Exam Questions 70-767

Implementing a SQL Data Warehouse (beta)

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

NEW QUESTION 1
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 2
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 deleteing.
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 3
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 ]
```

A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

```
UPDATE dbo.LastKeyByTable
SET   [ LastKey =? | LastKey = @A | LastKey = @B | LastKey = @LastKey ]
WHERE [ TableName =? | TableName = @A | TableName = @B | TableName = @TableName ]
```

NEW QUESTION 4
You are building a server to host a data warehouse.
The planned disk activity for-khe 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 5
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.

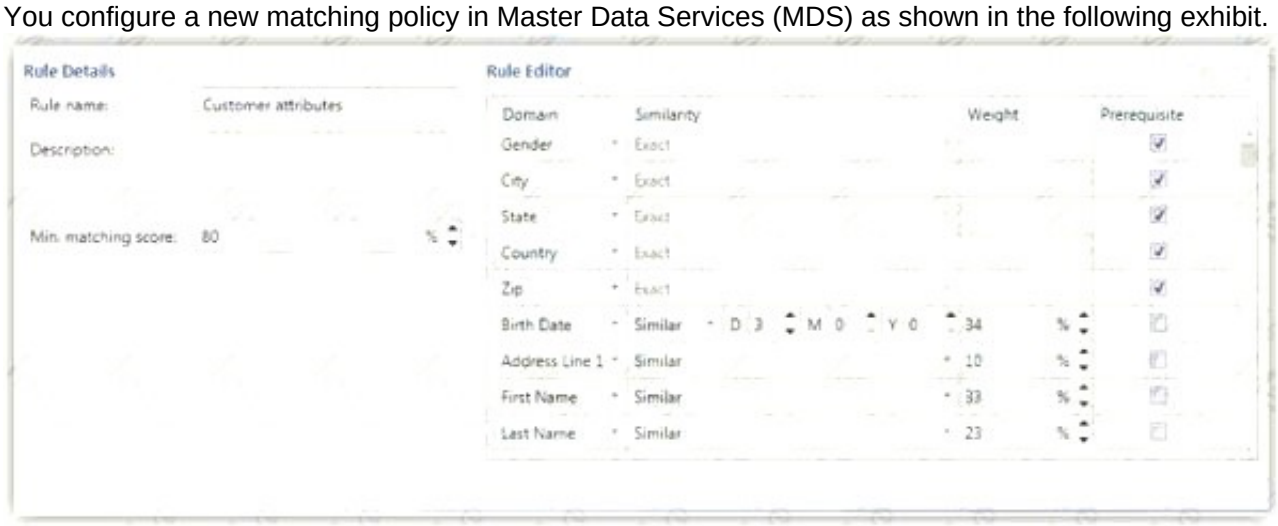

You review the Matching Results of the policy and find that the number of new values matches the new values.
You verify that the data contains multiple records that have similar address values, and you expect some of the records to match. You need to increase the likelihood that the records will match 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 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 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 7
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 8
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 9
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.

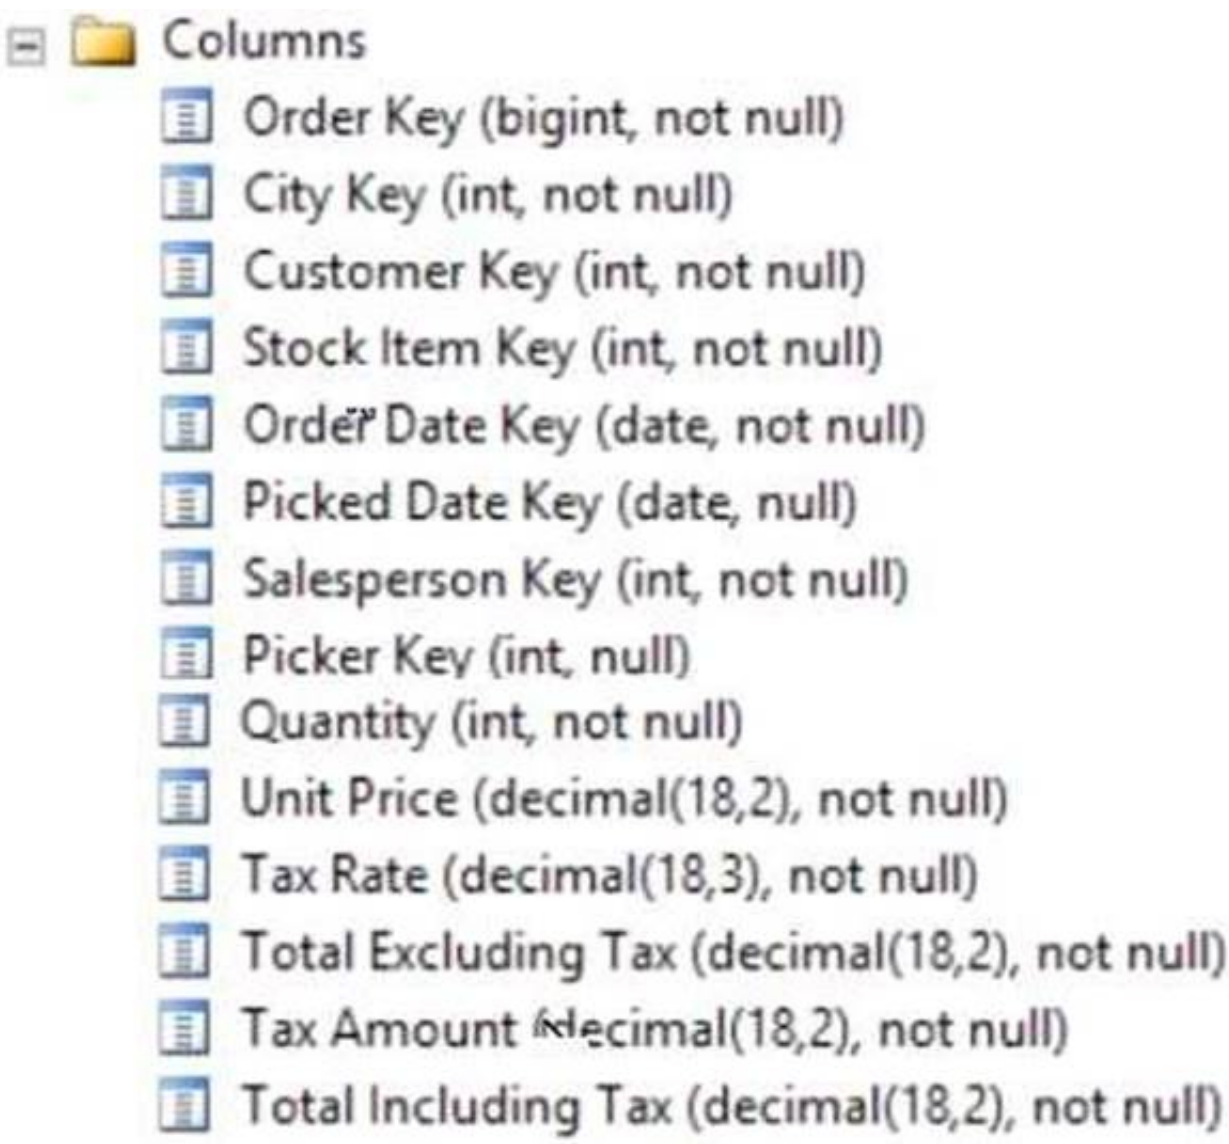


```
SELECT AVG([Tax Amount]) AS [Average Tax Amount]
FROM Fact.Order
WHERE [Order Date Key] BETWEEN '20150701' AND '20151231'

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 10
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 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.
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.)

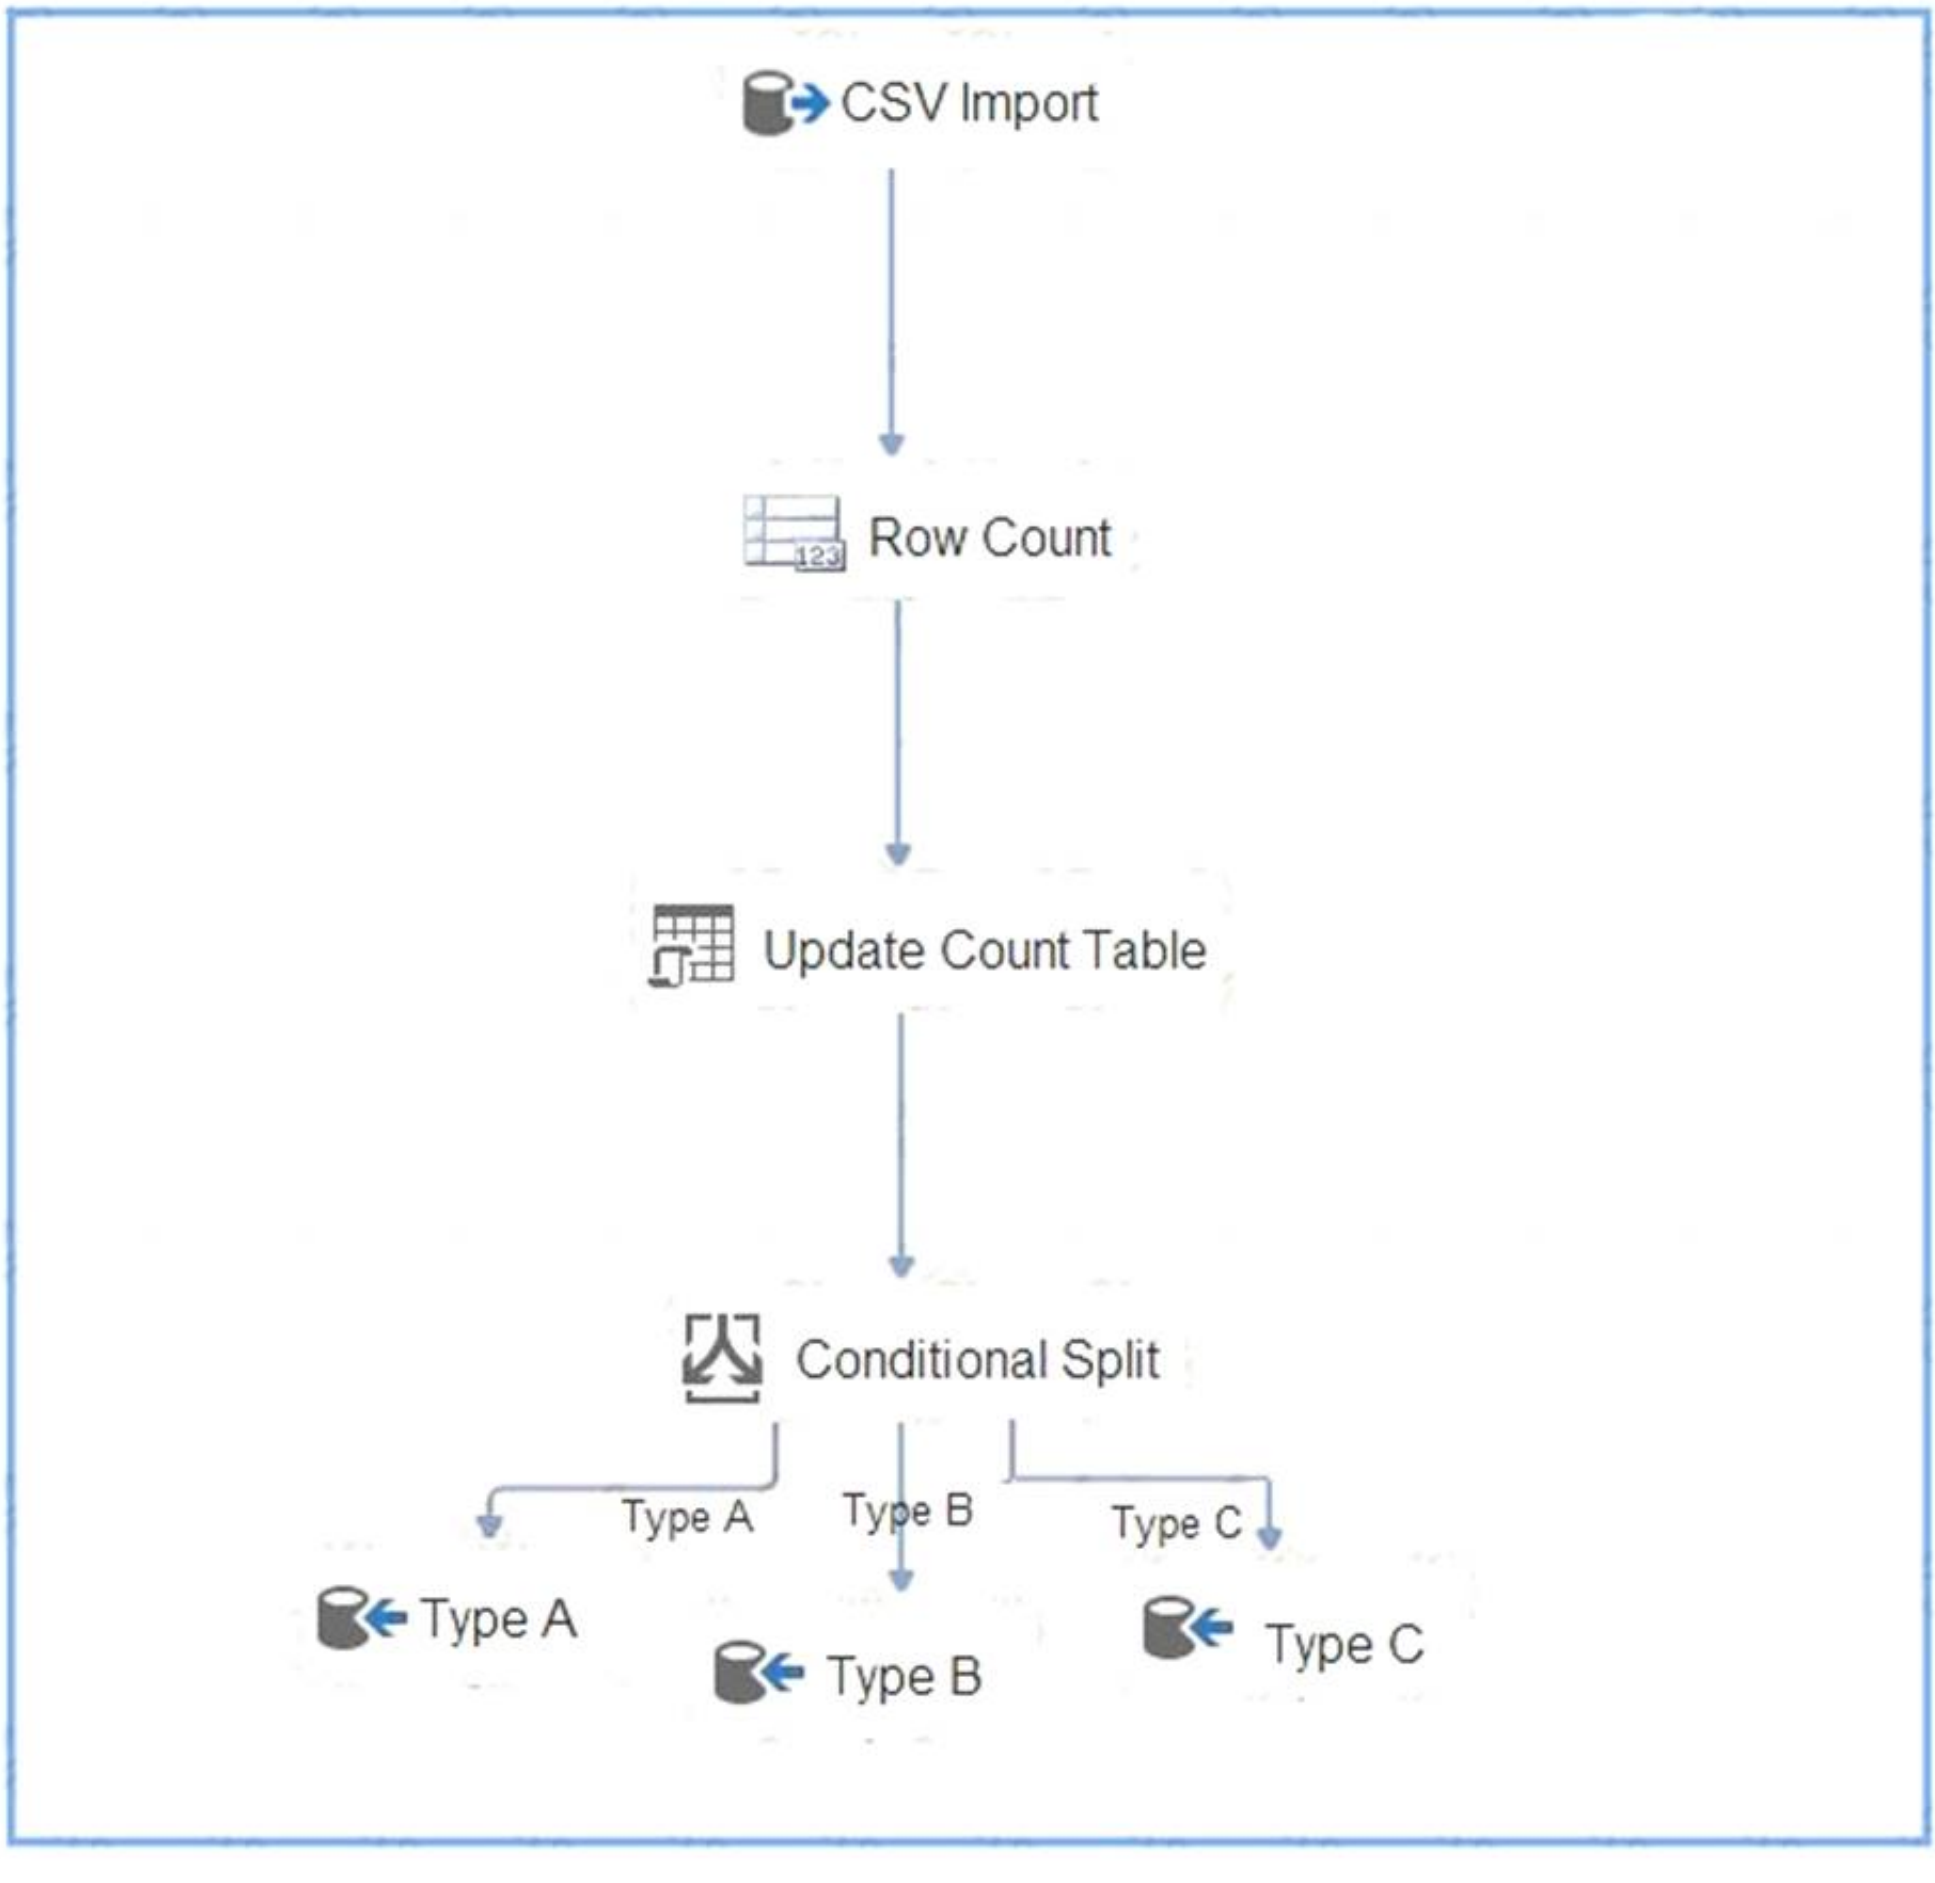

Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
The conditional split must be before the count.

NEW QUESTION 15
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 18
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.

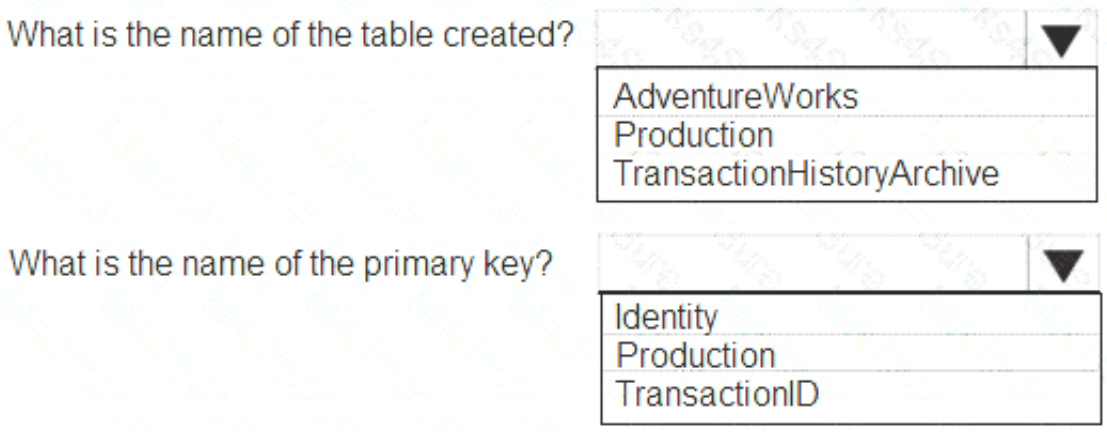

A. Mastered
B. Not Mastered

Answer: A

Explanation:

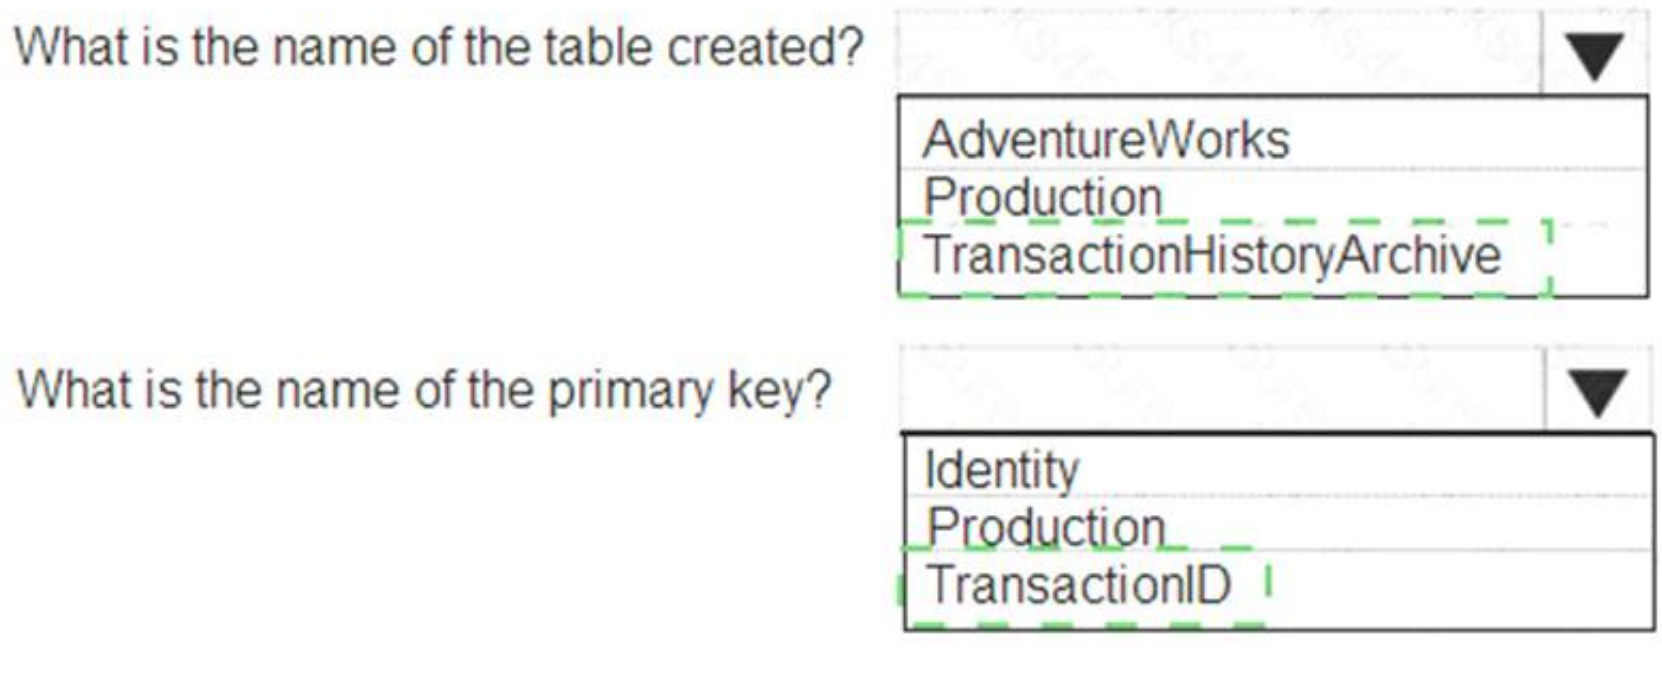

NEW QUESTION 23
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After 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.

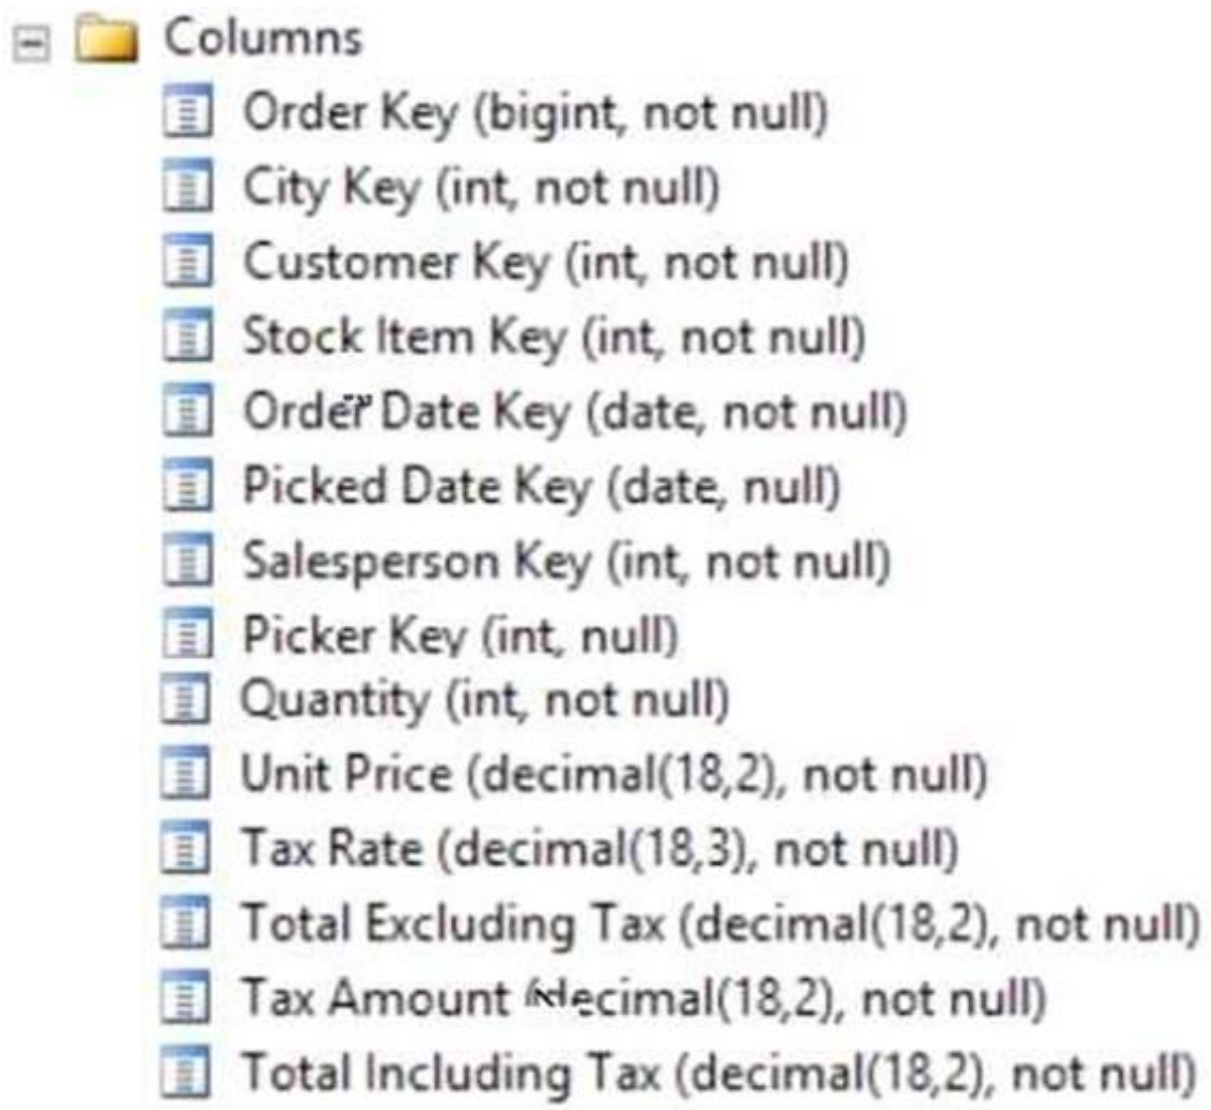

You must minimize the amount of space that indexes for the Fact.Order table consume. You run the following queries frequently. Both queries must be able to use a columnstore index:

```
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.
SolutionvYou 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 26
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 30
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 32
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.

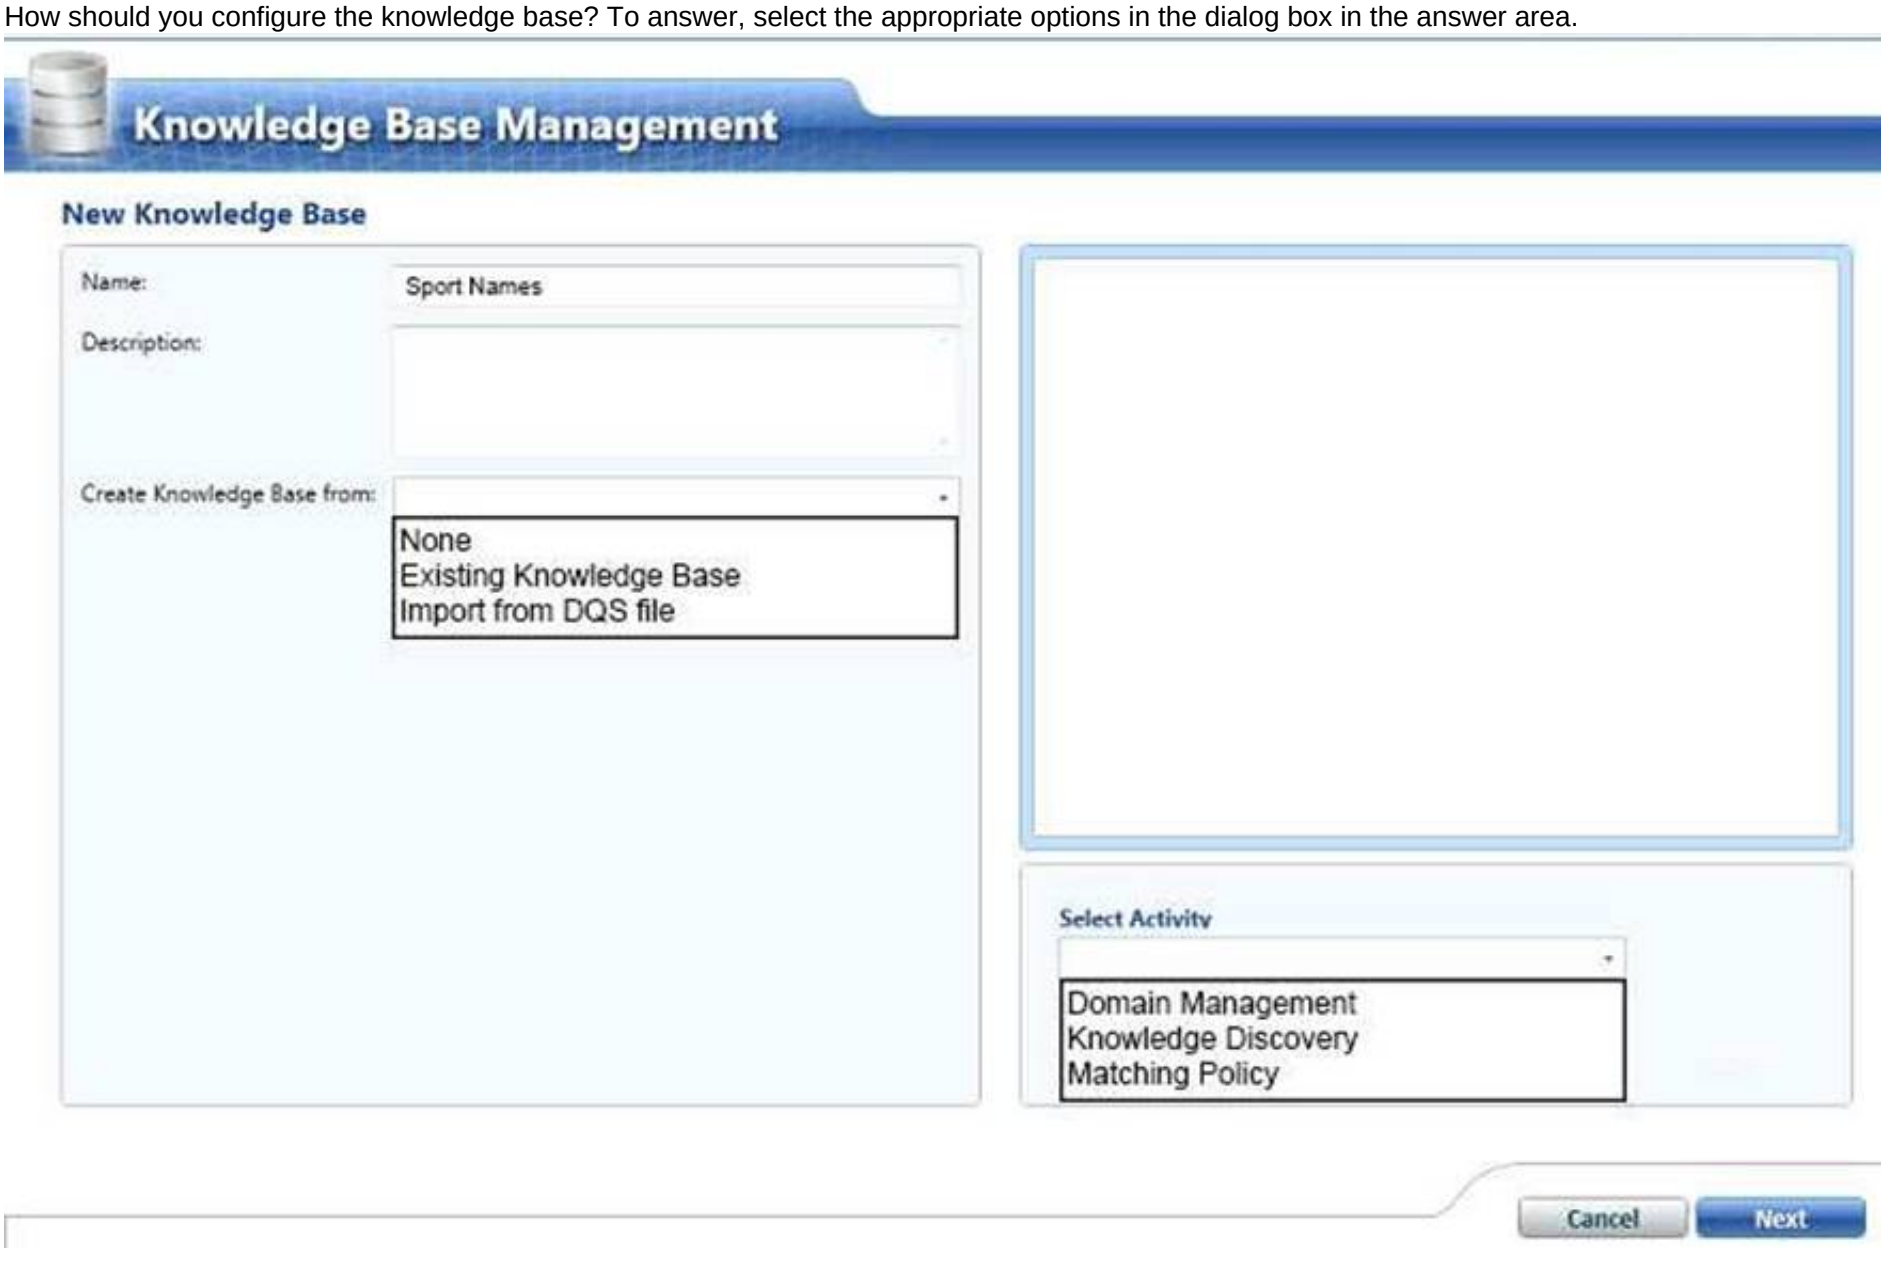

A. Mastered
B. Not Mastered

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 36
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 remove the Integration Services job.
You need to stop tracking changes to the database. The solution must remove all the change data capture configurations from DB1.
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: F

Explanation:
sys.sp_cdc_disable_db disables change data capture for all tables in the database currently enabled. All system objects related to change data capture, such as change tables, jobs, stored procedures and functions, are dropped.
References:
https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-disable-db-transa

NEW QUESTION 38
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 43
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 44
You create a Microsoft SQL Server Integration Services (SSIS) package as shown in the SSIS Package exhibit. (Click the Exhibit button.)

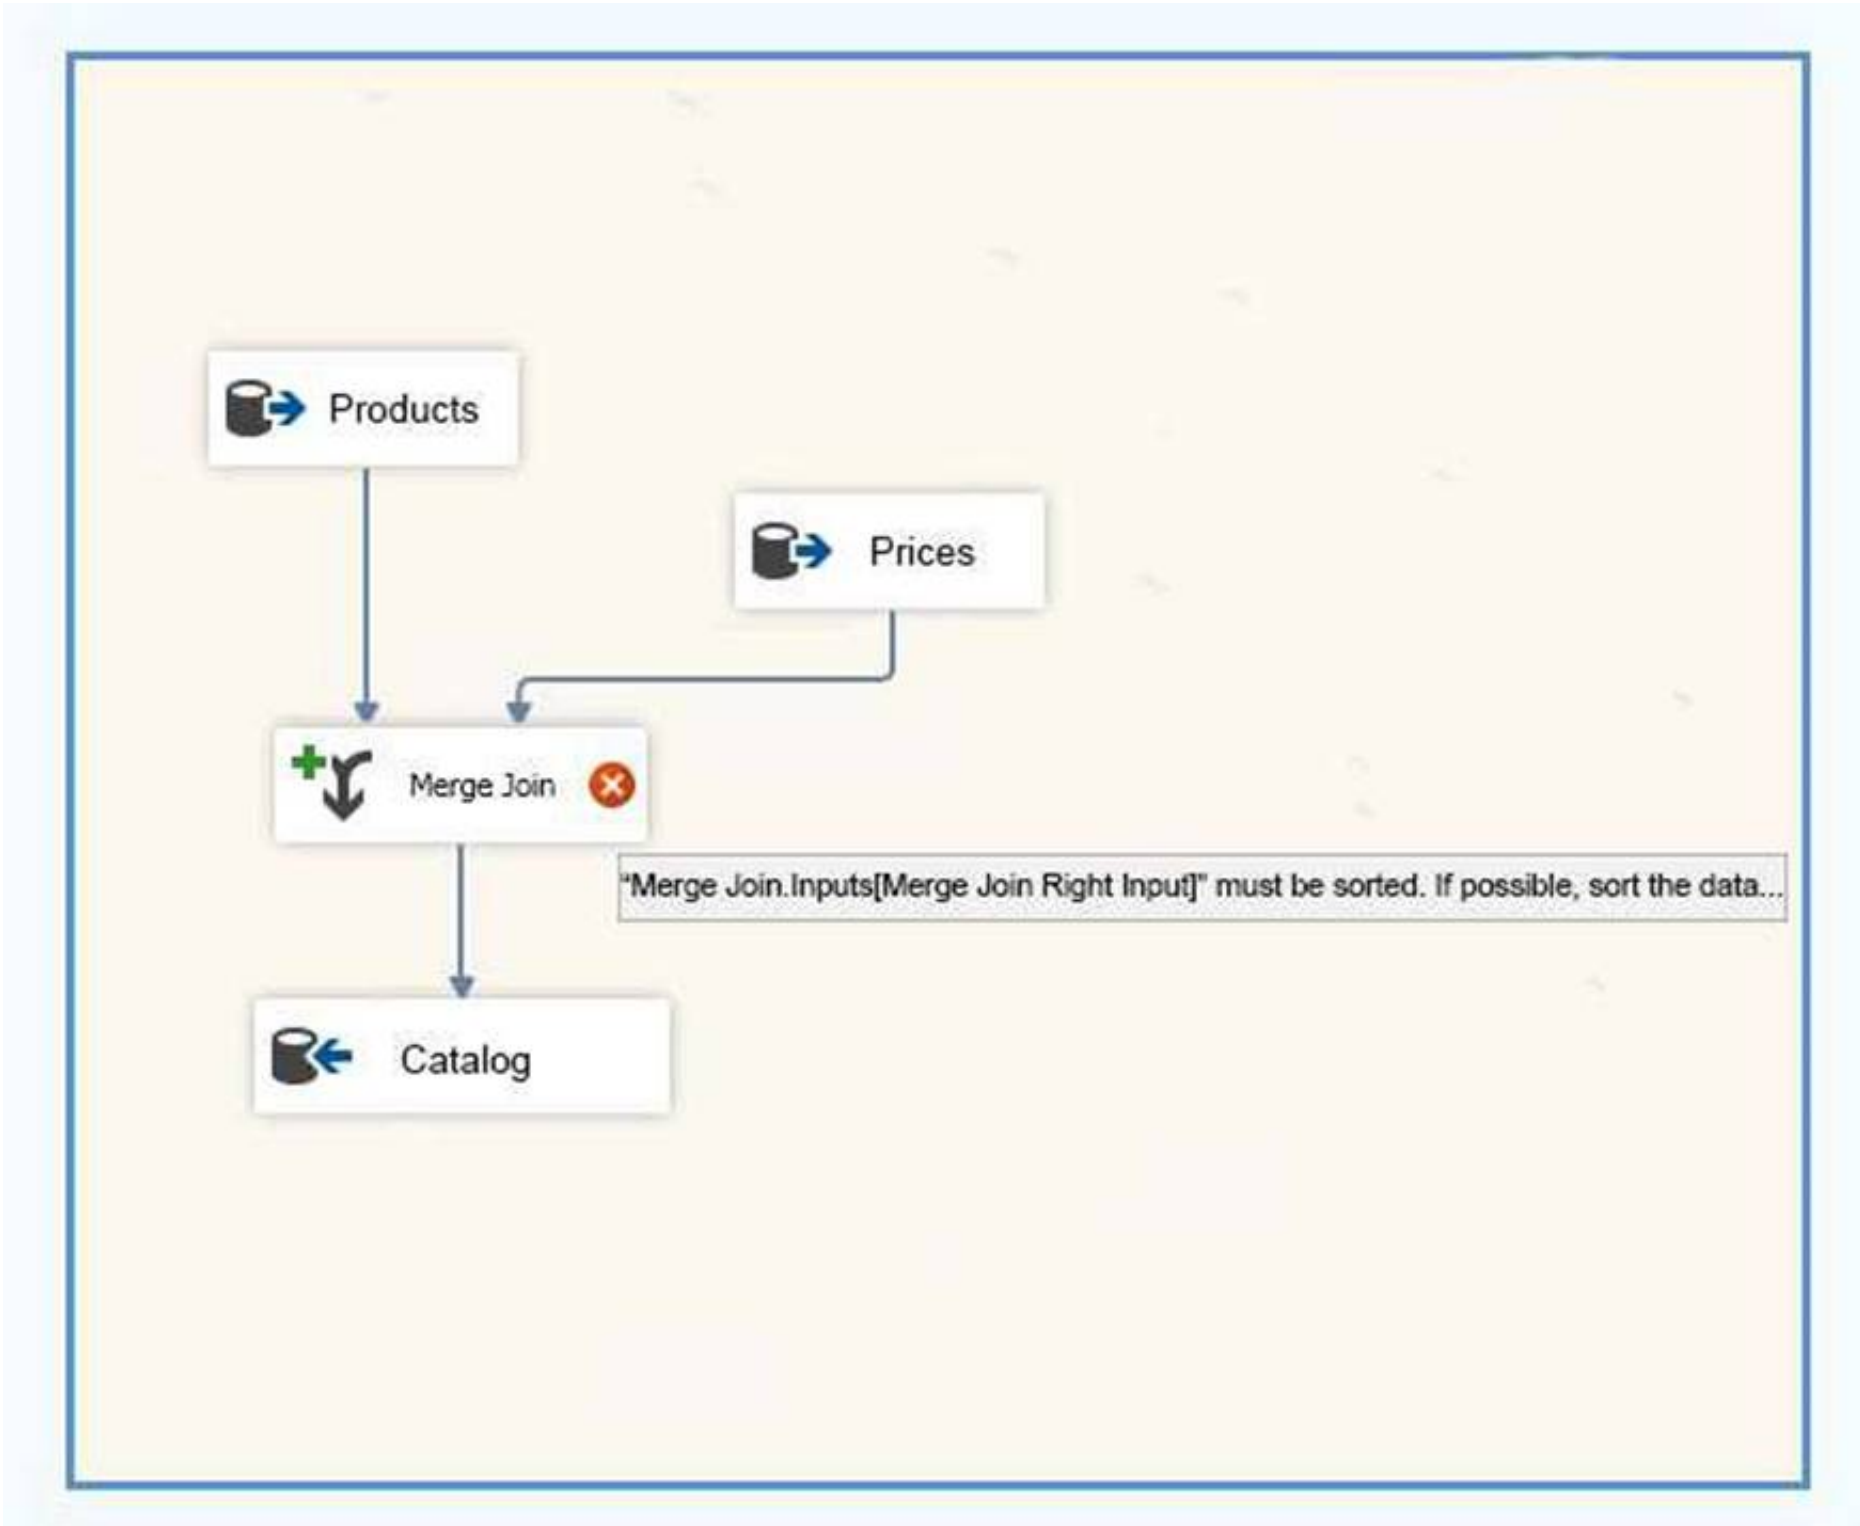

The package uses data from the Products table and the Prices table. Properties of the Prices source are shown in the OLE DB Source Editor exhibit (Click the Exhibit Button.) and the Advanced Editor for Prices exhibit (Click the Exhibit button.)

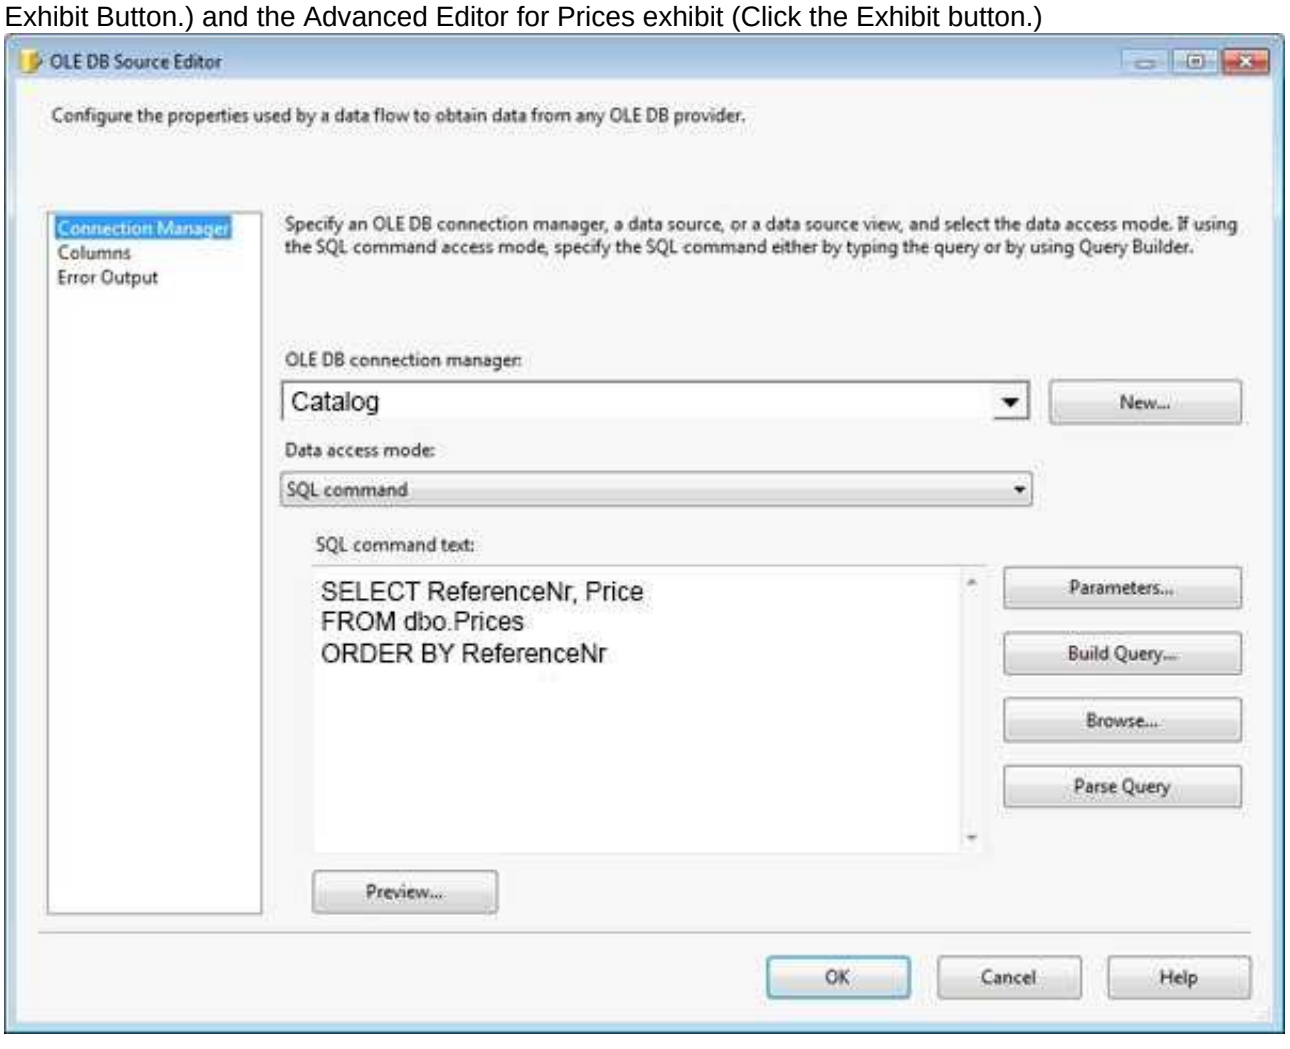

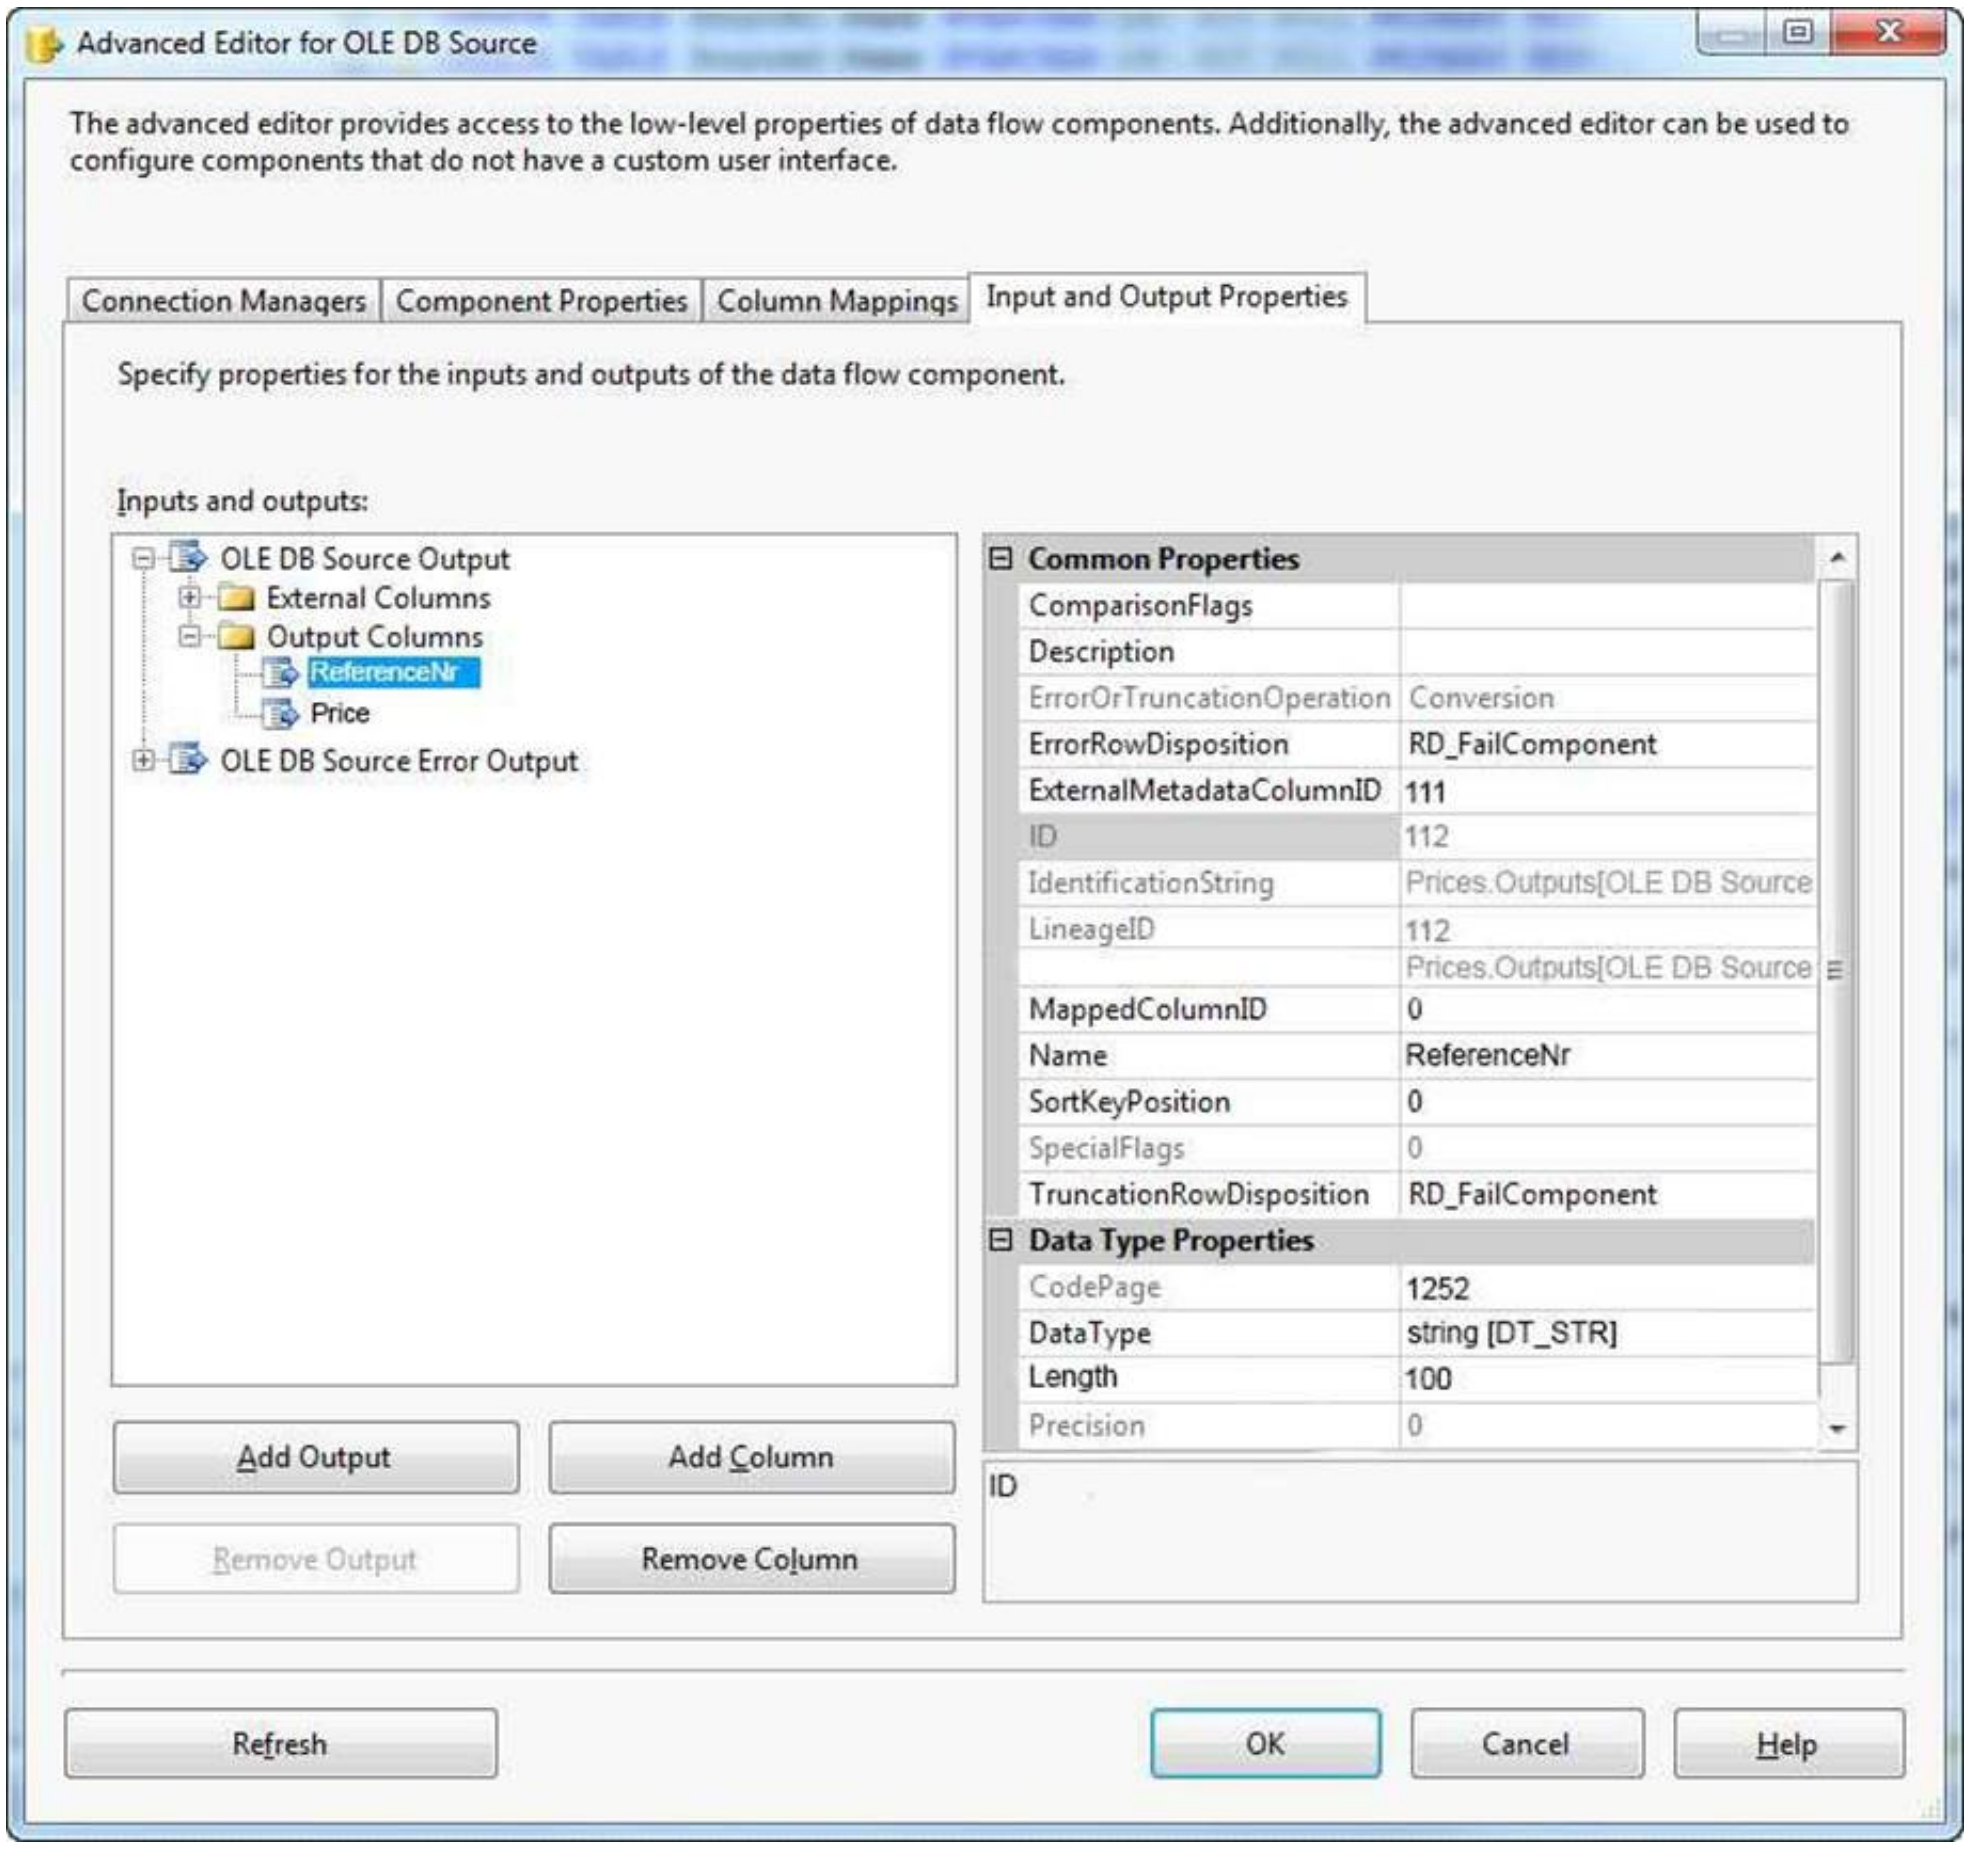

You join the Products and Prices tables by using the ReferenceNr column. You need to resolve the error with the package.
For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

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 49
You have a data quality project that focuses on the Products catalog for the company. The data includes a product reference number.
The product reference should use the following format: Two letters followed by an asterisk and then four or five numbers. An example of a valid number is XX*55522. Any reference number that does not conform to the format must be rejected during the data cleansing.
You need to add a Data Quality Services (DQS) domain rule in the Products domain. Which rule should you use?

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

Answer: A

Explanation:
For a pattern matching rule:
Any letter (A…Z) can be used as a pattern for any letter; case insensitive Any digit (0…9) can be used as a pattern for any digit
Any special character, except a letter or a digit, can be used as a pattern for itself Brackets, [], define optional matching
Example: ABC:0000
This rule implies that the data will contain three parts: any three letters followed by a colon (:), which is again followed by any four digits.

NEW QUESTION 51
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 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 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 60
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 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 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 64
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 67
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 remove the Integration Services job.
You need to stop tracking changes to the database. The solution must remove all the change data capture configurations from DB1.
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: F

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 are developing a Microsoft SQL Server Integration Services (SSIS) package.
You need to ensure that the packa
ge records the current Log Sequence Number (LSN) in the source database before the package begins reading source tables.
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: A

Explanation:
The CDC Control task is used to control the life cycle of change data capture (CDC) packages. It handles CDC package synchronization with the initial load package, the management of Log Sequence Number (LSN) ranges that are processed in a run of a CDC package.
References: https://docs.microsoft.com/en-us/sql/integration-services/control-flow/cdc-control-task

NEW QUESTION 75
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 80
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	376	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
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

A. Mastered
B. Not Mastered

Answer: A

Explanation:

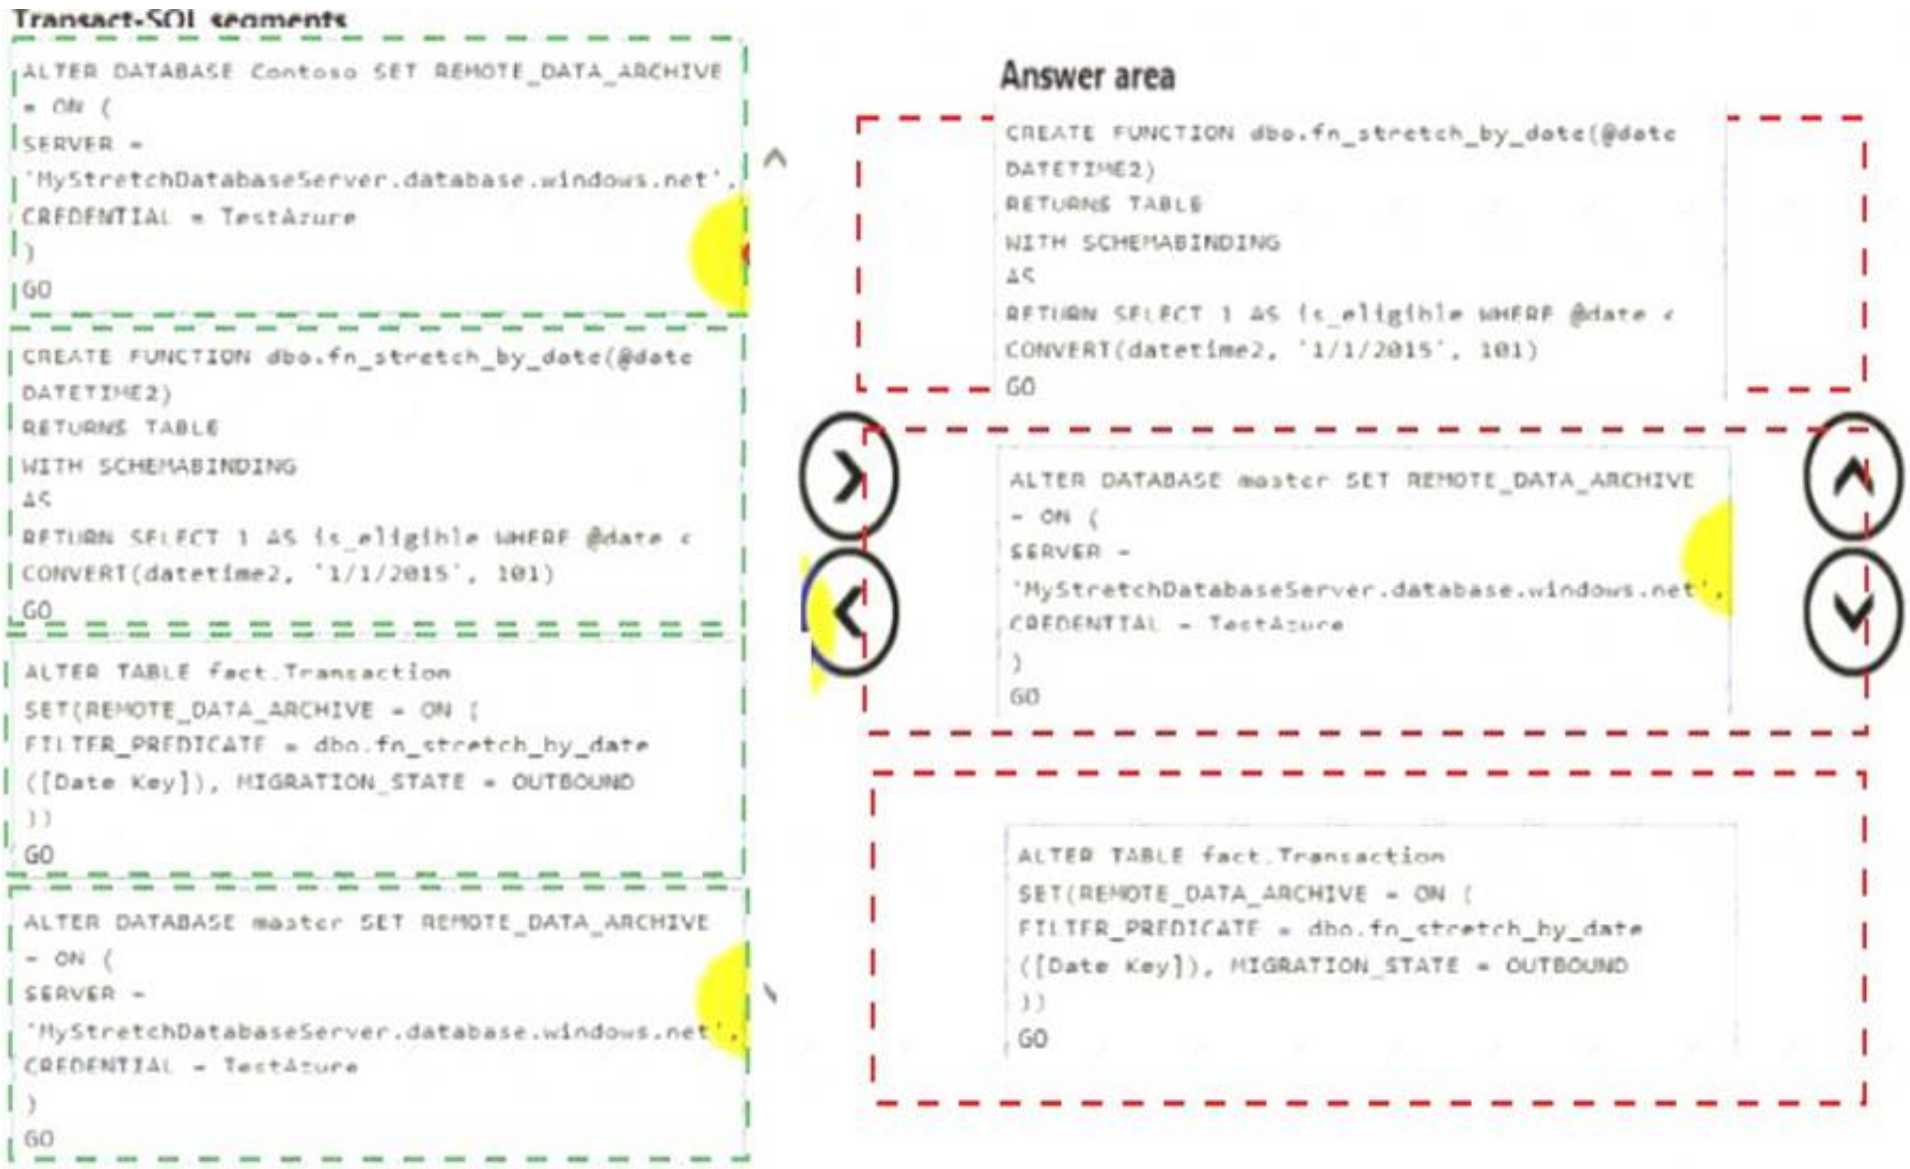

NEW QUESTION 82
Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question 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.
A new version of that integration Services package is released that introduces several errors in the loading process.
You need to roll back the Integration Services package to the previous version. 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

Answer: B

Explanation:
catalog.restore_project restores a project in the Integration Services catalog to a previous version. References:
https://docs.microsoft.com/en-us/sql/integration-services/system-stored-procedures/catalog-restore-project-ssisd

NEW QUESTION 85
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 86
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 91
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 92
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 96
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 the data load process for a data warehouse.
The data warehouse uses daily partitions to store data added or modified during the last 60 days. Older data is stored in monthly partitions.
You need to ensure that the ETL process can modify the partition scheme during the data load process. 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

Answer: E

NEW QUESTION 97
.....

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 FirstTry

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