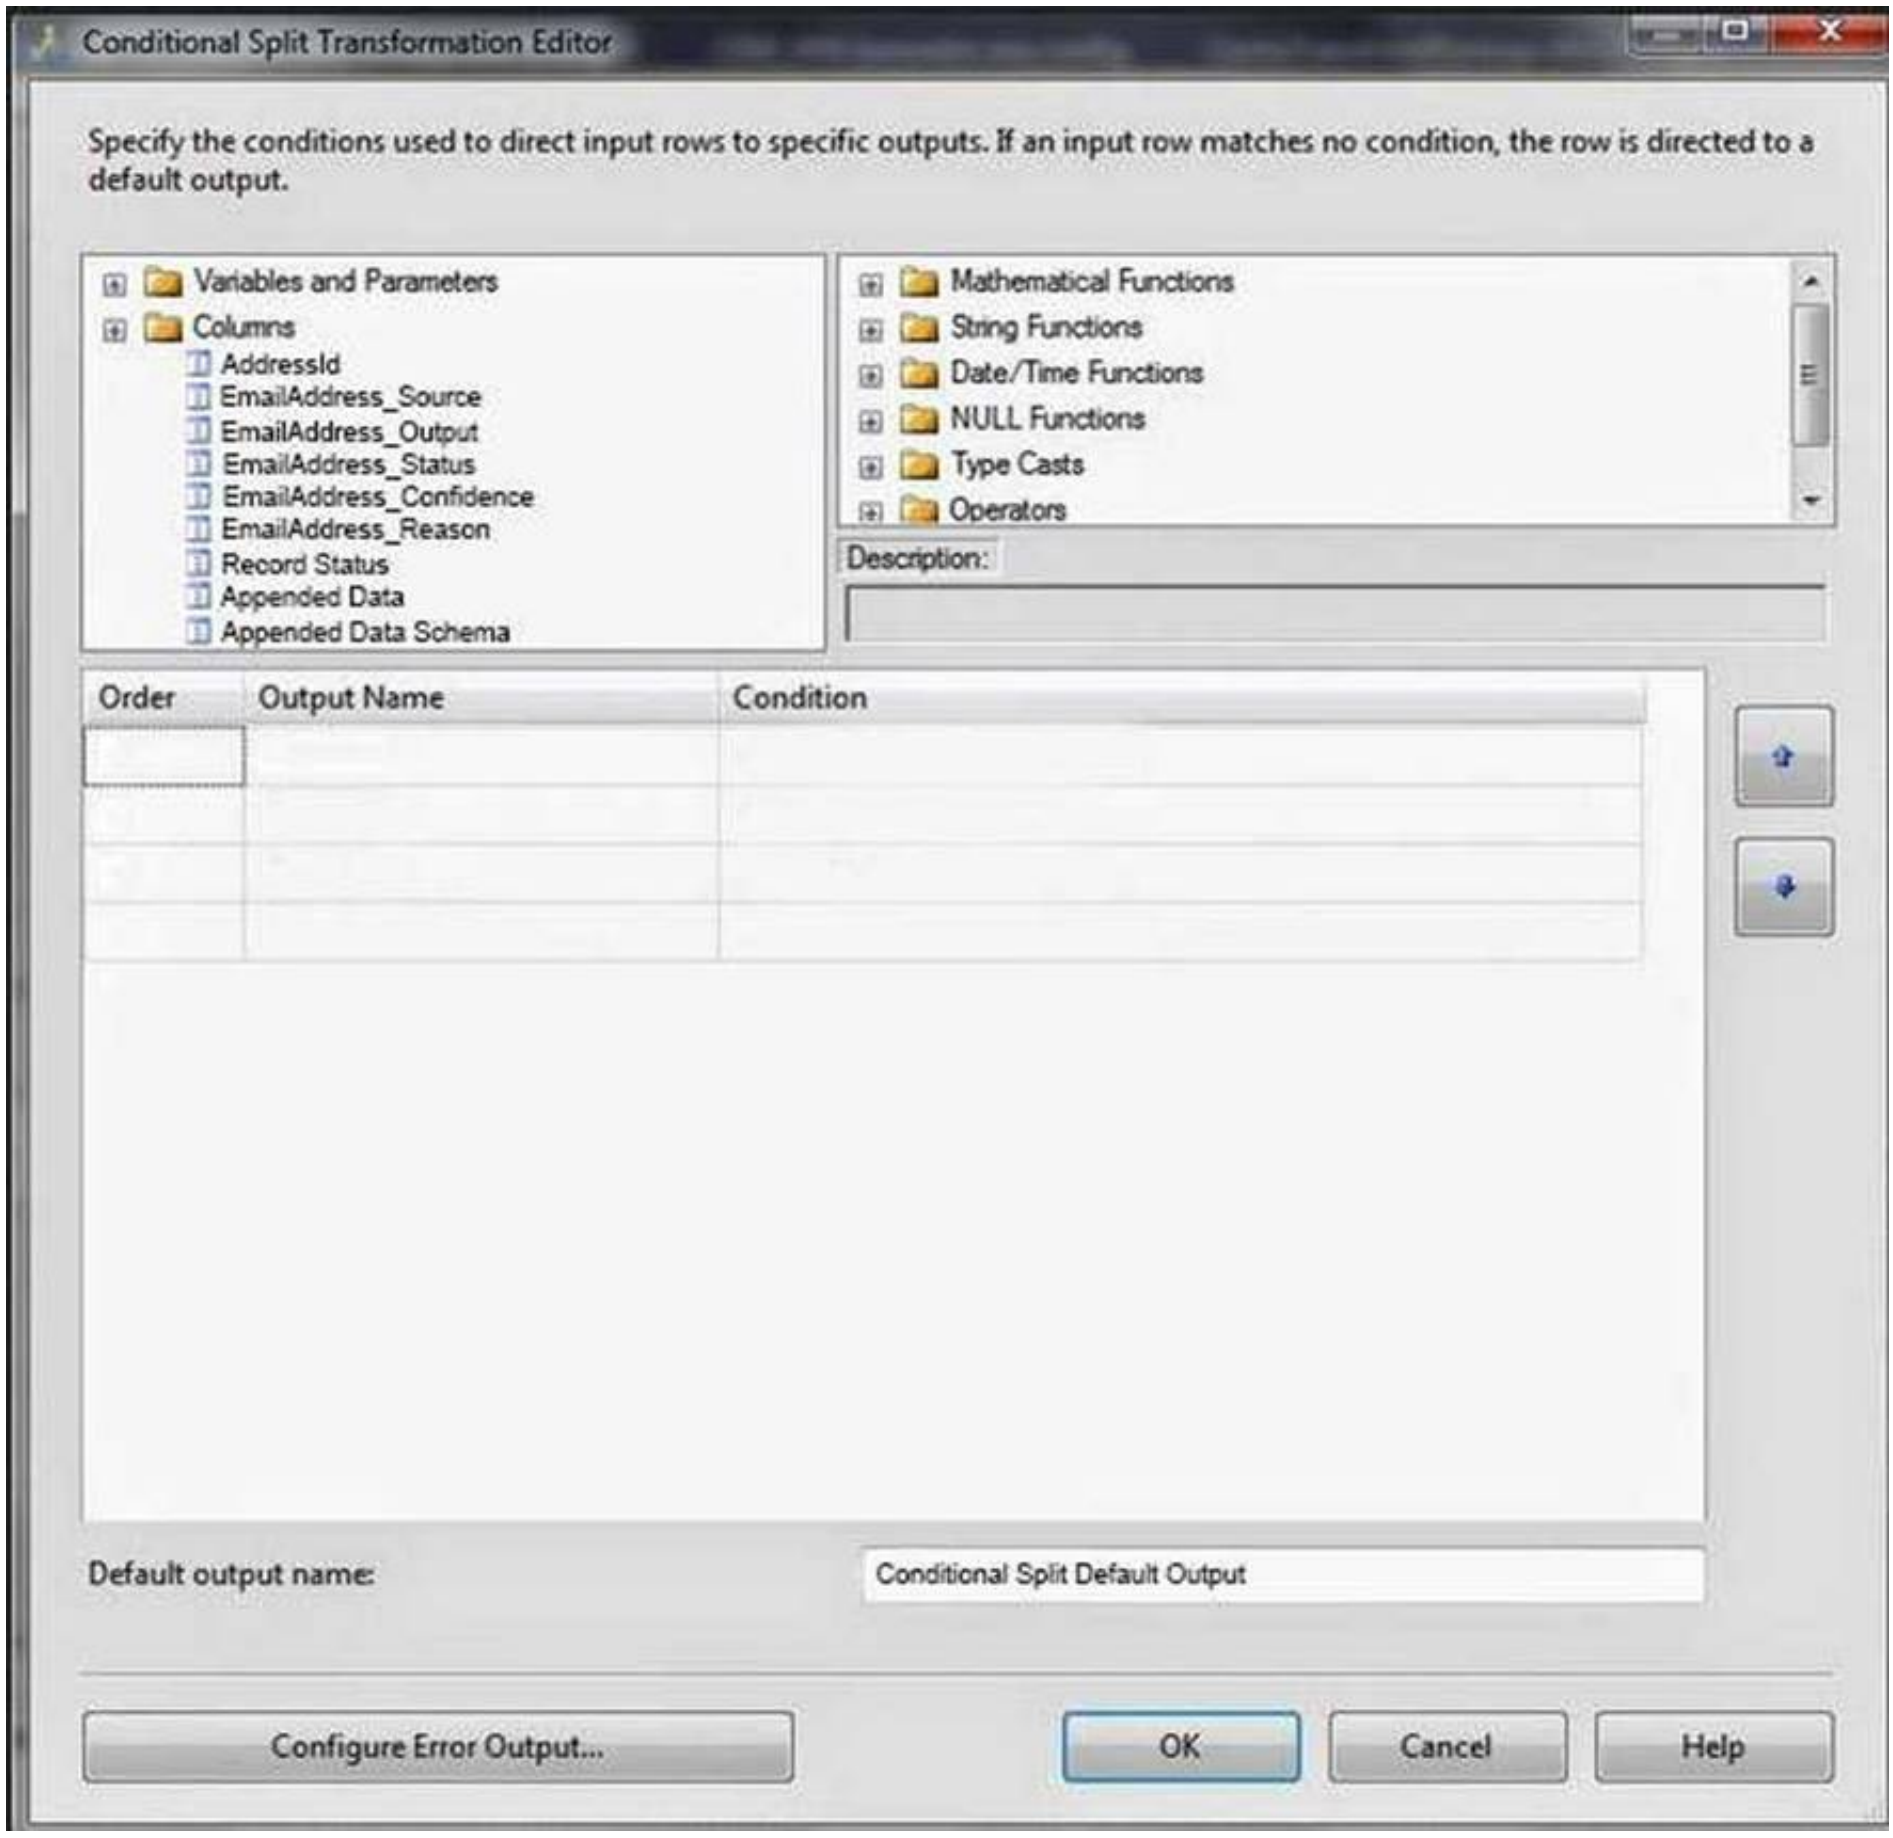


You need to split the output of the DQS Cleansing task to obtain only Correct values from the EmailAddress column.
For each of the following statements, select Yes if the statement is true. Otherwise, select No.

**Answer Area**

| | Yes | No |
|---|---|---|
| You can use the EmailAddress_Output column to split the output. | ○ | ○ |
| You can use the EmailAddress_Status solumn to split the output. | ○ | ○ |
| You can use the EmailAddress_Reason column to split the output. | ○ | ○ |

A. Mastered
B. Not Mastered

**Answer:** A

**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. Is 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 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 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 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 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 alter the package to log the start and completion times for a task to a table in the on-premises SQL Server instance.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** A

**NEW QUESTION 54**
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 57**
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.

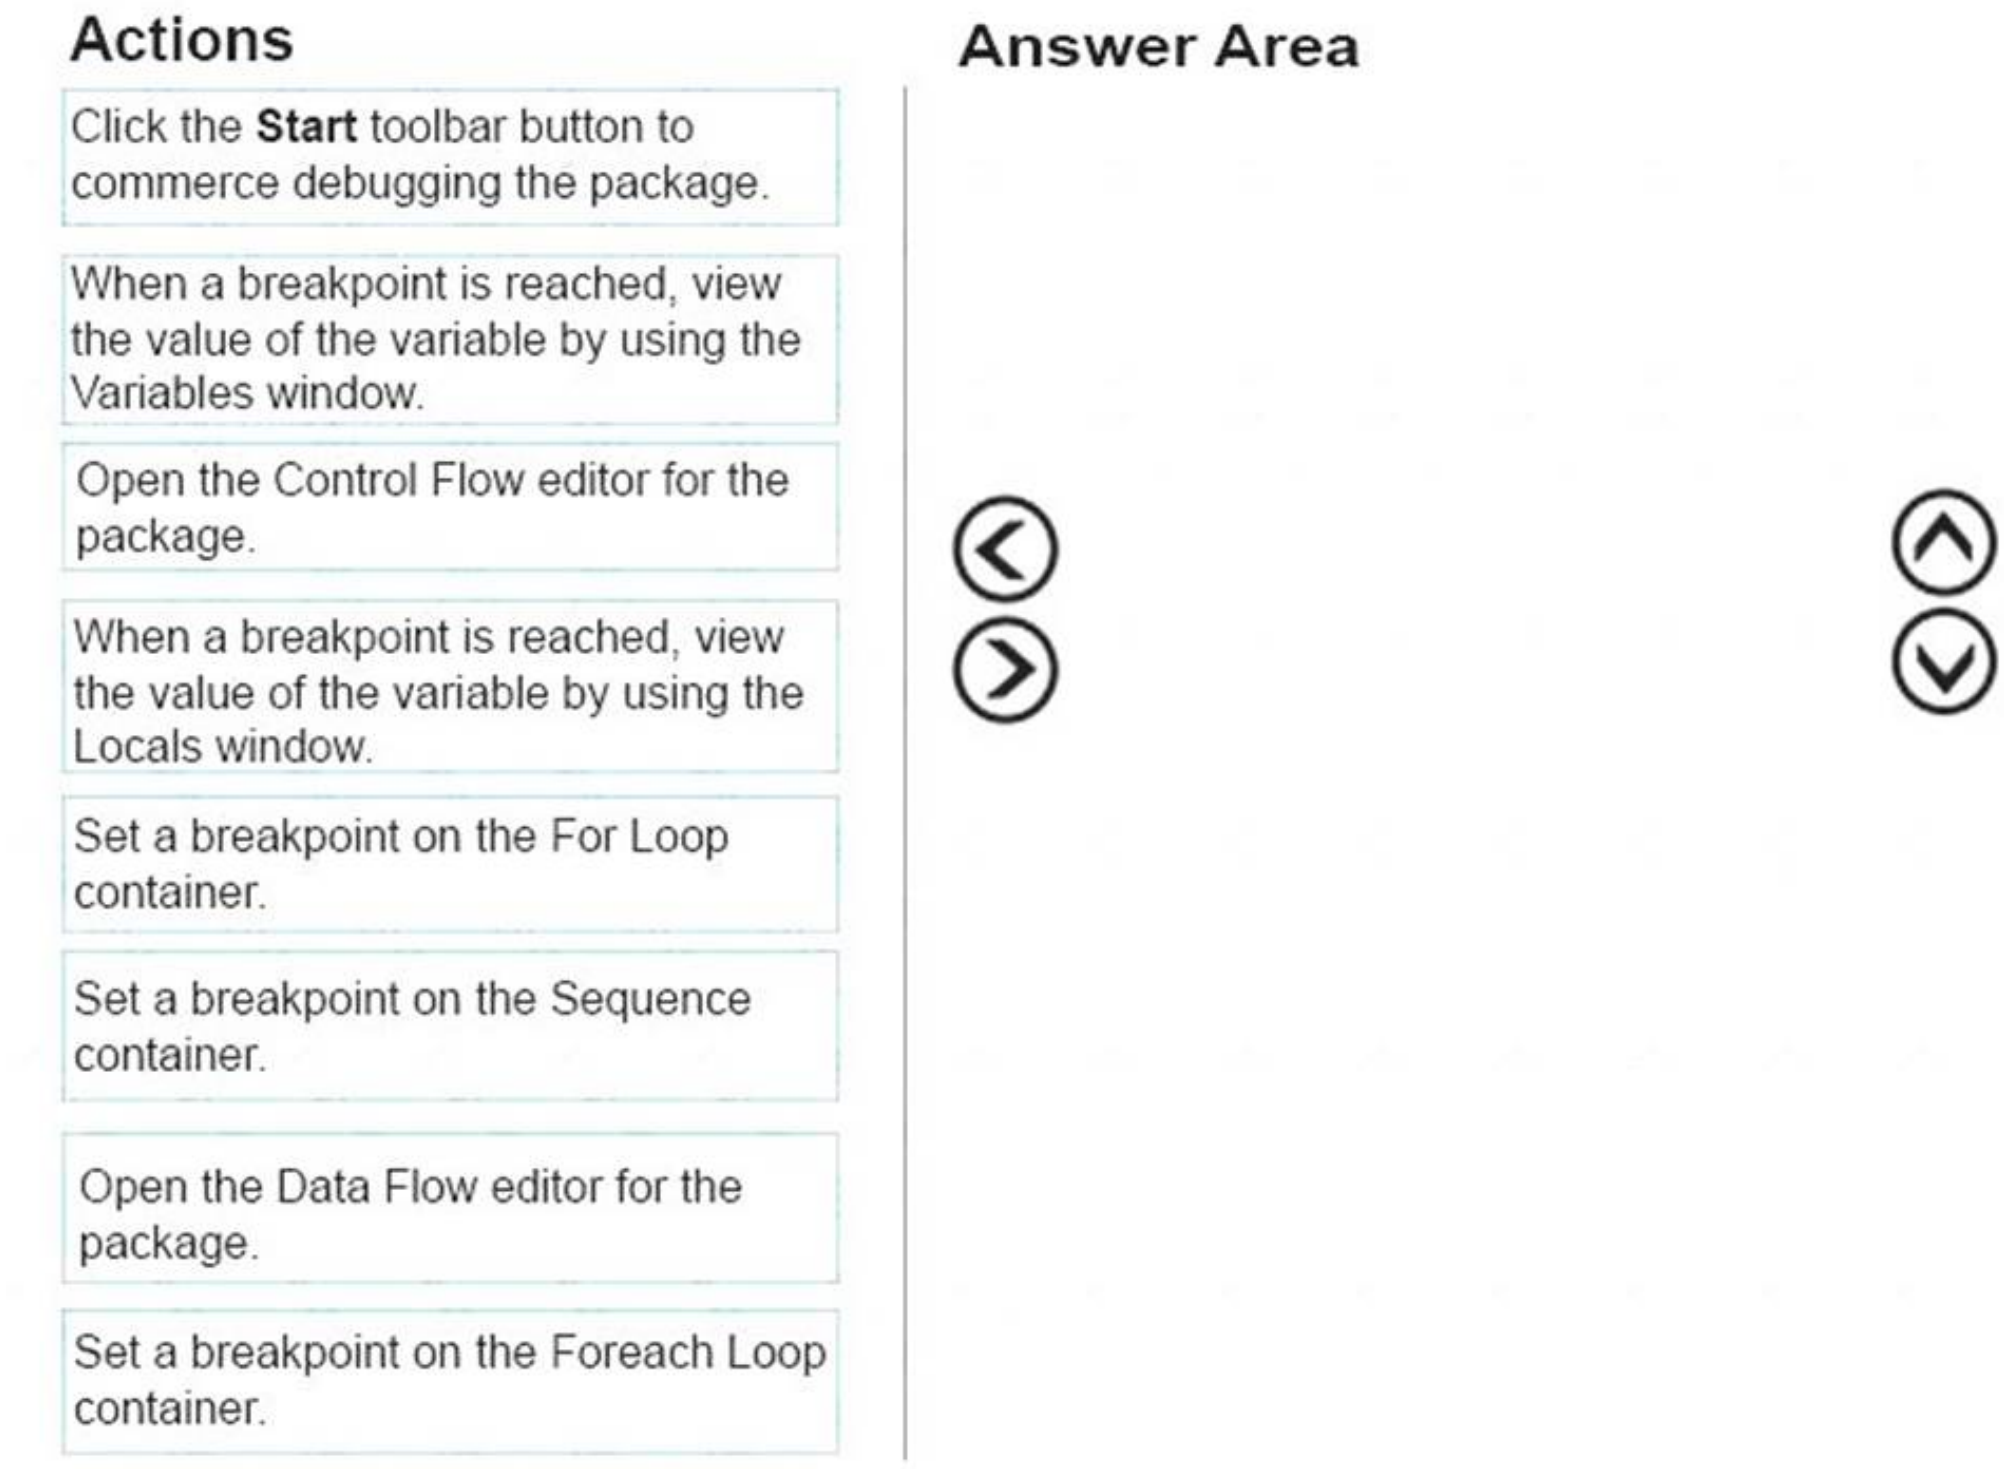


A. Mastered
B. Not Mastered

**Answer:** A

**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 62**
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. You need to use XPath to extract information from documents.
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:** B

**NEW QUESTION 64**
You have a data warehouse named DW1 that contains 20 years of data. DW1 contains a very large fact table. New data is loaded to the fact table monthly. Many reports query DW1 for the past year of data. Users frequently report that the reports are slow.
You need to modify the fact table to minimize the amount of time it takes to run the reports. The solution must ensure that other reports can continue to be generated from DW1.
What should you do?

A. Move the historical data to SAS disks and move the data from the past year to SSD disk
B. Run the ALTERTABLE statement.
C. Move all the data to SSD disk
D. Load and archive the data by using partition switching.
E. Move all the data to SAS disk
F. Load and archive the data by using partition switching.
G. Move the historical data to SAS disks and move the data for the past year to SSD disk
H. Create a distributed partitioned view.

**Answer:** A

**Explanation:**
We use ALTER TABLE to partition the table.

**NEW QUESTION 69**
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.

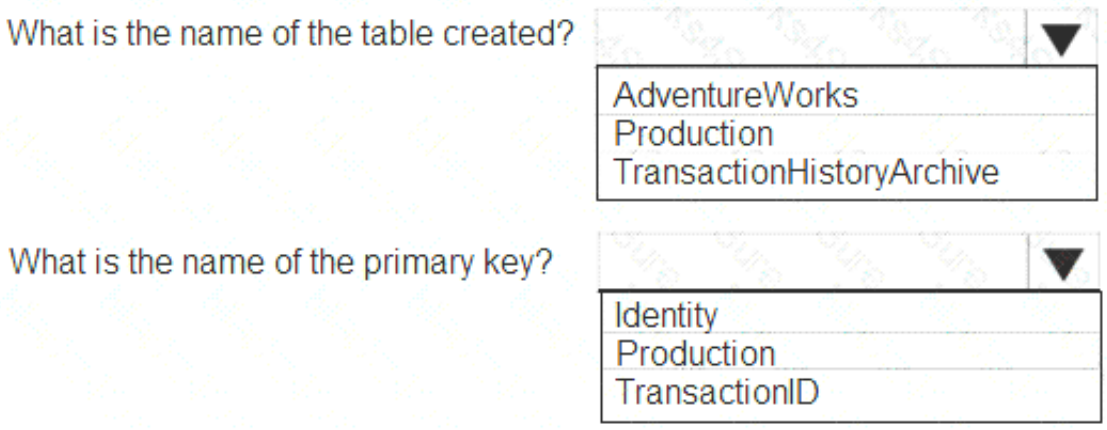


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

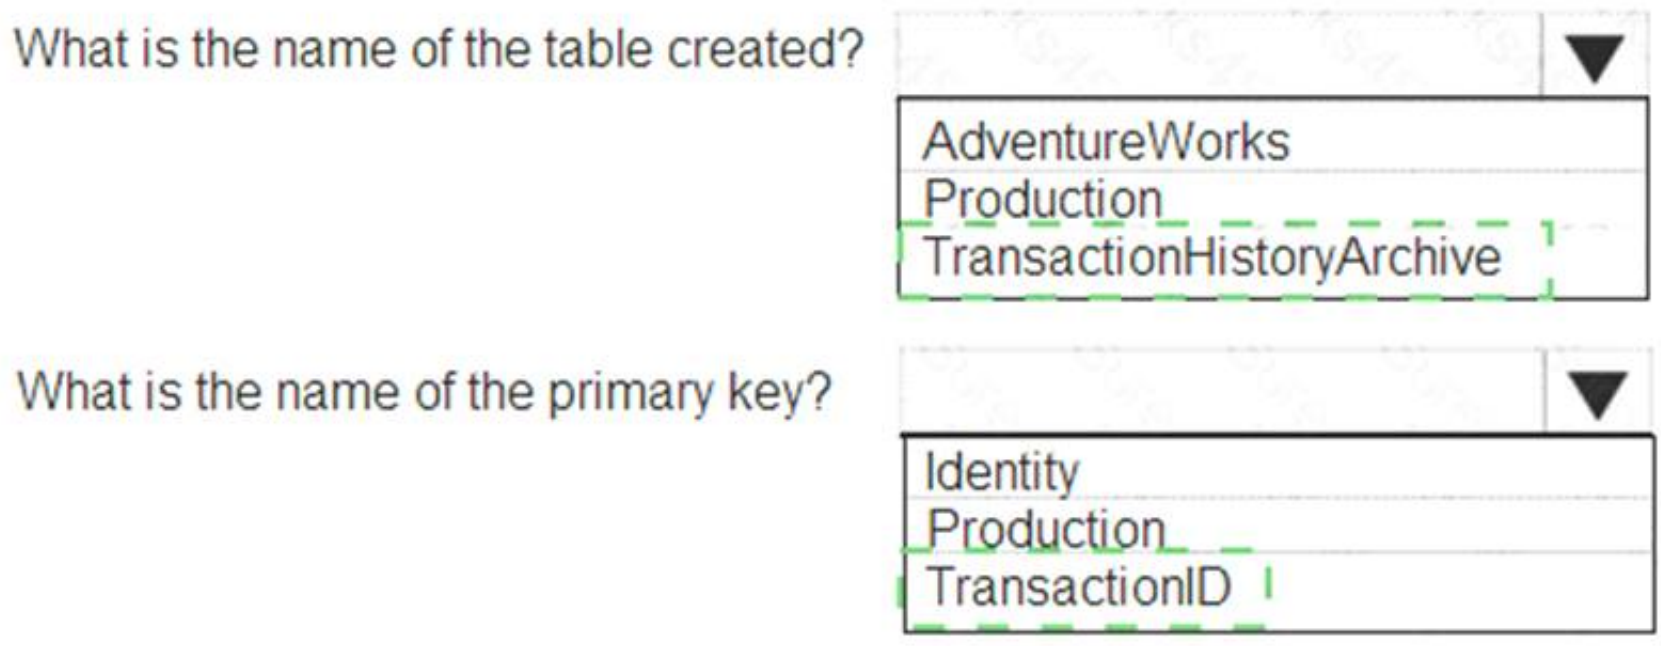


**NEW QUESTION 72**
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question 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 76**
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 warenouse 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 Diamension.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.

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 | 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
```

A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**

**Answer area**

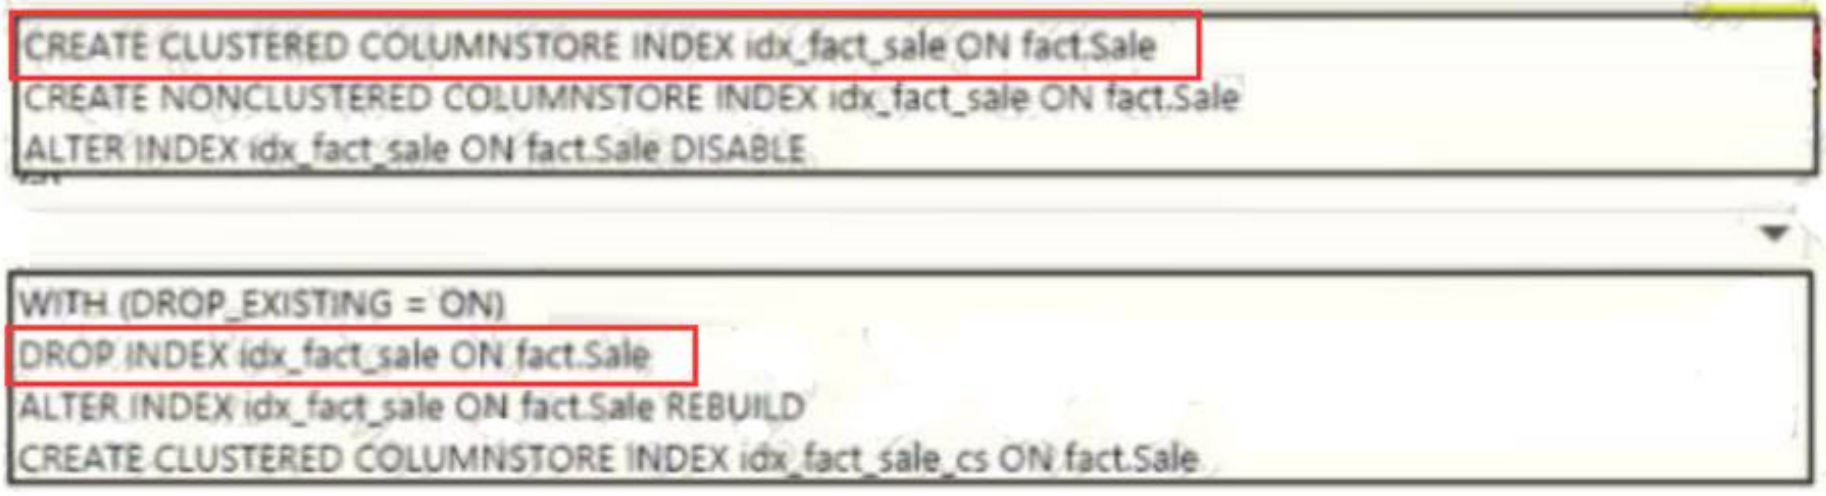


**NEW QUESTION 78**
You have a database named DB1 that contains millions of rows. You plan to perform a weekly audit of the changes to the rows.
You need to ensure that you can view which rows were modified and the hour that the modification occurred. What should you do?

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 83**
You have a series of analytic data models and reports that provide insights into the participation rates for sports at different schools. Users enter information about sports and participants into a client application. The application stores this transactional data in a Microsoft SQL Server database. A SQL Server Integration Services (SSIS) package loads the data into the models.
When users enter data, they do not consistently apply the correct names for the sports. The following table shows examples of the data entry issues.

| Sport | Variations entered by users |
|---|---|
| baseball | baseball, ball, play ball |
| football | soccer, football |

You need to create a new knowledge base to improve the quality of the sport name data.
How should you configure the knowledge base? To answer, select the appropriate options in the dialog box in the answer area.

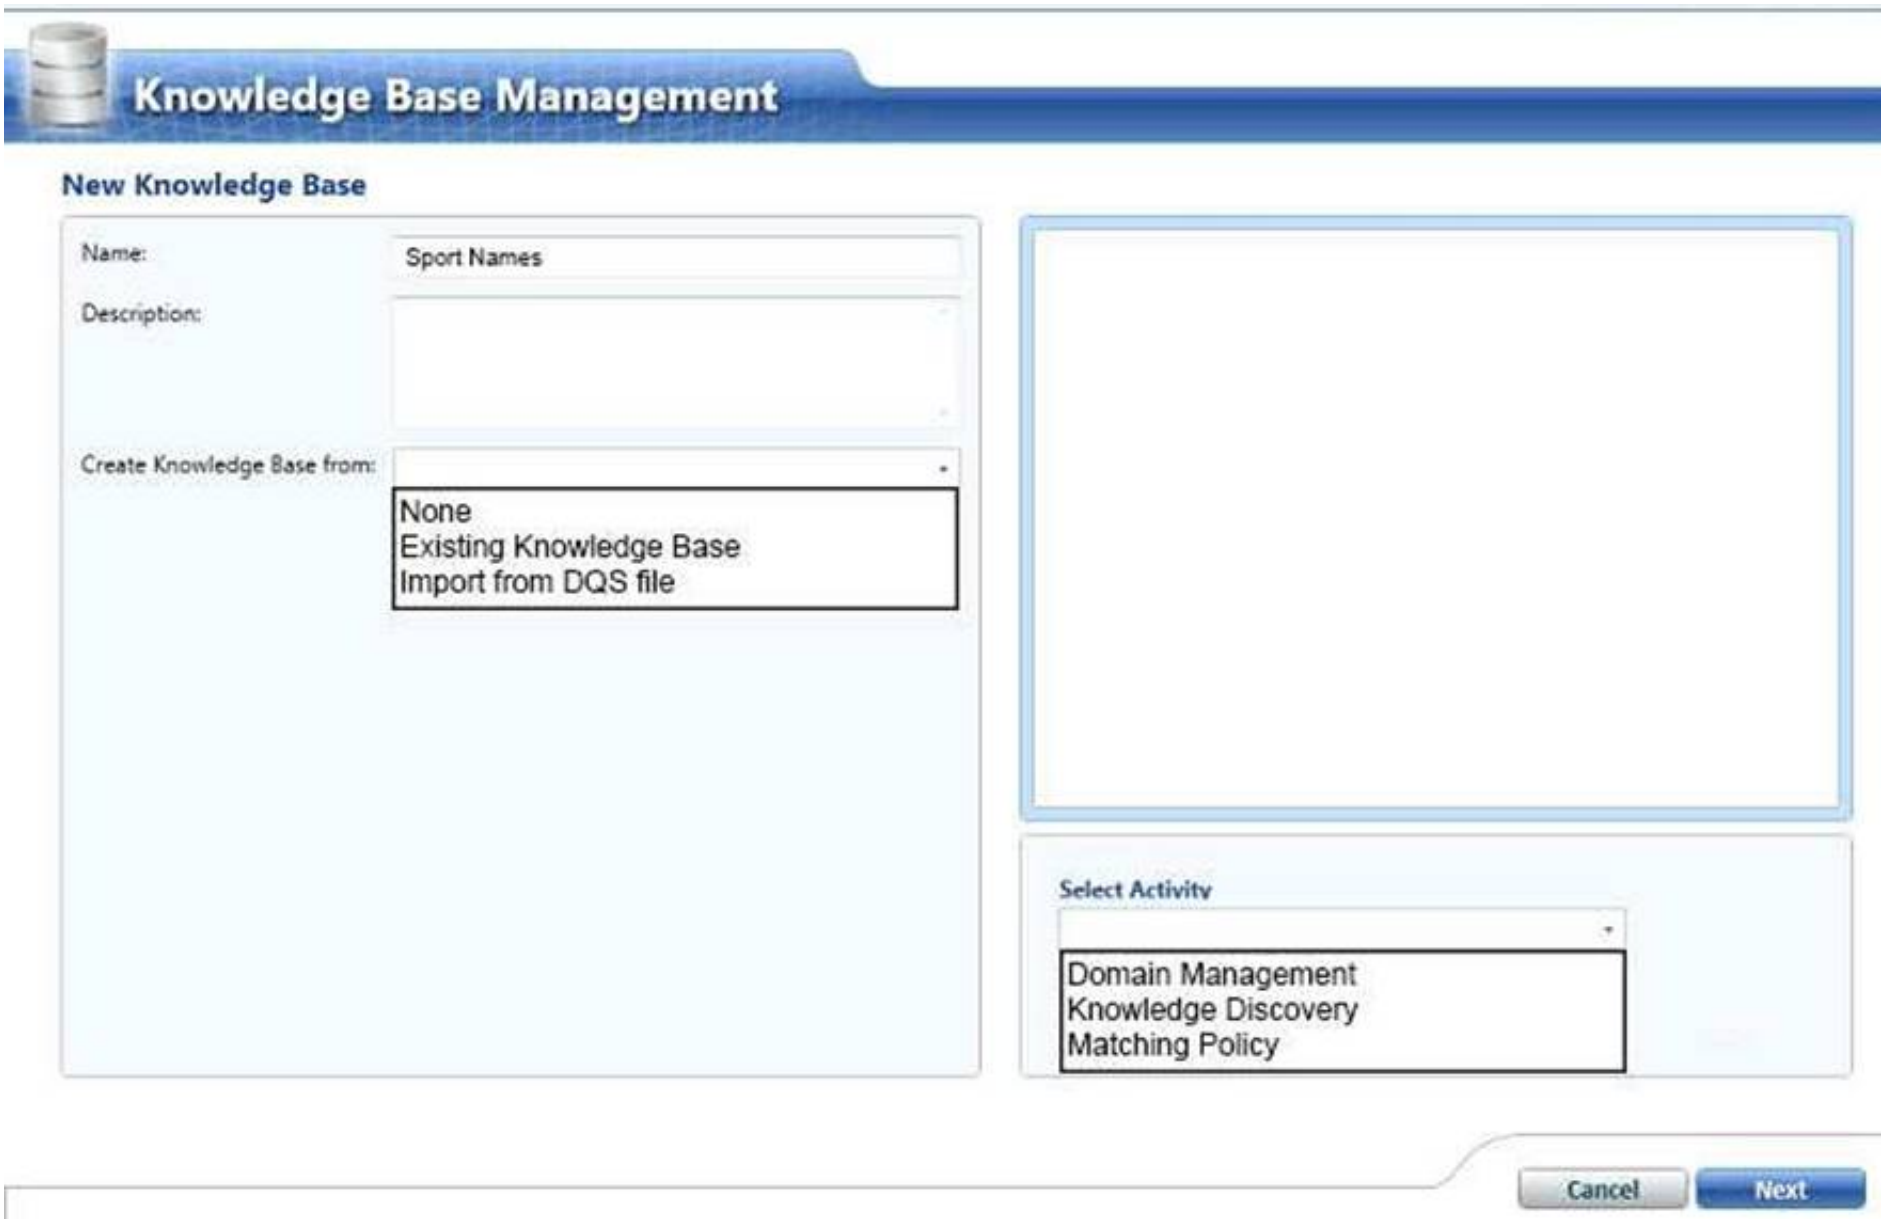


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
Spot 1: Create Knowledge base from: None
Select None if you do not want to base the new knowledge base on an existing knowledge base or data file.

**NEW QUESTION 88**
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 configure the Fact.Order table.
Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

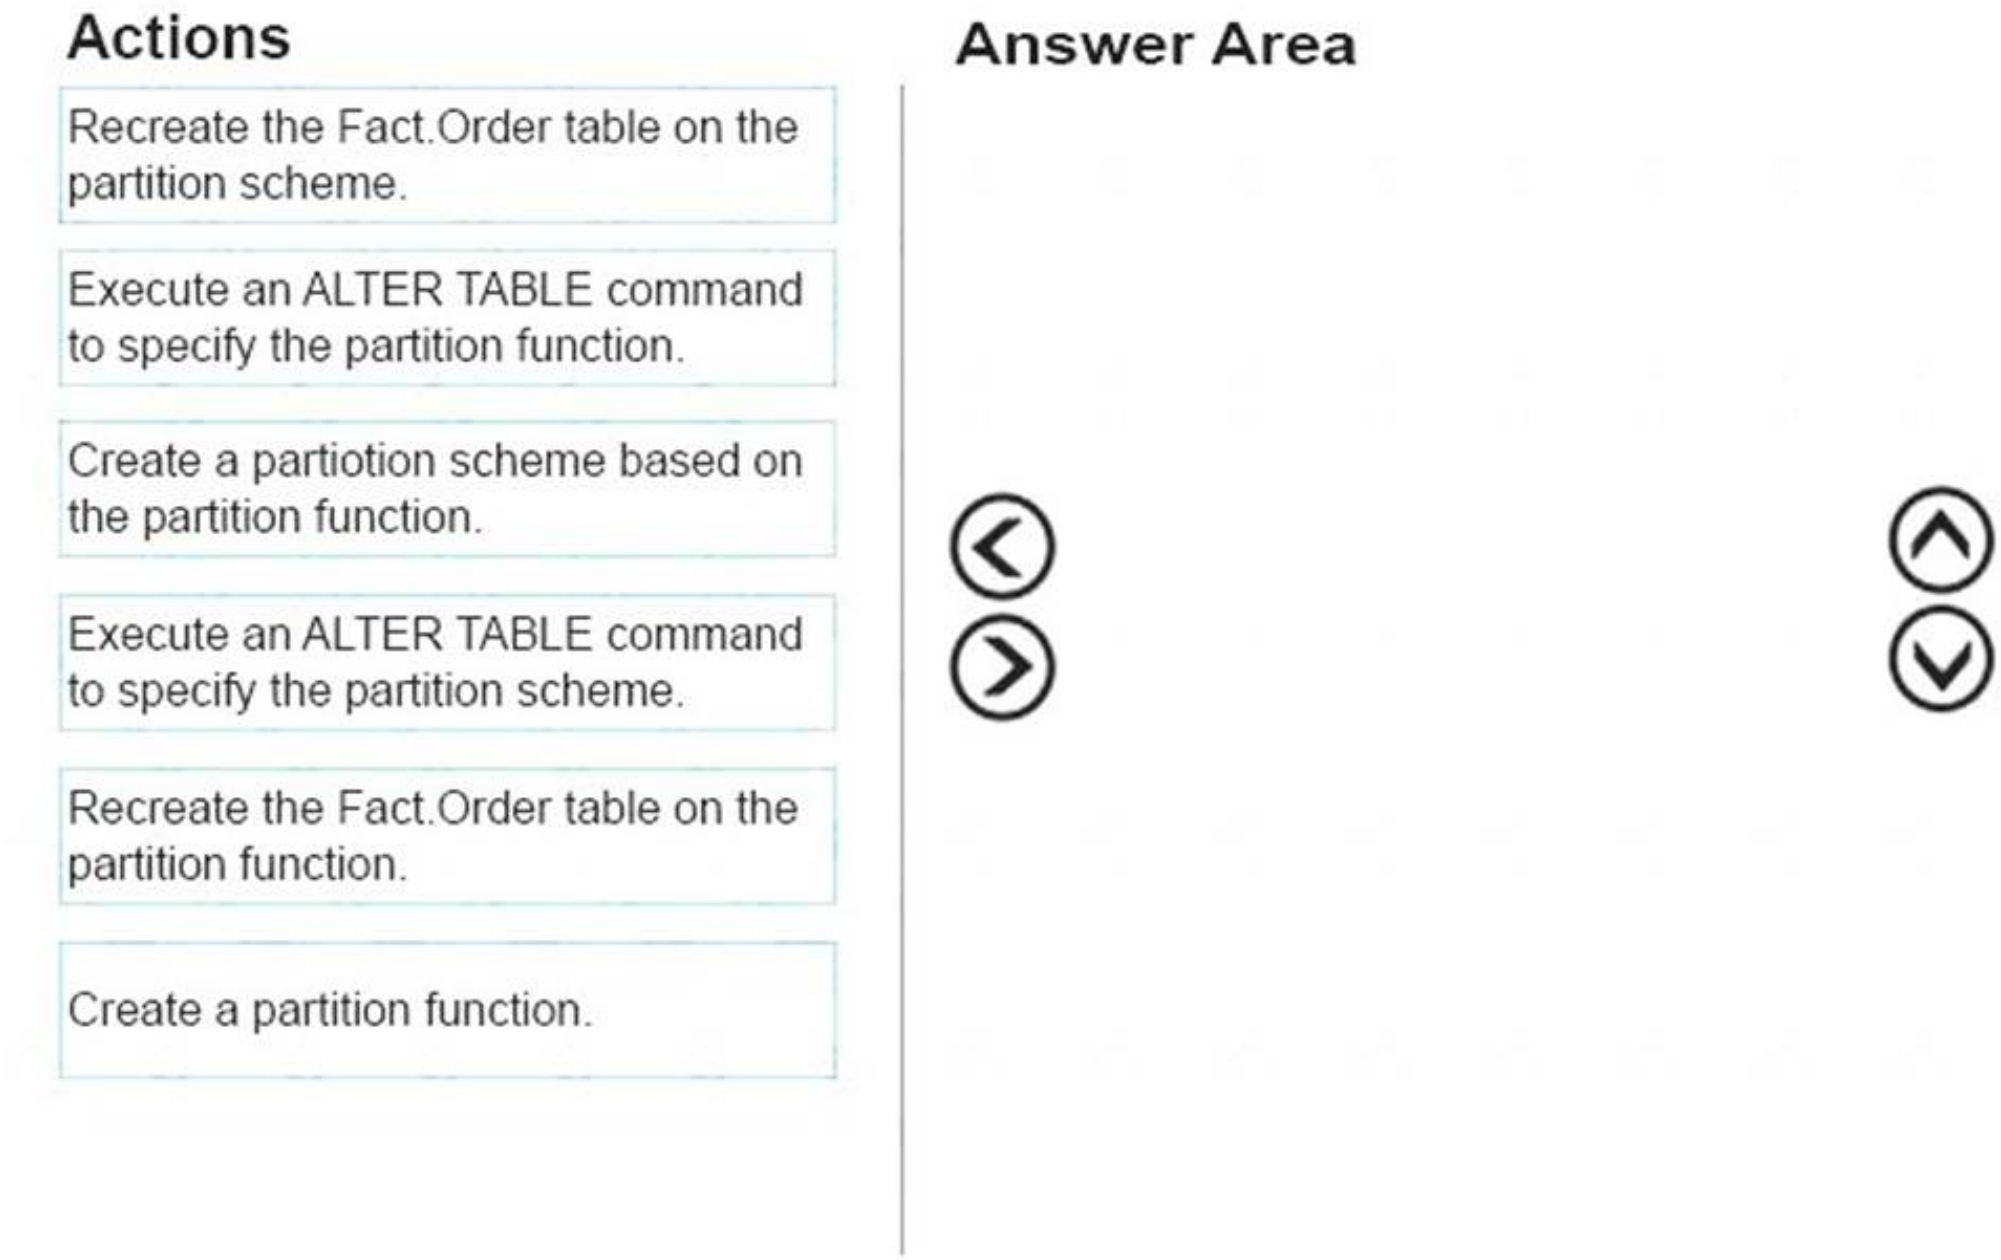


A. Mastered
B. Not Mastered

**Answer:** A

**Explanation:**
From scenario: Partition the Fact.Order table and retain a total of seven years of data. Maximize the performance during the data loading process for the Fact.Order partition.
Step 1: Create a partition function.
Using CREATE PARTITION FUNCTION is the first step in creating a partitioned table or index. Step 2: Create a partition scheme based on the partition function.
To migrate SQL Server partition definitions to SQL Data Warehouse simply: Step 3: Execute an ALTER TABLE command to specify the partition function.
References: https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-partition

**NEW QUESTION 89**
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 93**
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 95**
You create a Microsoft SQL Server Integration Services (SSIS) package as shown in the SSIS Package exhibit. (Click the Exhibit button.)

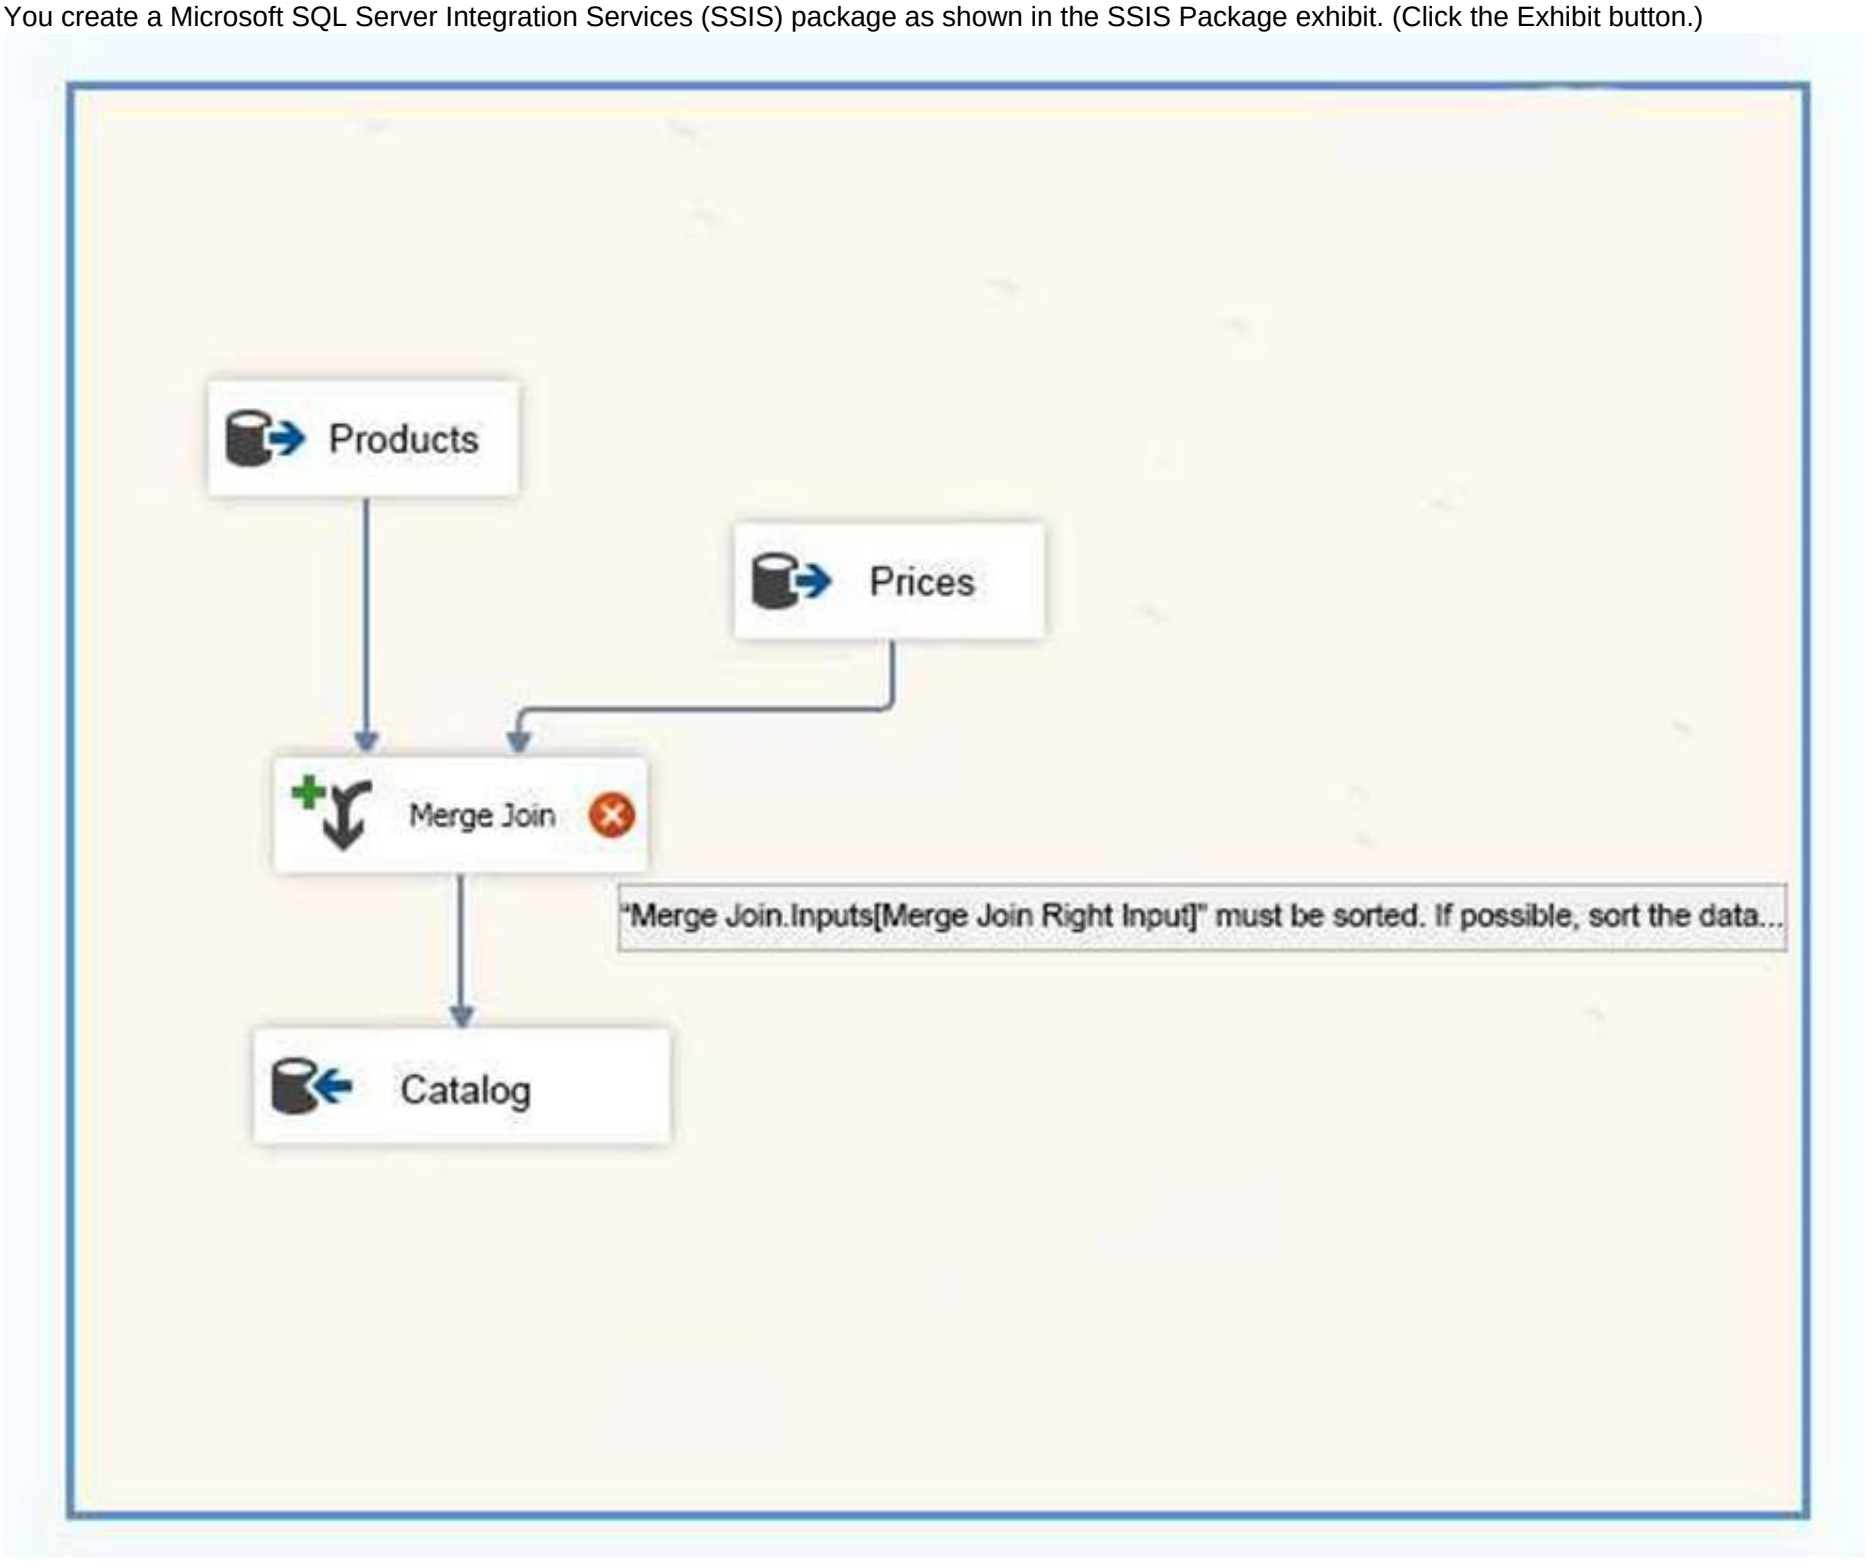


The package uses data from the Products table and the Prices table. Properties of the Prices source are shown in the OLE DB Source Editor exhibit (Click the Exhibit Button.) and the Advanced Editor for Prices exhibit (Click the Exhibit button.)

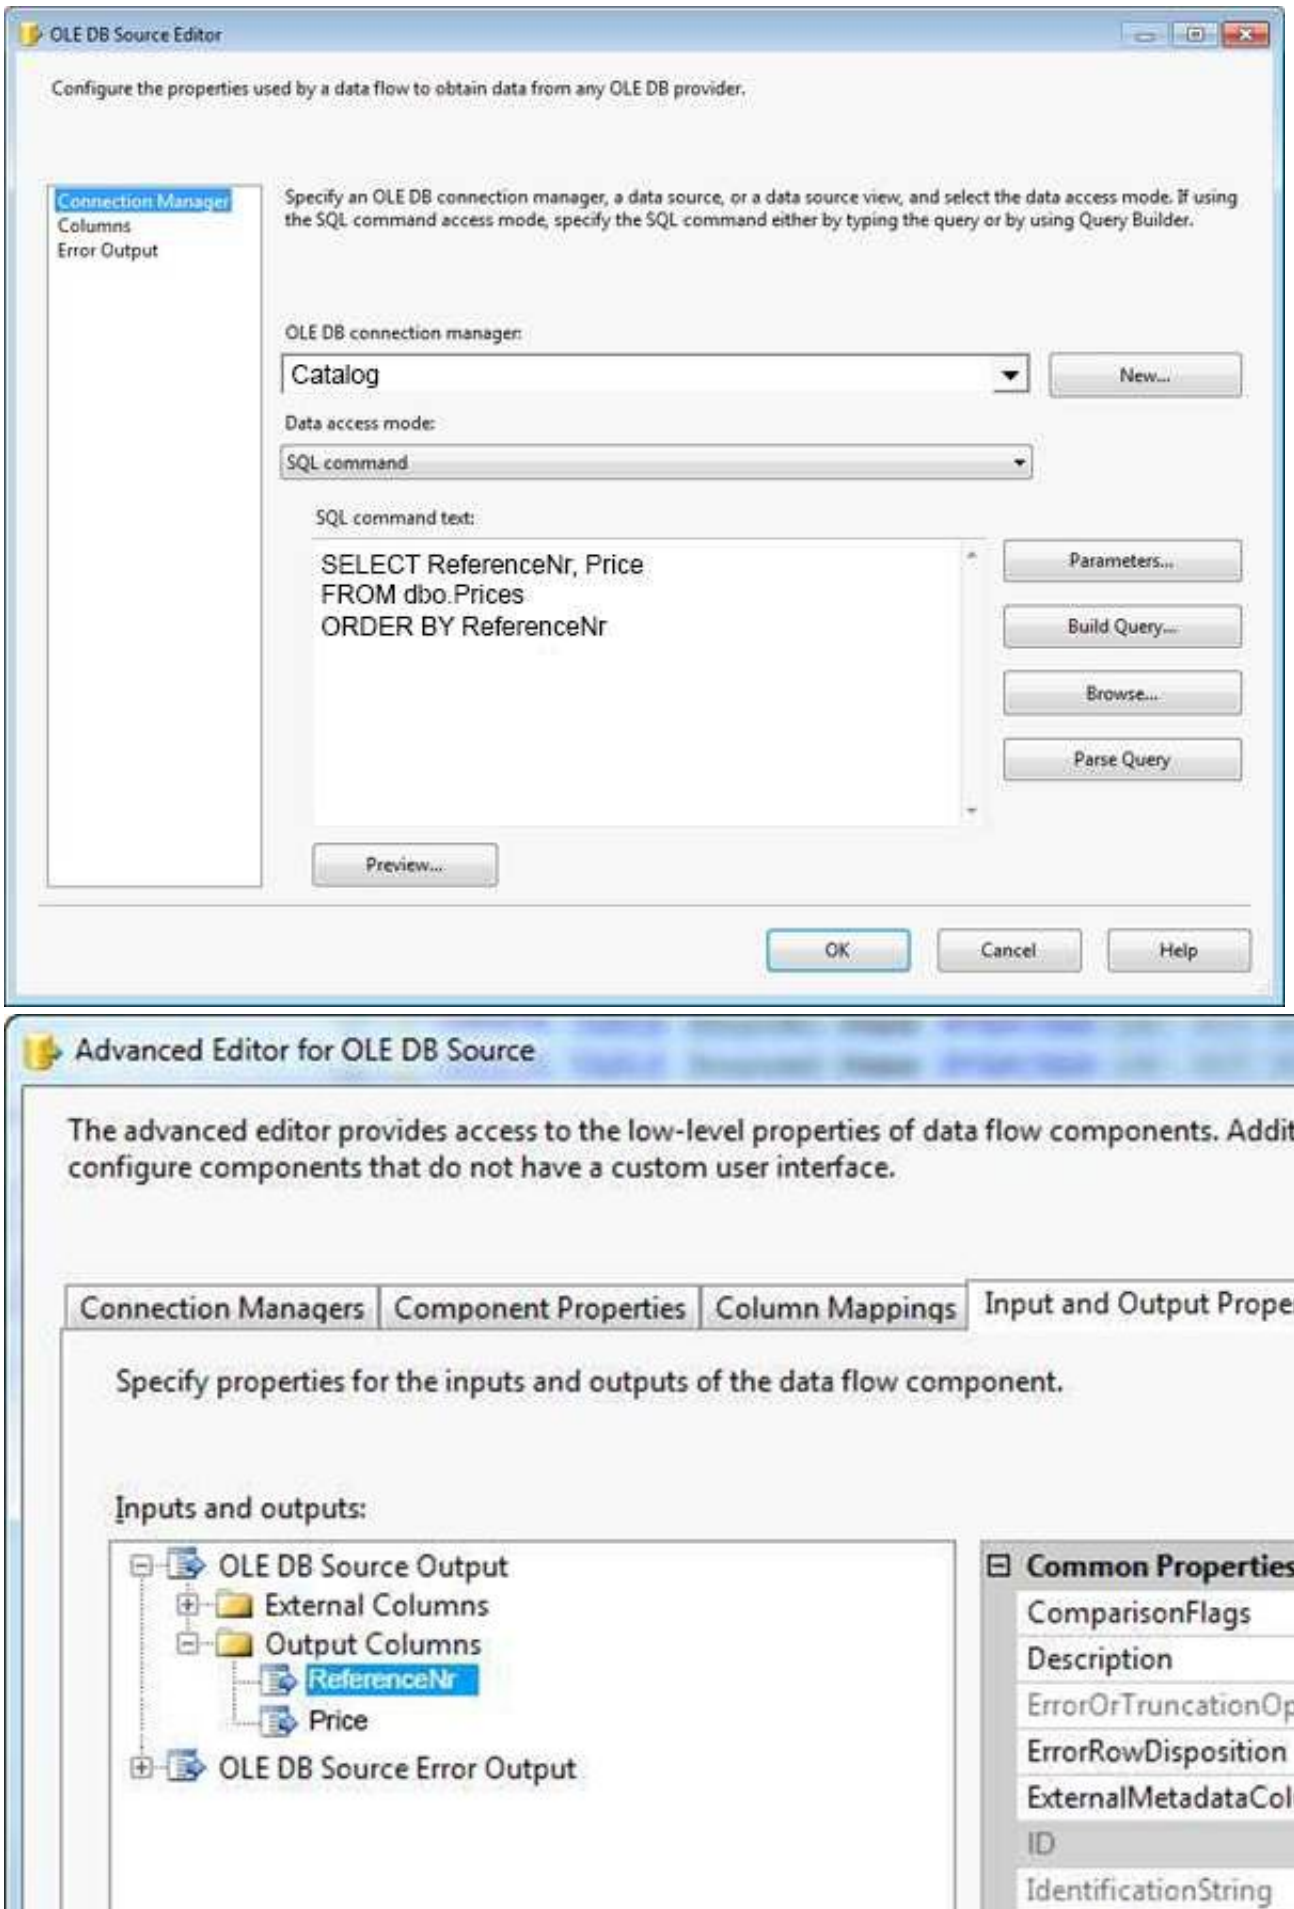


You join the Products and Prices tables by using the ReferenceNr column. You need to resolve the error with the package.

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

**Answer Area**

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

A. Mastered
B. Not Mastered

**Answer:** A

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

**NEW QUESTION 96**
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 99**
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 103**
You are developing a Microsoft SQL Server Integration Services (SSIS) package that loads a data warehouse. You need to inspect the data that is being processed by the package. What should you do first?

A. Set a break point on the Control Flow path.
B. Enable SQL Trace.
C. Enable logging on the Data Flow path.
D. Enable a data viewer on the Data Flow path.

**Answer:** A

**NEW QUESTION 107**
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 109**
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-berOfTrades | 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)

**NEW QUESTION 110**
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.)

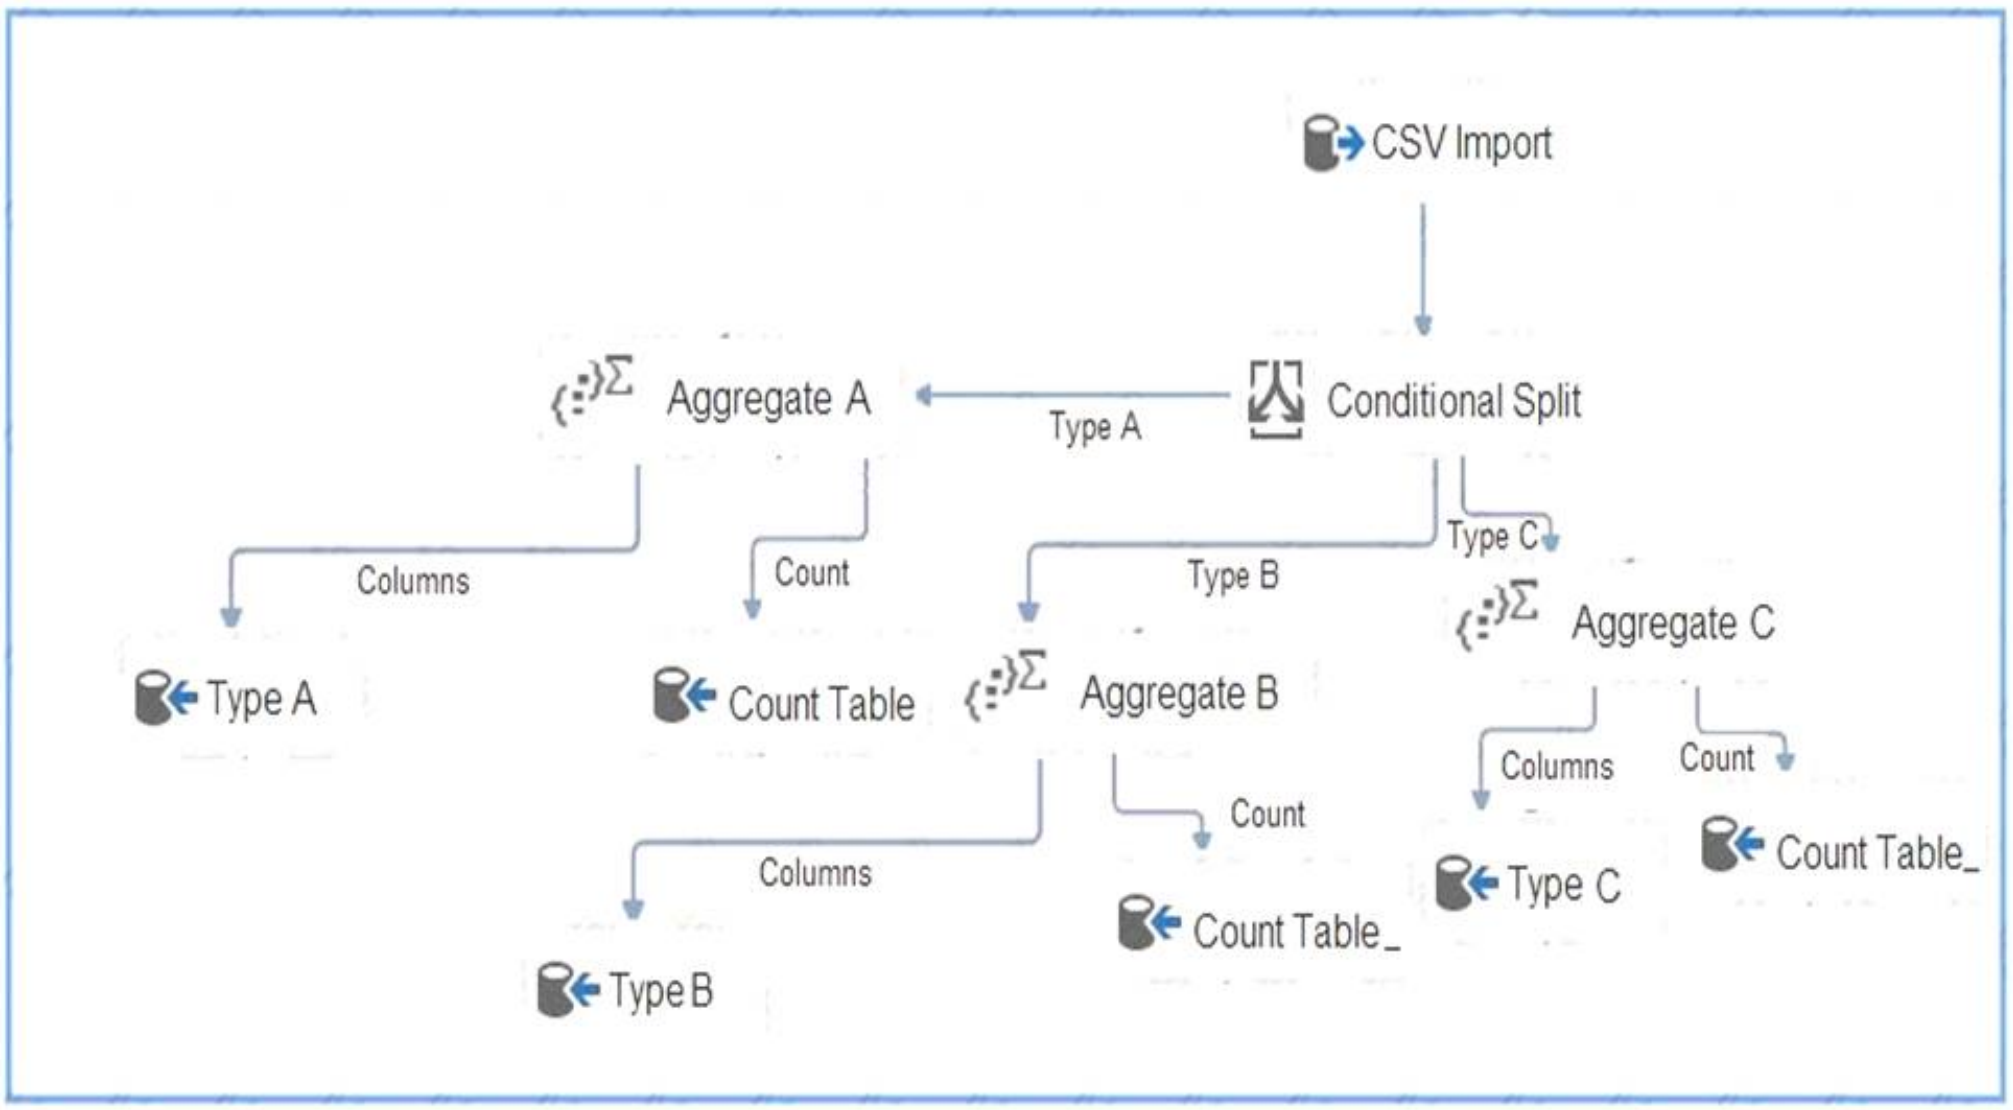


Does the solution meet the goal?

A. Yes
B. NO

**Answer:** A

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

**NEW QUESTION 114**
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.).


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. | ○ | ○ |
| Creating a filtered unique index on the EmployeeSSN column solves the issue. | ○ | ○ |

A. Mastered
B. Not Mastered

**Answer:** A

**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 118**
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 123**
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].Tabid to DW.Tablel.

• All rows from [source].Table1 must be captured in DW.Tablel for error.Tablel.
• The table error.Tablel must accept rows that fail upon insertion into DW.Tablel 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 [sourceJ.Tablel 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.

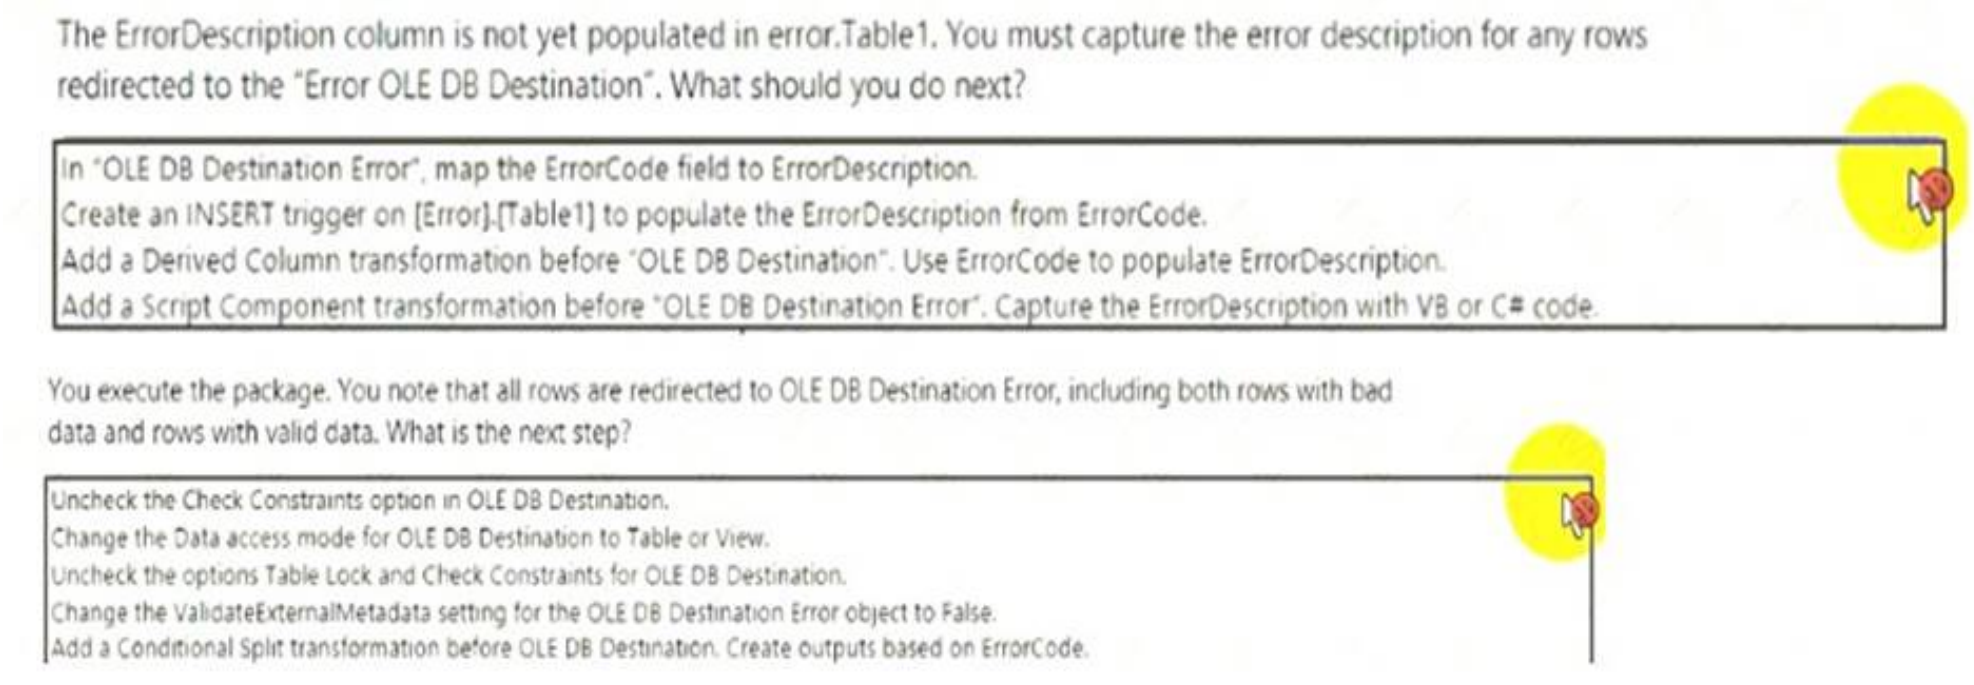


A. Mastered
B. Not Mastered

**Answer:** A

**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 125**
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.timmitchell.net/post/2016/01/18/getting-started-with-change-tracking-in-sql-server/

**NEW QUESTION 127**
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.

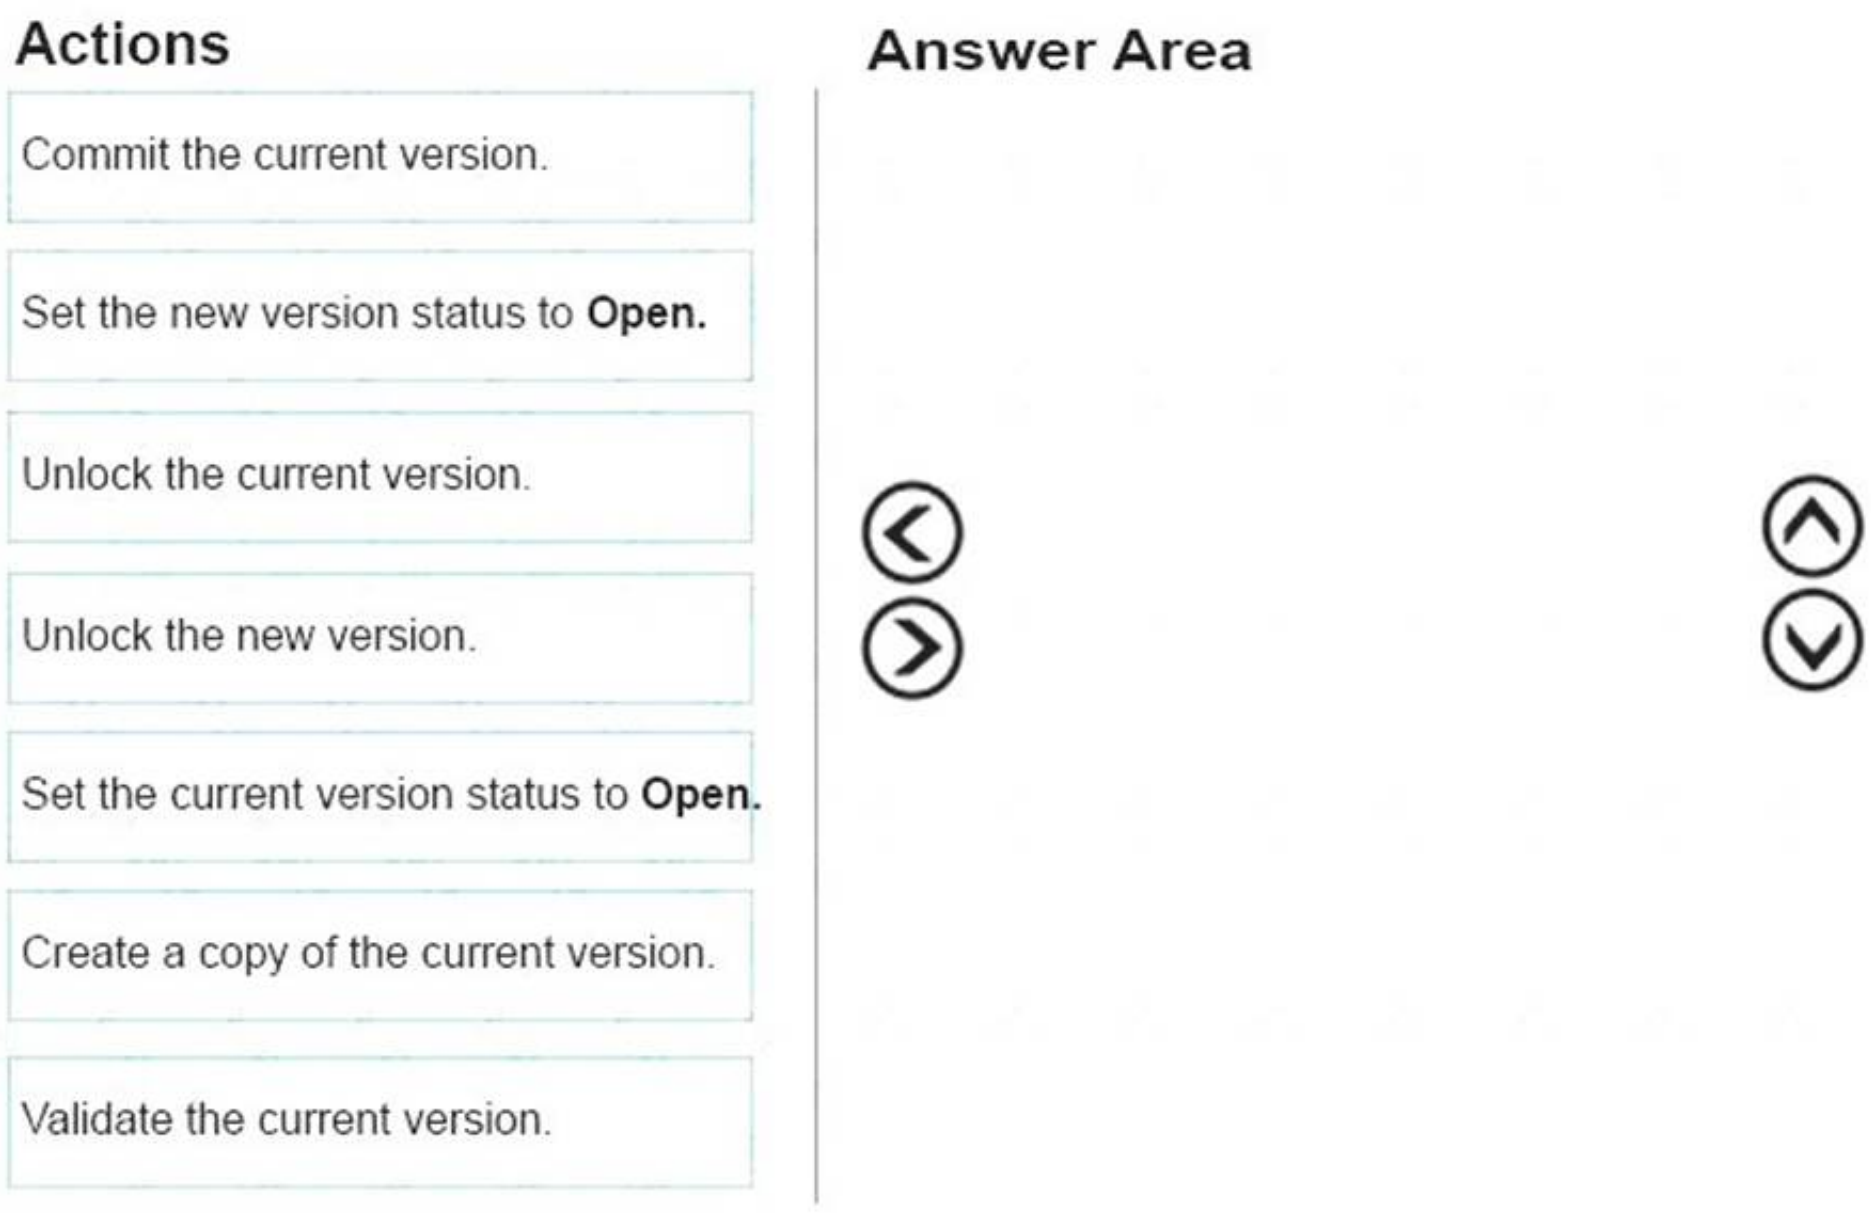


A. Mastered
B. Not Mastered

**Answer:** A

**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 129**
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 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.
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.)

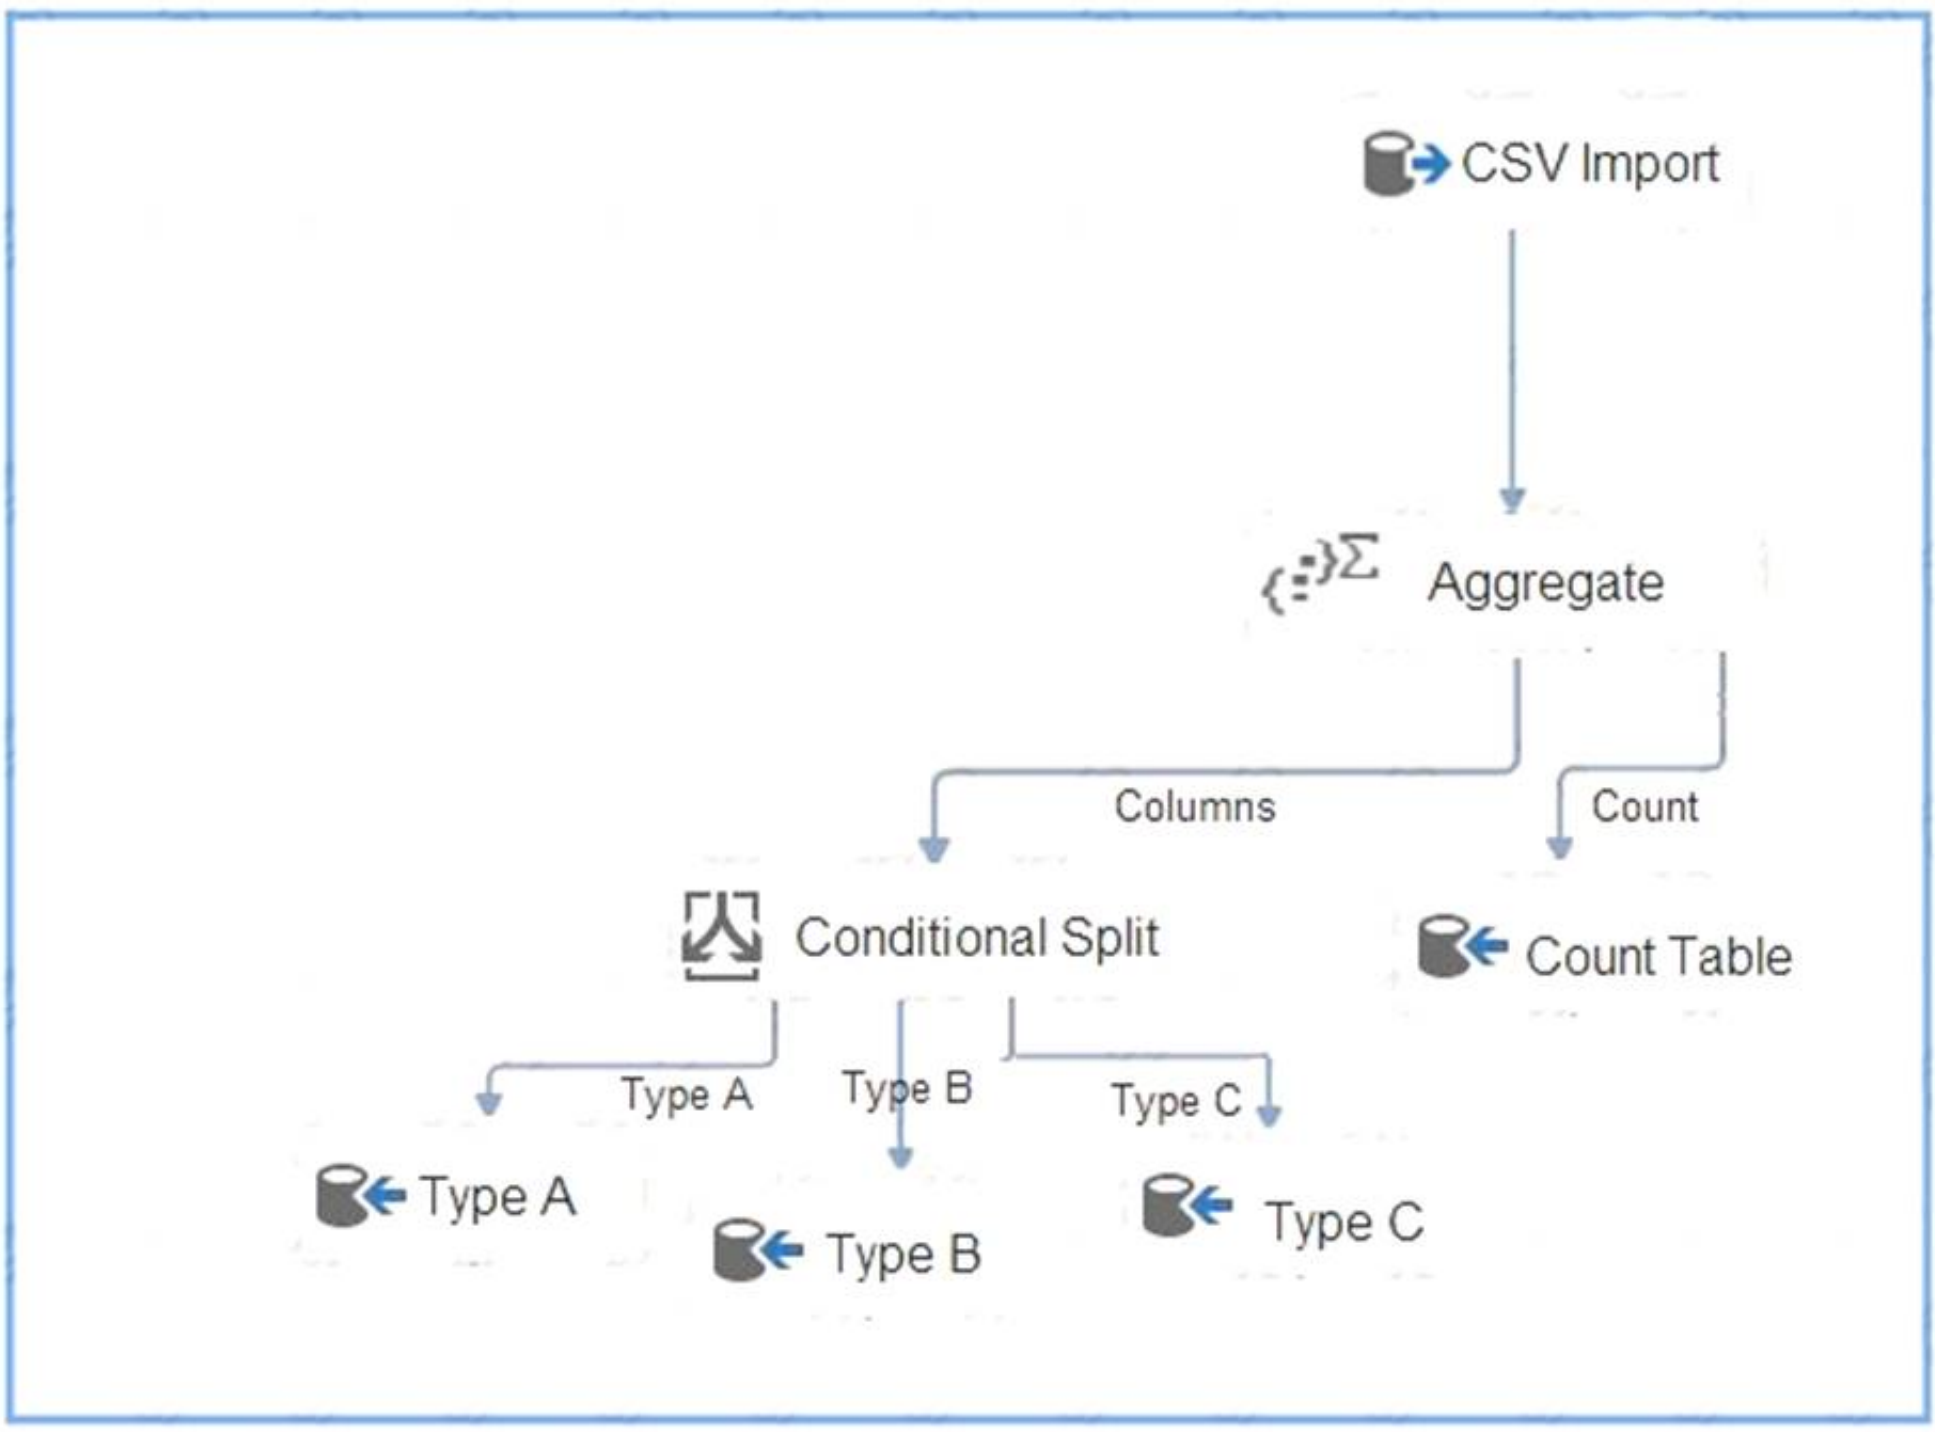


Does the solution meet the goal?

A. Yes
B. NO

**Answer:** B

**Explanation:**
The conditional split must be before the count.

**NEW QUESTION 131**
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 toDomainbased.

**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 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
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.

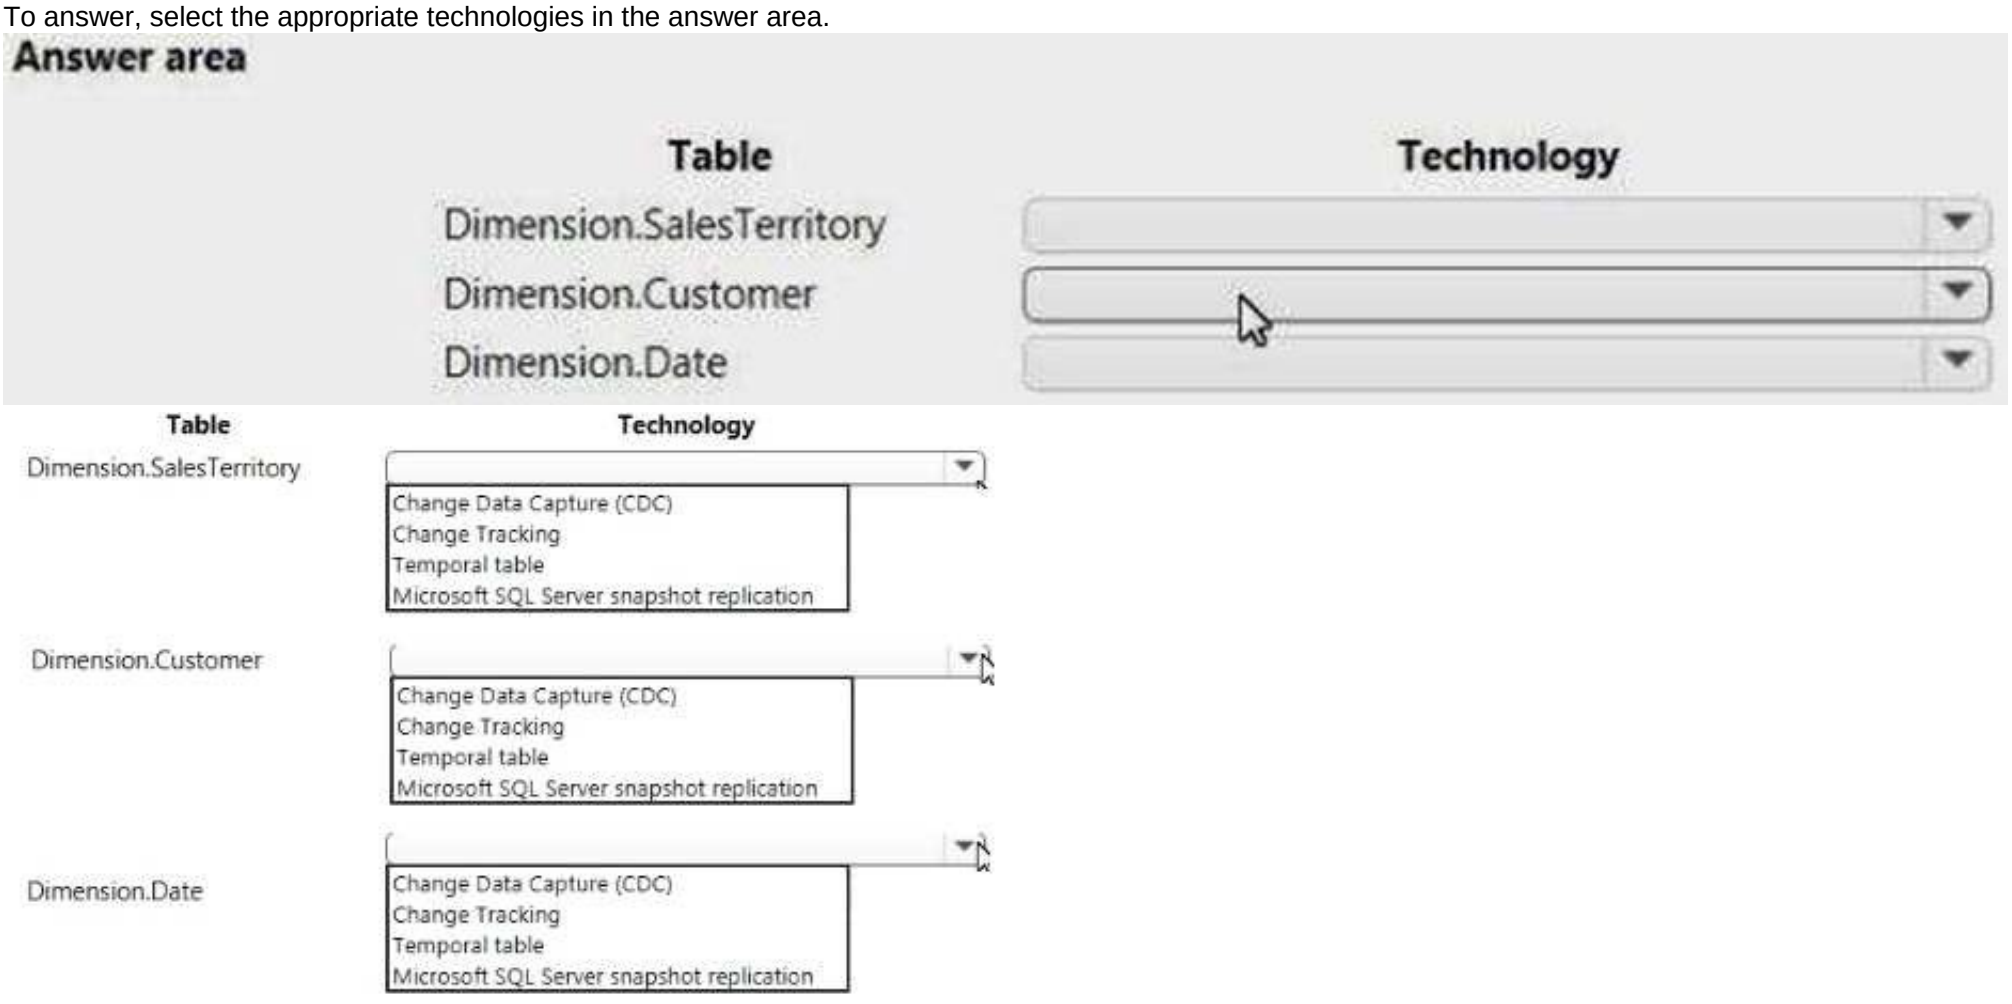


A. Mastered
B. Not Mastered

**Answer:** A

**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 136**
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 SOL 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: In the Geography model, publish a business rule with a Change Value action.
Does the solution meet the goal?

A. Yes
B. No

**Answer:** B

**NEW QUESTION 139**
You have a data warehouse that contains a fact table named Table1 and a Product table named Dim1. Dim1 is configured as shown in the following table.

| Column name | Column data type |
|---|---|
| ProductID | Integer identity |
| ProductKey | Char(10) |
| Name | Varchar(50) |
| Color | Varchar(20) |
| Weight | Decimal (13, 1) |

You are adding a second OLTP system to the data warehouse as a new fact table named Table2. The Product table of the OLTP system is configured as shown in the following table

| Column name | Column data type |
|---|---|
| ProductIdentifier | Char (8) |
| ProductName | Varchar(35) |
| SalesUnit | varchar(25) |
| Weight | Decimal(19,2) |

You need to modify Dim1 to ensure that the table can be used for both fact tables.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

A. Modify the data type of the Weight column in Dim1 to decimal (19, 2).
B. Add the SalesUnit column to Dim1.
C. Modify the data type of the Name column in Dim1 to varchar (85).
D. Drop the ProductKey column from Dim1 and replace the column with the ProductIdentifier column.
E. Drop the Color column from Dim1.
F. Modify the data type of the ProductKey column in Dim1 to char (18).

**Answer:** AD

**NEW QUESTION 142**
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 configuring a Microsoft SQL server named ow1 for a new data warehouse. The server contains eight drives and eight processor cores. Each drive uses a separate physical disk.
You need to configure storage for the tempdb database. The solution must minimize the amount of time it takes to process daily ETL jobs.
Solution: You configure eight files for the tenpdb database. You place the files on a drive that contains the operating system files.
Does this meet the goal?

A. Yes
B. No

**Answer:** B

**NEW QUESTION 143**
.....

# Thank You for Trying Our Product

## We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

## 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 FirstTry

* 70-767 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year