

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 happen 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 until the data warehouse instance is resumed. |

What will happen 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

You are designing the data warehouse to import data from three different environments. The sources for the data warehouse will be loaded every hour.

Scenario A includes tables in a Microsoft Azure SQL Database:

- Millions of updates and inserts occur per hour
- A periodic query of the current state of rows that have changed is needed.
- The change detection method needs to be able to ignore changes to some columns in a table.
- The source database is a member of an AlwaysOn Availability group.

Scenario B includes tables with status update changes:

- Tracking the duration between workflow statuses.
- All transactions must be captured, including before/after values for UPDATE statements.
- To minimize impact to performance, the change strategy adopted should be asynchronous.

Scenario C includes an external source database:

- Updates and inserts occur regularly.
- No changes to the database should require code changes to any reports or applications.
- Columns are added and dropped to tables in the database periodically. These schema changes should not require any interruption or reconfiguration of the change detection method chosen.
- Data is frequently queried as the entire row appeared at a past point in time. All tables have primary keys.

You need to load each data source. You must minimize complexity, disk storage, and disruption to the data sources and the existing data warehouse.

Which change detection method should you use for each scenario? To answer, drag the appropriate loading methods to the correct scenarios. Each source may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer Area

| Loading methods | Scenario | Loading method |
|---------------------------------|----------|----------------|
| Change Tracking | A | |
| Change Data Capture | B | |
| System-Versioned Temporal Table | C | |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

| Scenario | Loading method |
|----------|---------------------------------|
| A | System-Versioned Temporal Table |
| B | Change Tracking |
| C | Change Data Capture |

Box A: System-Versioned Temporal Table

System-versioned temporal tables are designed to allow users to transparently keep the full history of changes for later analysis, separately from the current data, with the minimal impact on the main OLTP workload.

Box B: Change Tracking Box C: Change Data Capture

Change data capture supports tracking of historical data, while that is not supported by change tracking. References:

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

NEW QUESTION 3

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

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

You are developing a Microsoft SQL Server Integration Services (SSIS) projects. The project consists of several packages that load data warehouse tables.

You need to extend the control flow design for each package to use the following control flow while minimizing development efforts and maintenance:



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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 4

Your company has a Microsoft SQL Server data warehouse instance. The human resources department assigns all employees a unique identifier. You plan to

store this identifier in a new table named Employee.

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

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

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

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

Answer Area

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

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

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

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

NEW QUESTION 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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have a Microsoft Azure SQL Data Warehouse instance that must be available six months a day for reporting. You need to pause the compute resources when the instance is not being used. Solution: You use SQL Server Configuration Manager. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

To pause a SQL Data Warehouse database, use any of these individual methods. Pause compute with Azure portal
 Pause compute with PowerShell
 Pause compute with REST APIs
 References:
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-compute-overview>

NEW QUESTION 6

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 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 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 8

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

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

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

What should you do?

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

Answer: D

NEW QUESTION 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.

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 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 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 10

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

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

You have a 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 12

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.

Transact-SQL segments

Answer Area

```
EXEC sys.sp_cdc_enable_table
@source_schema = N 'schema',
@source_name   = N 'Dimension.Customer',
@role_name     = NULL,
@supports_net_changes = 1
```

```
EXEC sys.sp_cdc_enable_db
```

```
USE DB2
```

```
EXEC sys.sp_cdc_enable_table
```

```
USE DB1
```

```
EXEC sys.sp_cdc_enable_db
@source_schema = N 'schema',
@source_name   = N 'Dimension.Customer',
@role_name     = NULL,
@supports_net_changes = 1
```



- 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 16

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 improve the quality of 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.

Actions

Add an external link from the Data Quality Service (DQS) knowledge base to the SQL Server table of replacement values.

Publish the knowledge base and modify the ETL package to call it by using a .NET Script transformation.

Create a Data Quality Service (DQS) knowledge base.

Add a Lookup transformation to the ETL package to replace incorrect values.

Import the raw data from the users to perform discovery.

Create a table to store a list of incorrect values and the correct values to which they should map.

Publish the knowledge base and modify the ETL package to call it by using a Data Quality Service (DQS) Client transformation.

Map alternative values for entries that have been identified as being incorrectly entered.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 21

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.

| Actions | Answer Area |
|------------------------------------|--|
| Perform Network Discovery. |  |
| Configure a matching policy. | |
| Configure reference data services. | |
| Perform Domain Management. | |
| Perform Knowledge Discovery. | |

- 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 26

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

You have a Microsoft SQL Server data warehouse instance that supports several client applications. The data warehouse includes the following tables:

Dimension.SalesTerritory, Dimension.Customer,

Dimension.Date, Fact.Ticket, and Fact.Order. The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. The Fact.Order table is optimized for weekly reporting, but the company wants to change it daily. The Fact.Order table is loaded by using an ETL process. Indexes have been added to the table over time, but the presence of these indexes slows data loading.

All data in the data warehouse is stored on a shared SAN. All tables are in a database named DB1. You have a second database named DB2 that contains copies of production data for a development environment. The data warehouse has grown and the cost of storage has increased. Data older than one year is accessed infrequently and is considered historical.

You have the following requirements:

- Implement table partitioning to improve the manageability of the data warehouse and to avoid the need to repopulate all transactional data each night. Use a partitioning strategy that is as granular as possible.
- Partition the Fact.Order table and retain a total of seven years of data.
- Partition the Fact.Ticket table and retain seven years of data. At the end of each month, the partition structure must apply a sliding window strategy to ensure that a new partition is available for the upcoming month, and that the oldest month of data is archived and removed.
- Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.
- Maximize the performance during the data loading process for the Fact.Order partition.
- Ensure that historical data remains online and available for querying.
- Reduce ongoing storage costs while maintaining query performance for current data. You are not permitted to make changes to the client applications.

You need to implement the data partitioning strategy. How should you partition the Fact.Order table?

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

Answer: C

Explanation:

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

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

NEW QUESTION 28

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

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

Answer: A

NEW QUESTION 30

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

You have a Microsoft SQL Server data warehouse instance that supports several client applications. The data warehouse includes the following tables:

Dimension.SalesTerritory, Dimension.Customer, Dimension.Date, Fact.Ticket, and Fact.Order. The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. The Fact.Order table is optimized for weekly reporting, but the company wants to change it daily. The Fact.Order table is loaded by using an ETL process. Indexes have been added to the table over time, but the presence of these indexes slows data loading.

All data in the data warehouse is stored on a shared SAN. All tables are in a database named DB1. You have a second database named DB2 that contains copies of production data for a development environment. The data warehouse has grown and the cost of storage has increased. Data older than one year is accessed infrequently and is considered historical.

You have the following requirements:

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

You are not permitted to make changes to the client applications. You need to implement partitioning for the Fact.Ticket table.

Which three actions should you perform in sequence? To answer, drag the appropriate actions to the correct locations. Each action may be used once, more than once or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: More than one combination of answer choices is correct. You will receive credit for any of the correct combinations you select.

Actions

- INSERT SELECT
- MERGE
- SWITCH
- DELETE
- SPLIT

Answer area

| First action | Second action |
|--------------|---------------|
| Action | |
| Action | Action |

- 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 31

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 35

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

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

You have the following line-of-business solutions:

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

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

The database for each solution has a table named Products that stored product information. The Products table in each database uses a separate and unique key for product records. Each table shares a column named ReferenceNr between the databases. This column is used to create queries that involve more than once solution.

You need to load data from the individual solutions into the data warehouse nightly. The following requirements must be met:

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

Solution: Perform the following actions: Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 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 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 37

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.

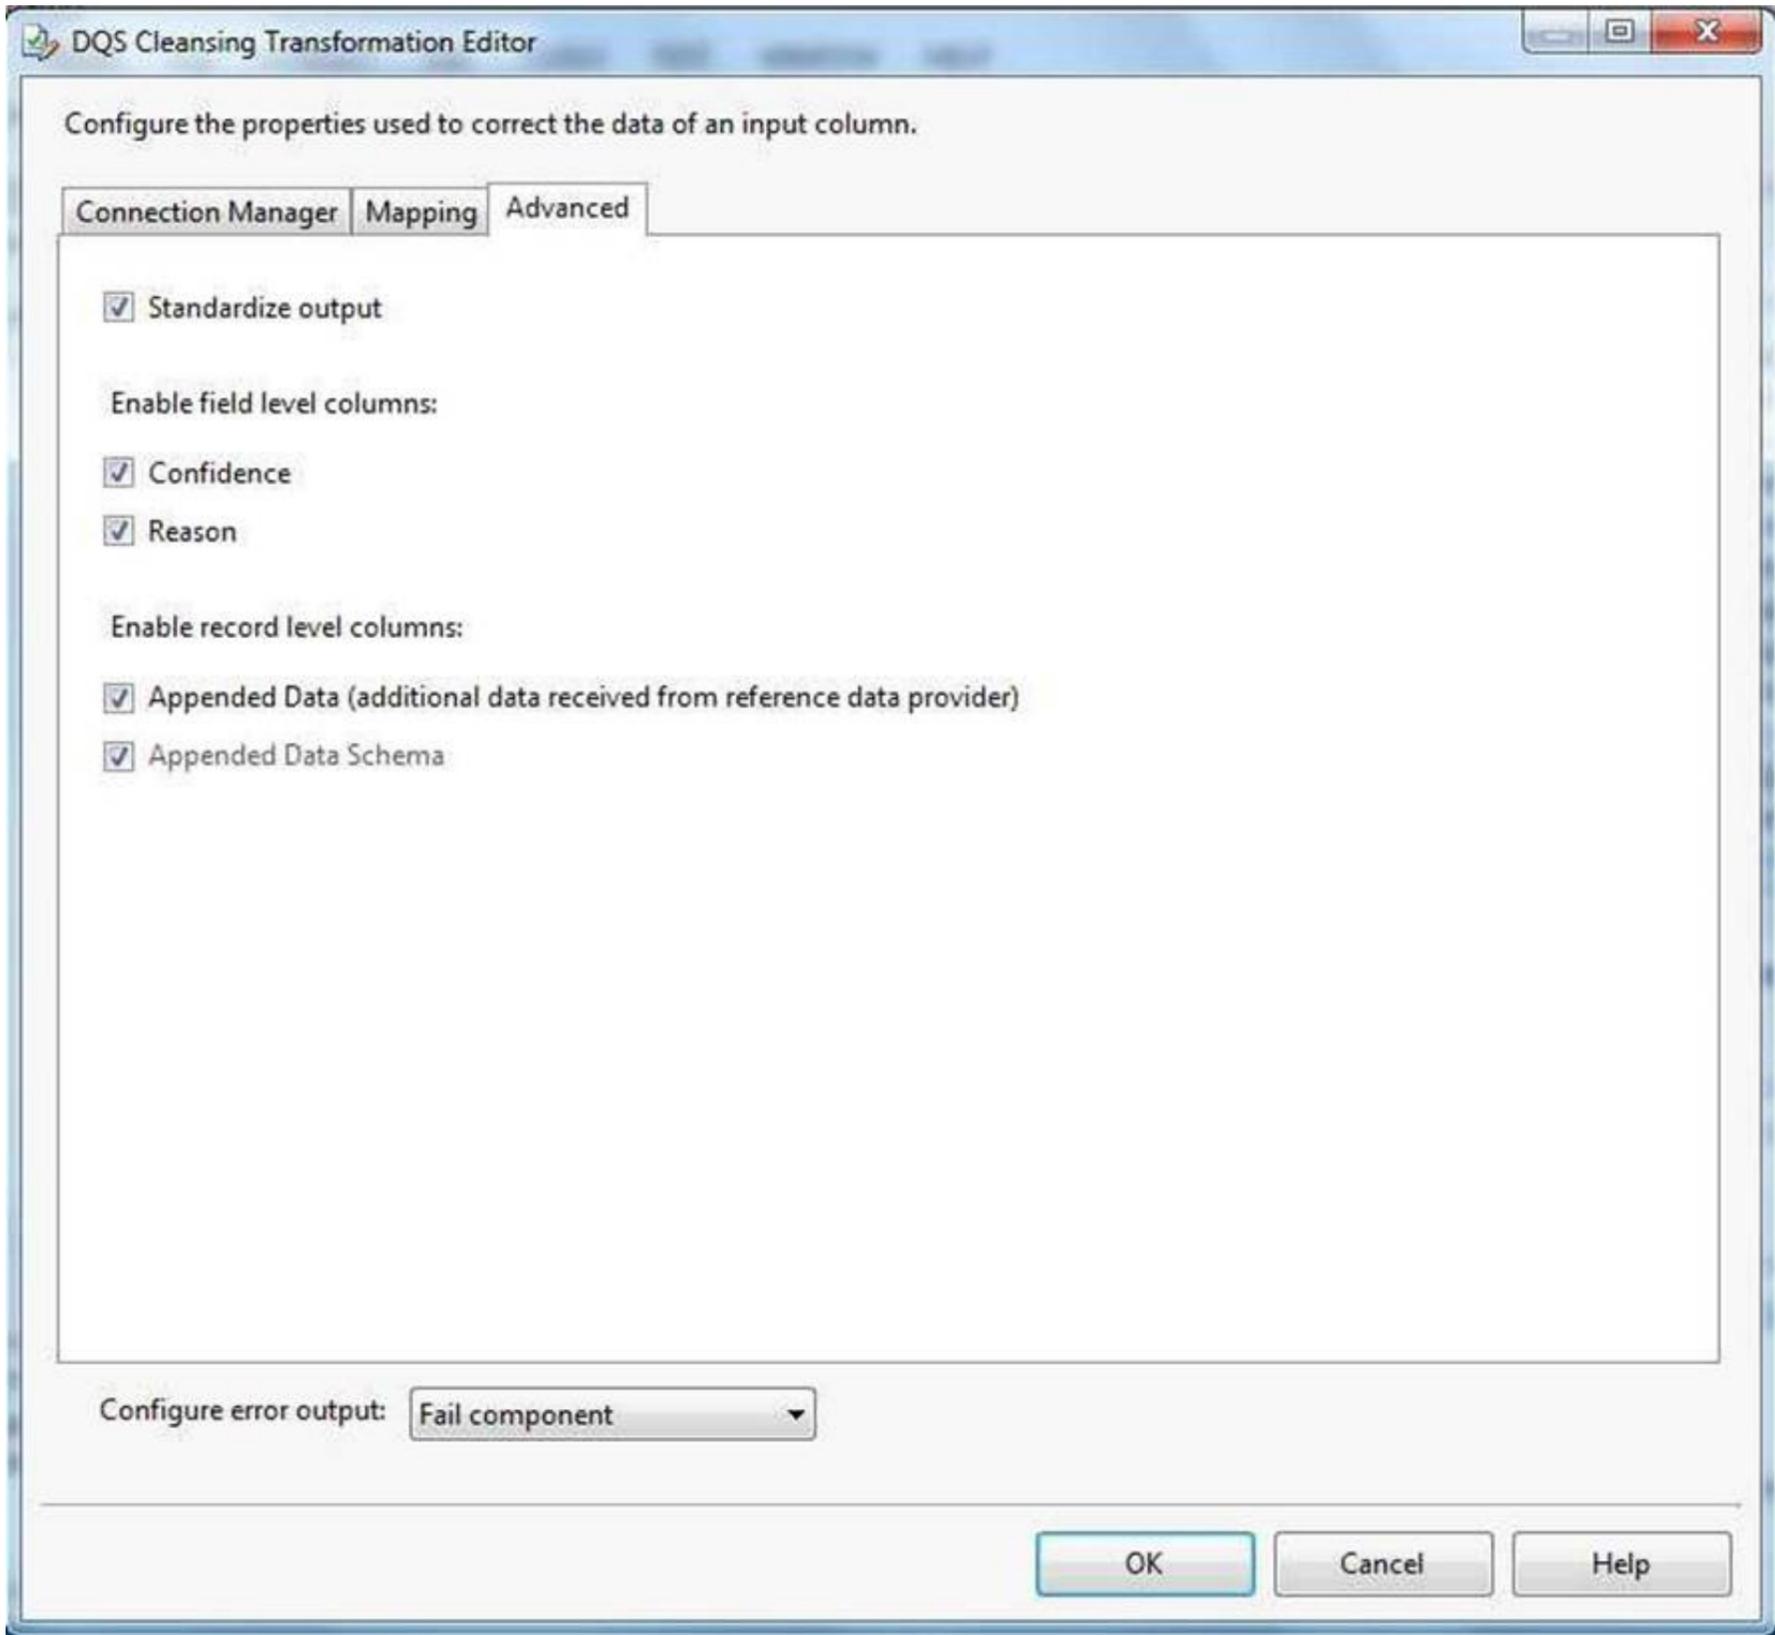
References: <https://docs.microsoft.com/en-us/sql/integration-services/control-flow/control-flow>

NEW QUESTION 39

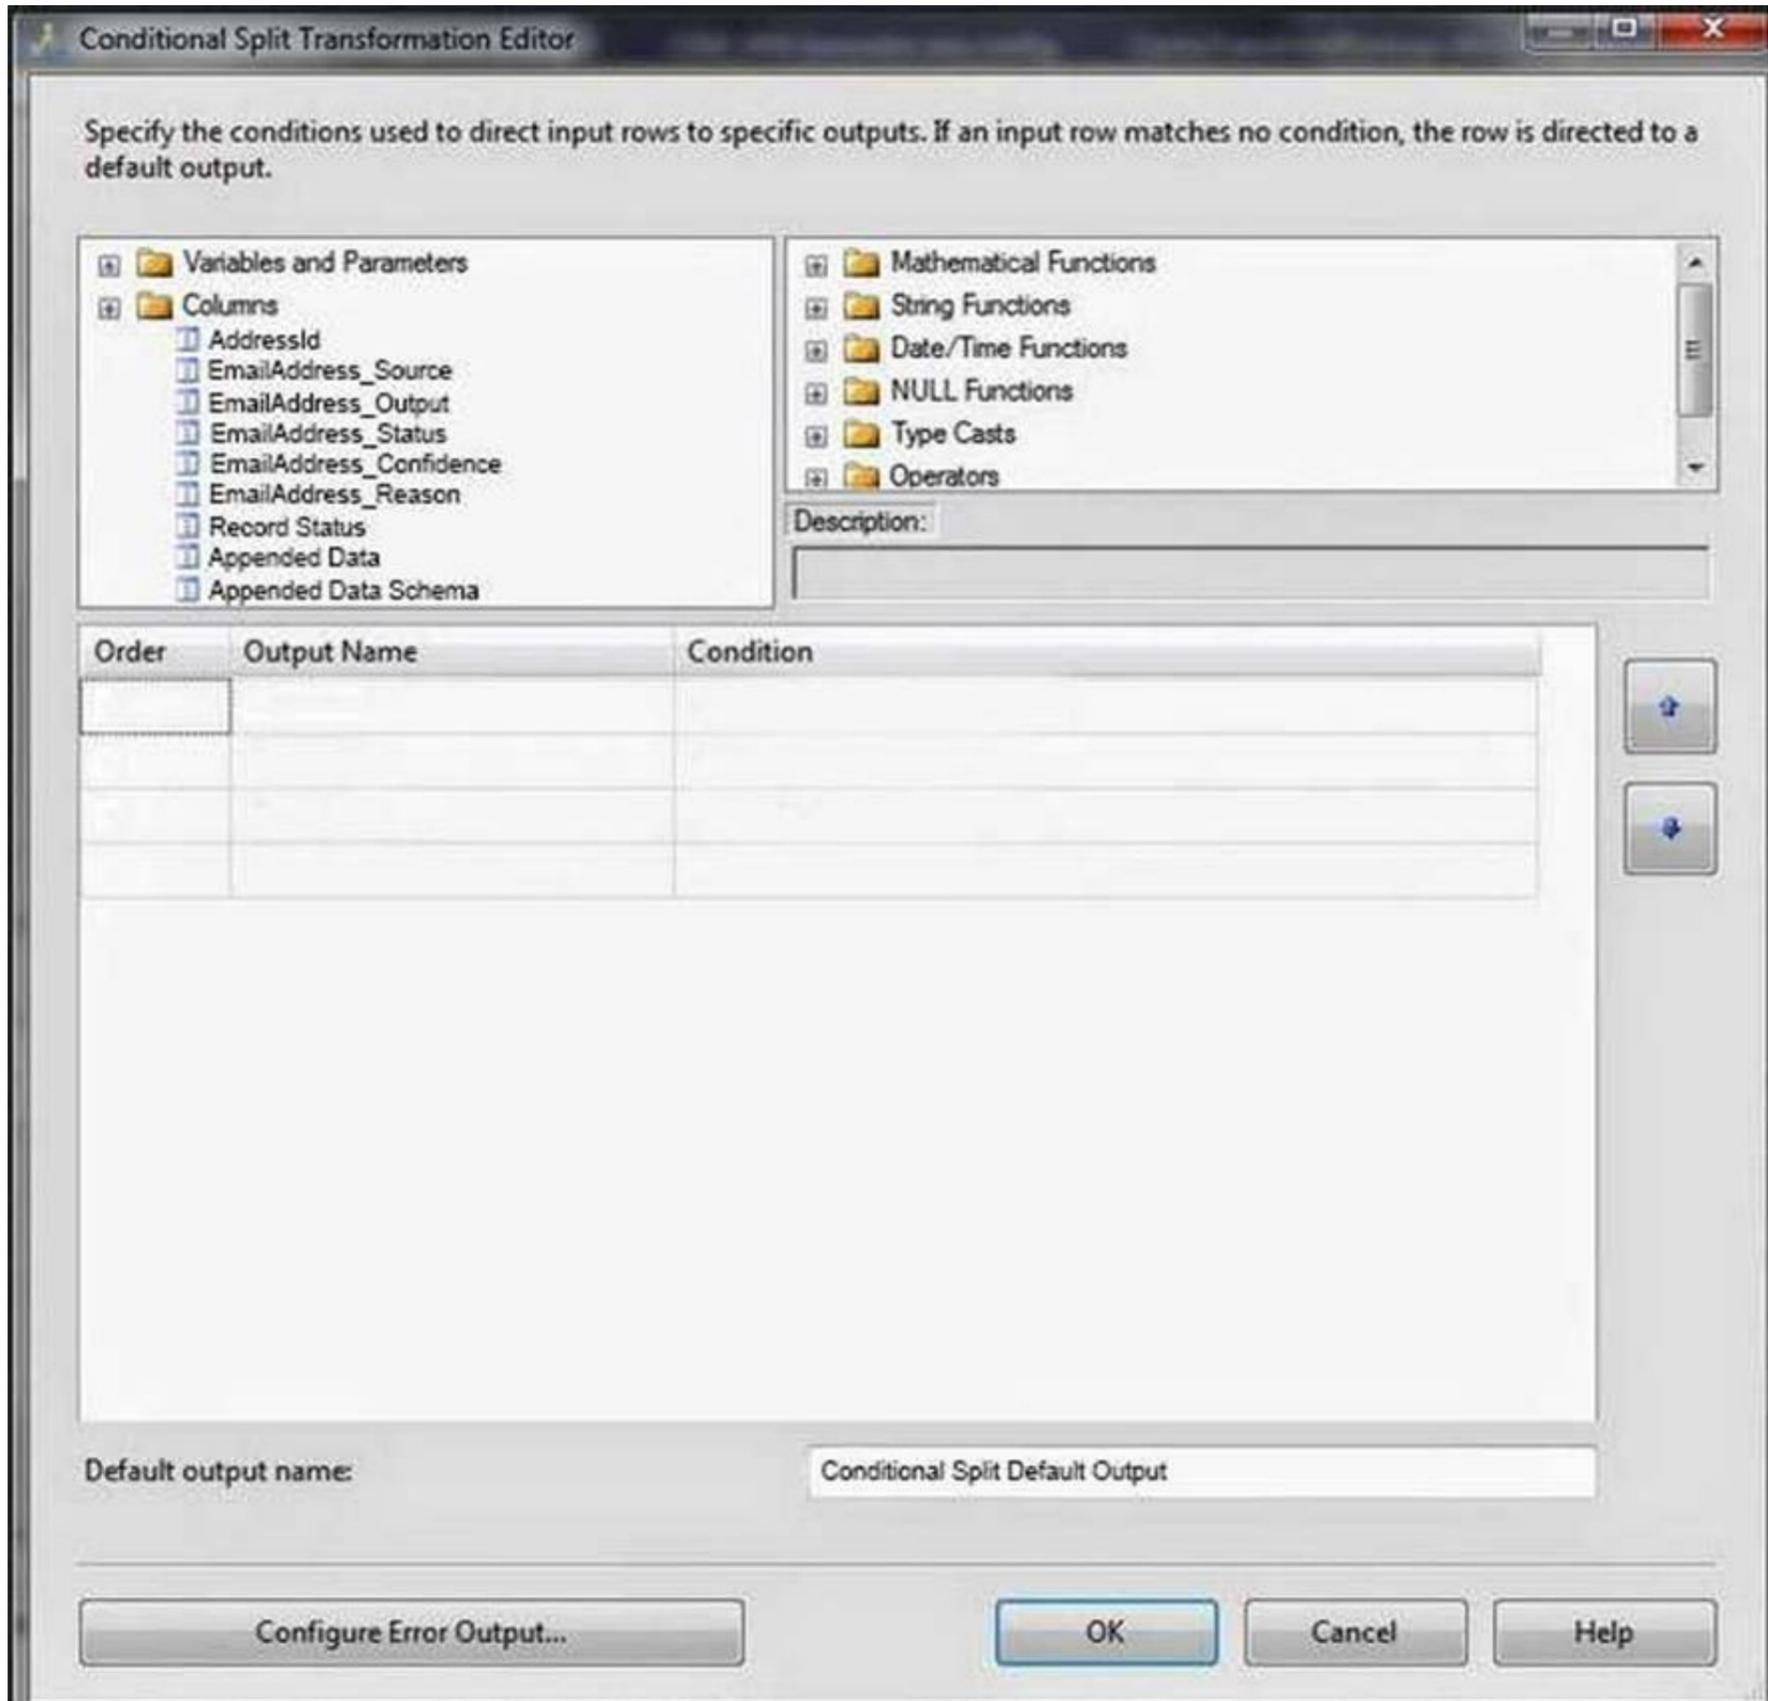
You have a Microsoft SQL Server Integration Services (SSIS) package that contains a Data Flow task as shown in the Data Flow exhibit. (Click the Exhibit button.)



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



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



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

Answer Area

Yes

No

You can use the EmailAddress_Output column to split the output.

You can use the EmailAddress_Status column to split the output.

You can use the EmailAddress_Reason column to split the output.

- 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. In such cases, it is recommended to use a Conditional Split component below the DQS Cleansing component, and configure it to split the records to groups based on the record status (or based on other columns such as specific field status).

References: <https://blogs.msdn.microsoft.com/dqs/2011/07/18/using-the-ssis-dqs-cleansing-component/>

NEW QUESTION 41

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

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

| Transact-SQL segments | | Answer Area | |
|--------------------------------|---------------------------------|----------------------|-------------------------------------|
| <input type="checkbox"/> BULK | <input type="checkbox"/> INSERT | <input type="text"/> | <input type="text"/> Sales.Invoices |
| <input type="checkbox"/> FROM | <input type="checkbox"/> WITH | <input type="text"/> | '\\share\data\file1.csv' |
| <input type="checkbox"/> MERGE | | <input type="text"/> | (FORMAT = 'CSV') |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

References:

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

NEW QUESTION 43

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

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

You have the following line-of-business solutions:

- ERP system
- Online WebStore
- Partner extranet

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

The database for each solution has a table named Products that stored product information. The Products table in each database uses a separate and unique key for product records. Each table shares a column named ReferenceNr between the databases. This column is used to create queries that involve more than once solution.

You need to load data from the individual solutions into the data warehouse nightly. The following requirements must be met:

- If a change is made to the ReferenceNr column in any of the sources, set the value of IsDisabled to True and create a new row in the Products table.
- If a row is deleted in any of the sources, set the value of IsDisabled to True in the data warehouse. Solution: Perform the following actions:
 - Enable the Change Tracking for the Product table in the source databases.
 - Query the CHANGETABLE function from the sources for the updated rows.
 - Set the IsDisabled column to True for the listed rows that have the old ReferenceNr value.
 - Create a new row in the data warehouse Products table with the new ReferenceNr value.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 46

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

Start of repeated scenario

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

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

The data warehouse includes the following tables:

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

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

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

Pal to create a measure that calculates the profit margin based on the existing measures.

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

End of repeated scenario

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

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Transact-SQL segments

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

Answer area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Transact-SQL segments

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

ALTER SCHEMA Dimension TRANSFER dbo.City

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

SELECT *
INTO Dimension.City
FROM dbo.dim_City

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

Answer area

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

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

NEW QUESTION 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 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 52

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 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 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.

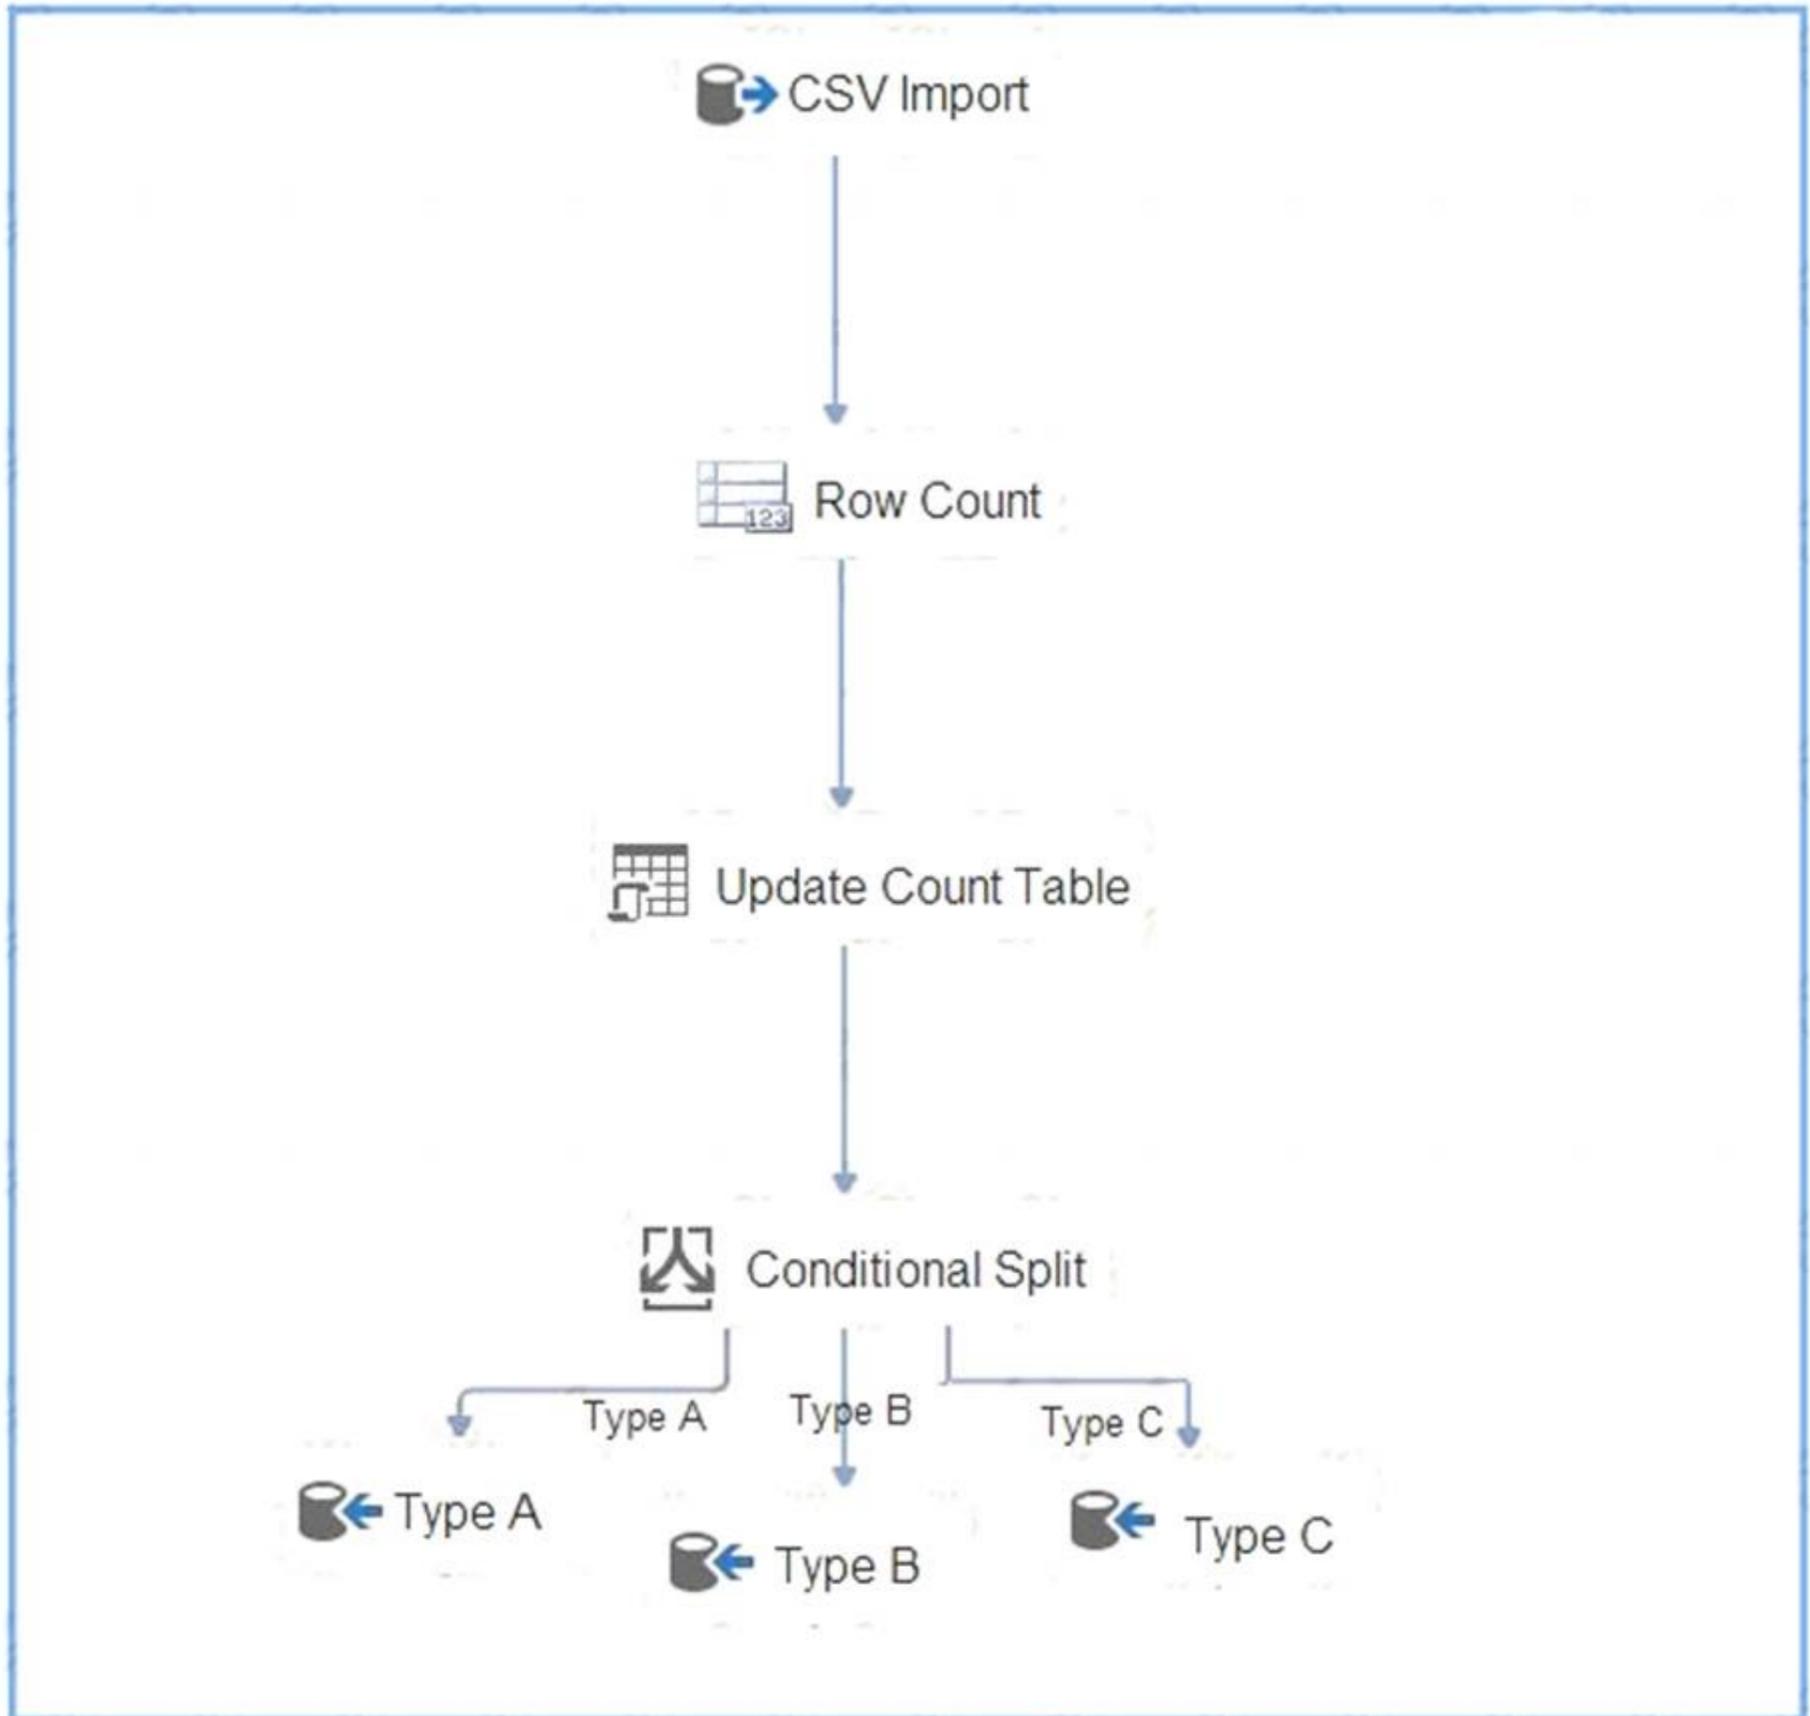
Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 61

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 66

You have a Microsoft SQL Server Integration Services (SSIS) package that loads data into a data warehouse each night from a transactional system. The package also loads data from a set of Comma-Separated Values (CSV) files that are provided by your company's finance department. The SSIS package processes each CSV file in a folder. The package reads the file name for the current file into a variable and uses that value to write a log entry to a database table. You need to debug the package and determine the value of the variable before each file is processed. Which four actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the

correct order.

Actions

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

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

Open the Control Flow editor for the package.

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

Set a breakpoint on the For Loop container.

Set a breakpoint on the Sequence container.

Open the Data Flow editor for the package.

Set a breakpoint on the Foreach Loop container.

Answer Area



- 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 71

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

You have a database named DB1.

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

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

Answer: D

Explanation:

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

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

Note: sys.sp_cdc_change_job

References:

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

NEW QUESTION 74

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 75

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 ALTER TABLE 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 80

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

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

You plan to deploy a Microsoft SQL server that will host a data warehouse named DB1. The server will contain four SATA drives configured as a RAID 10 array.

You need to minimize write contention on the transaction log when data is being loaded to the database. Solution: You replace the SATA disks with SSD disks.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

A data warehouse is too big to store on an SSD.

Instead you should place the log file on a separate drive. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on->

NEW QUESTION 84

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

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

Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

What is the name of the table created?

- AdventureWorks
- Production
- TransactionHistoryArchive

What is the name of the primary key?

- Identity
- Production
- TransactionID

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

What is the name of the table created?

- AdventureWorks
- Production
- TransactionHistoryArchive

What is the name of the primary key?

- Identity
- Production
- TransactionID

NEW QUESTION 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

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

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

The data warehouse includes the following tables:

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

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

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

Pal to create a measure that calculates the profit margin based on the existing measures.

You must implement a partitioning scheme 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.

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

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

End of repeated scenario

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

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Transact-SQL statements

- Description
- [Sale Key] ASC
- [Sale Key] DESC
- Description ASC, [Sale Key]

Answer area

```
CREATE NONCLUSTERED INDEX nc_fact_sale ON Fact.Sale
(
    Transact-SQL segments.
)
GO
CREATE UNIQUE CLUSTERED INDEX cl_fact_sale ON Fact.Sale
(
    Transact-SQL segments.
) GO
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Transact-SQL statements

Drag and drop the following Transact-SQL statements into the answer area:

- Description
- [Sale Key] ASC
- [Sale Key] DESC
- Description ASC, [Sale Key]

Answer area

```
CREATE NONCLUSTERED INDEX nc_fact_sale ON Fact.Sale
(
  Description ASC, [Sale Key]
)
GO
CREATE UNIQUE CLUSTERED INDEX cl_fact_sale ON Fact.Sale
(
  Description
)
GO
```

NEW QUESTION 91

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 93

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.

Actions

Recreate the Fact.Order table on the partition scheme.

Execute an ALTER TABLE command to specify the partition function.

Create a partition scheme based on the partition function.

Execute an ALTER TABLE command to specify the partition scheme.

Recreate the Fact.Order table on the partition function.

Create a partition function.

Answer Area



- 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 94

You have a Microsoft SQL Server Data Warehouse instance that uses SQL Server Analysis Services (SSAS). The instance has a cube containing data from an on-premises SQL Server instance. A measure named Measure1 is configured to calculate the average of a column.

You plan to change Measure1 to a full additive measure and create a new measure named Measure2 that evaluates data based on the first populated row.

You need to configure the measures.

What should you do? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

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

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

References:

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

NEW QUESTION 98

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 103

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?

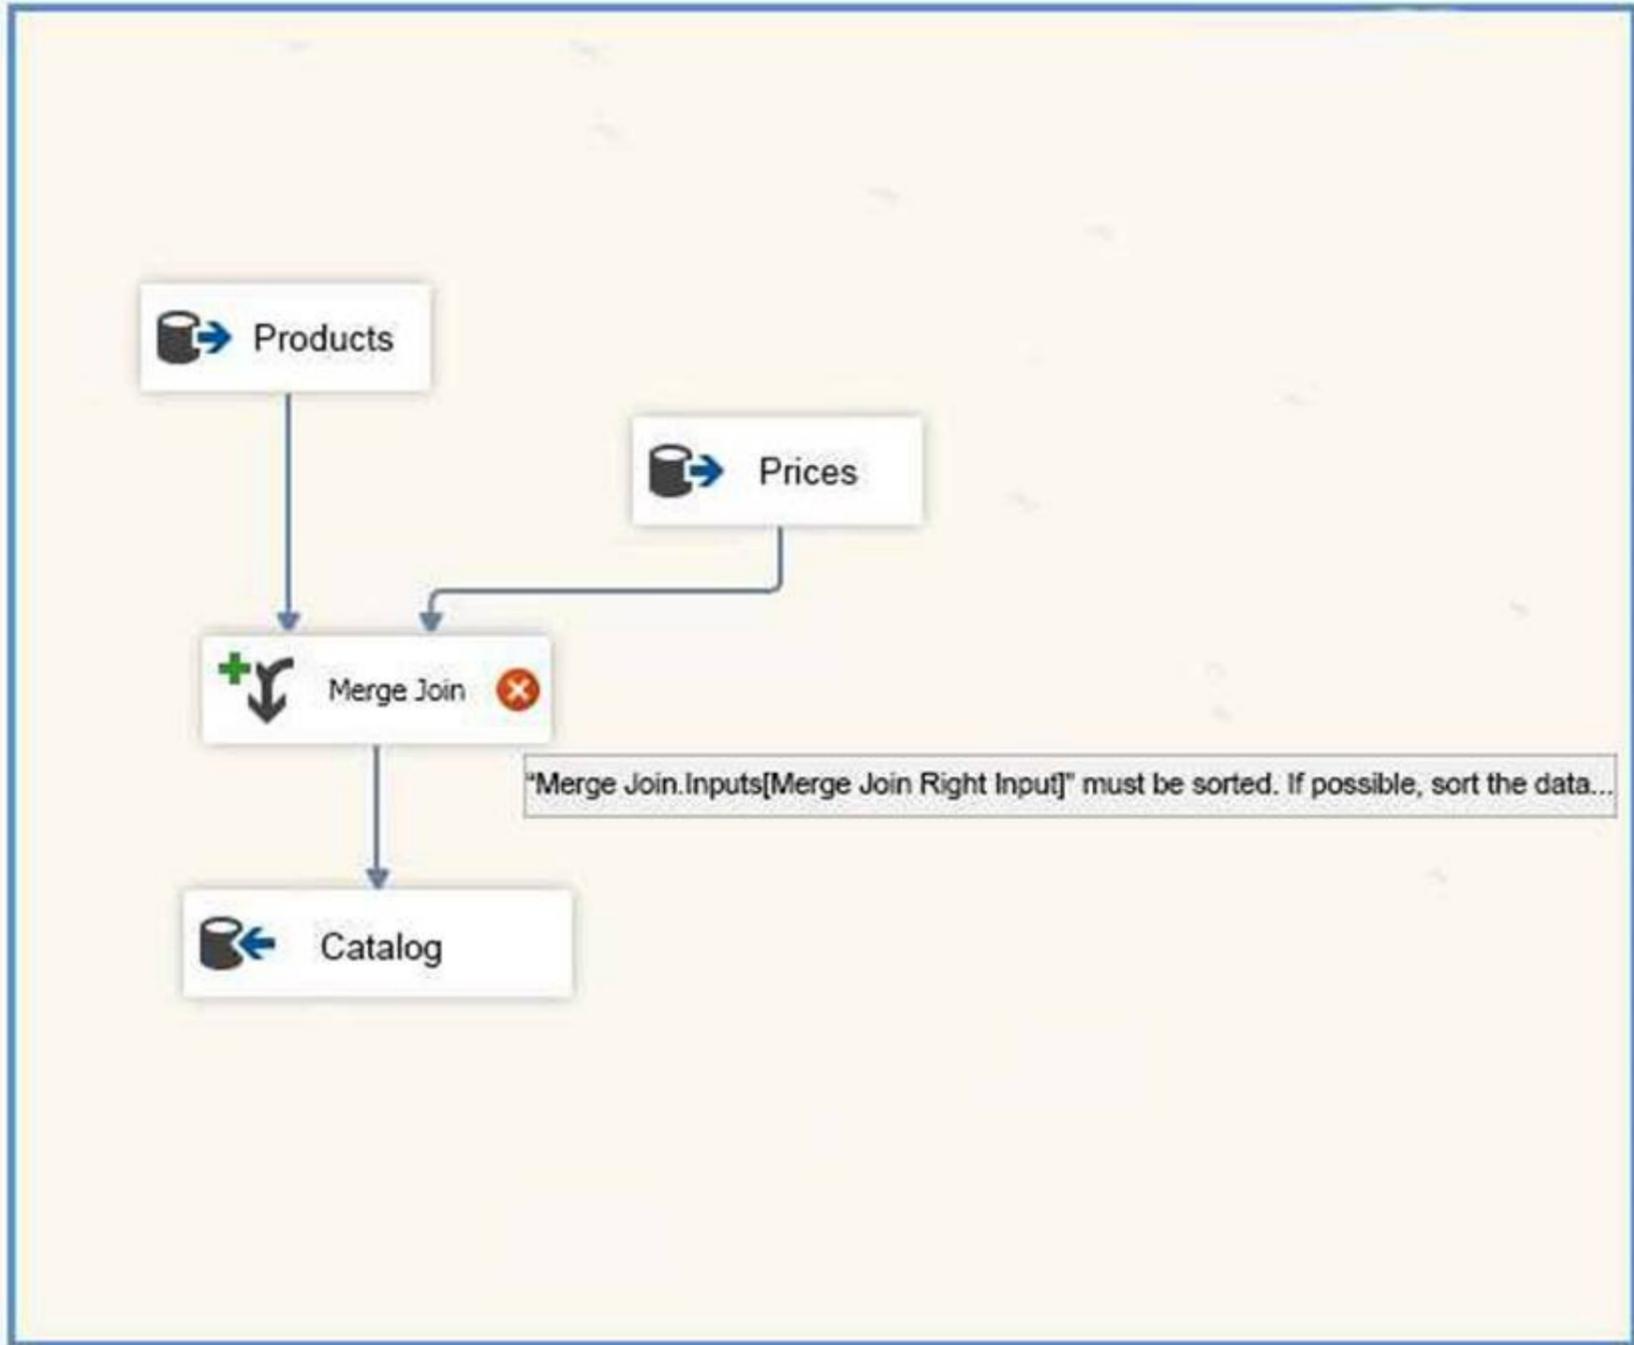
- A. star schema
- B. snowflake schema
- C. conformed dimension
- D. slowly changing dimension (SCD)

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

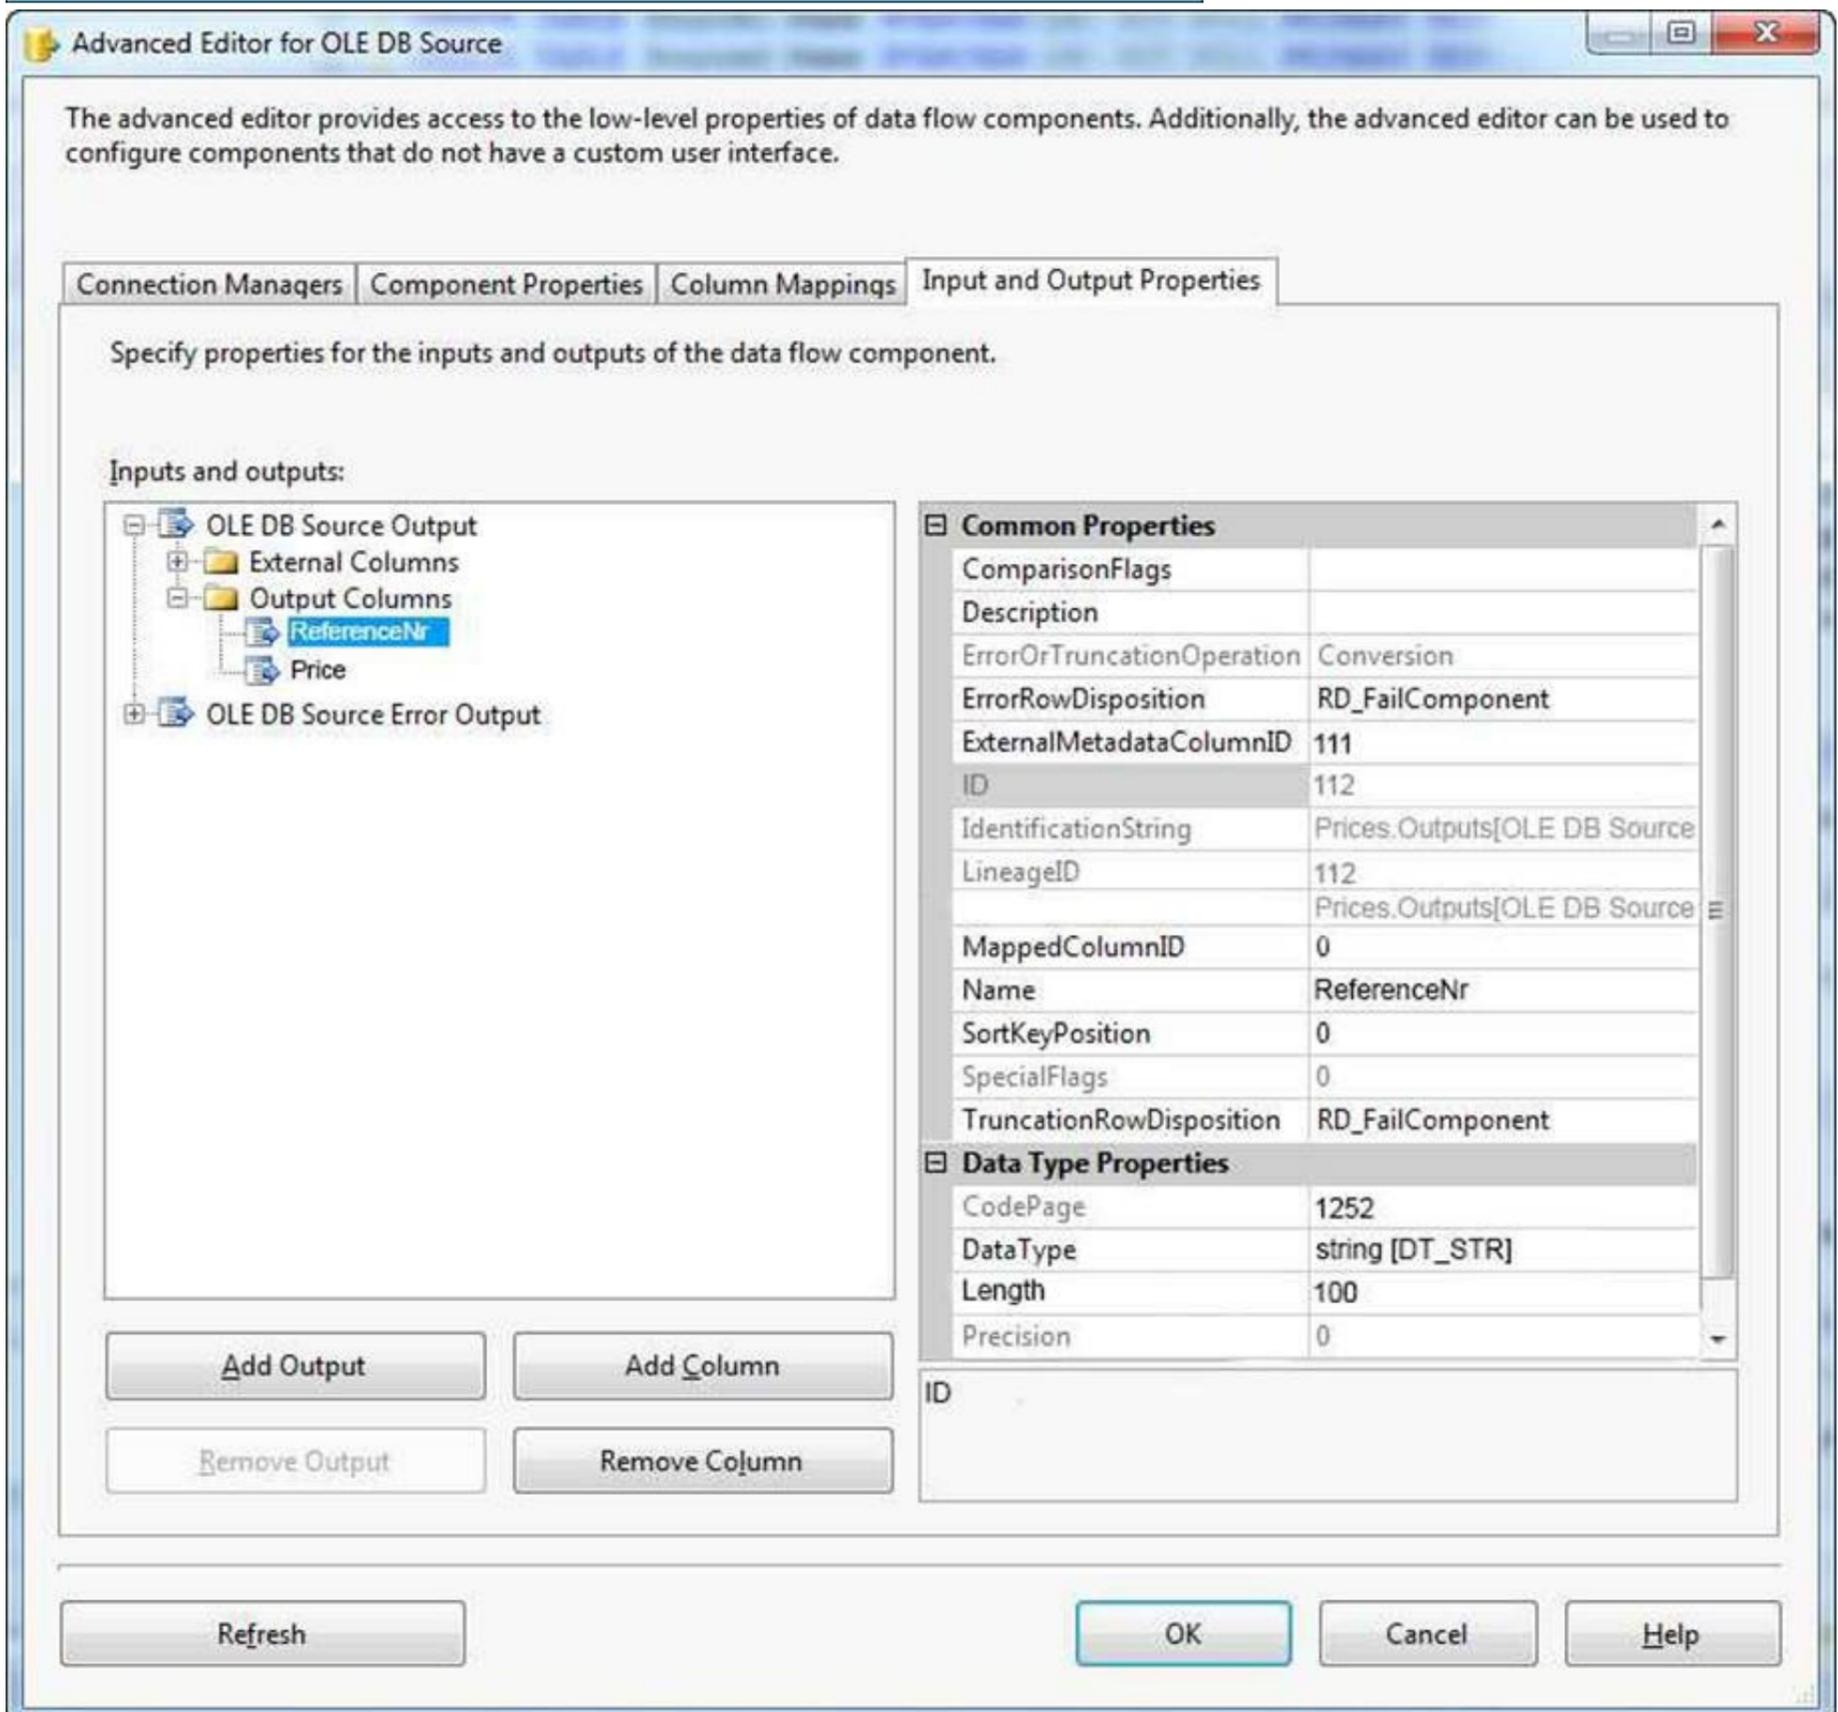
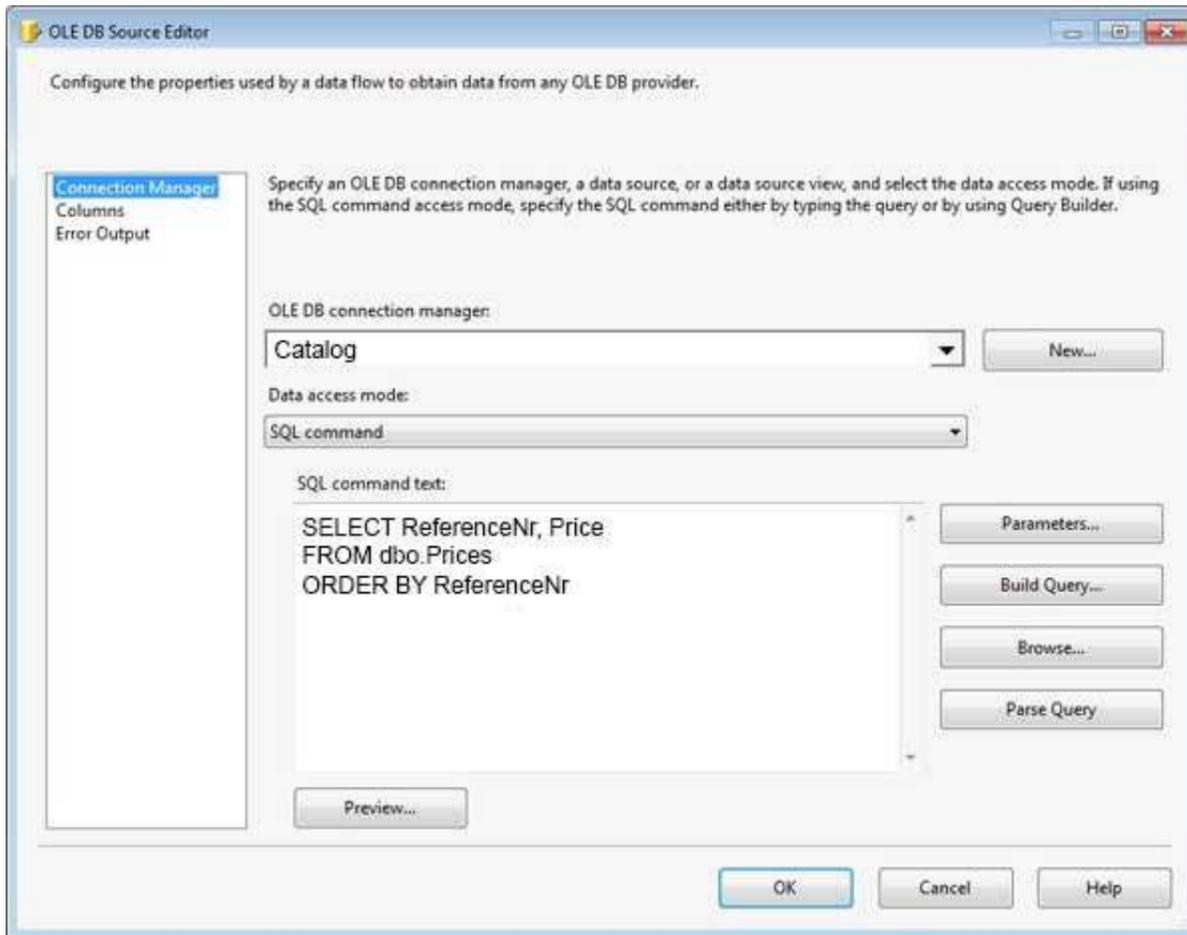
Answer: H

NEW QUESTION 107

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. | <input type="radio"/> | <input type="radio"/> |
| 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. | <input type="radio"/> | <input type="radio"/> |
| You can resolve the error by adding an Aggregate transform between the OLE DB source and the Merge Join transform. | <input type="radio"/> | <input type="radio"/> |
| You can resolve the error by replacing the Merge Join transform with a Lookup transform. | <input type="radio"/> | <input type="radio"/> |

- 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 112

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 113

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 115

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.

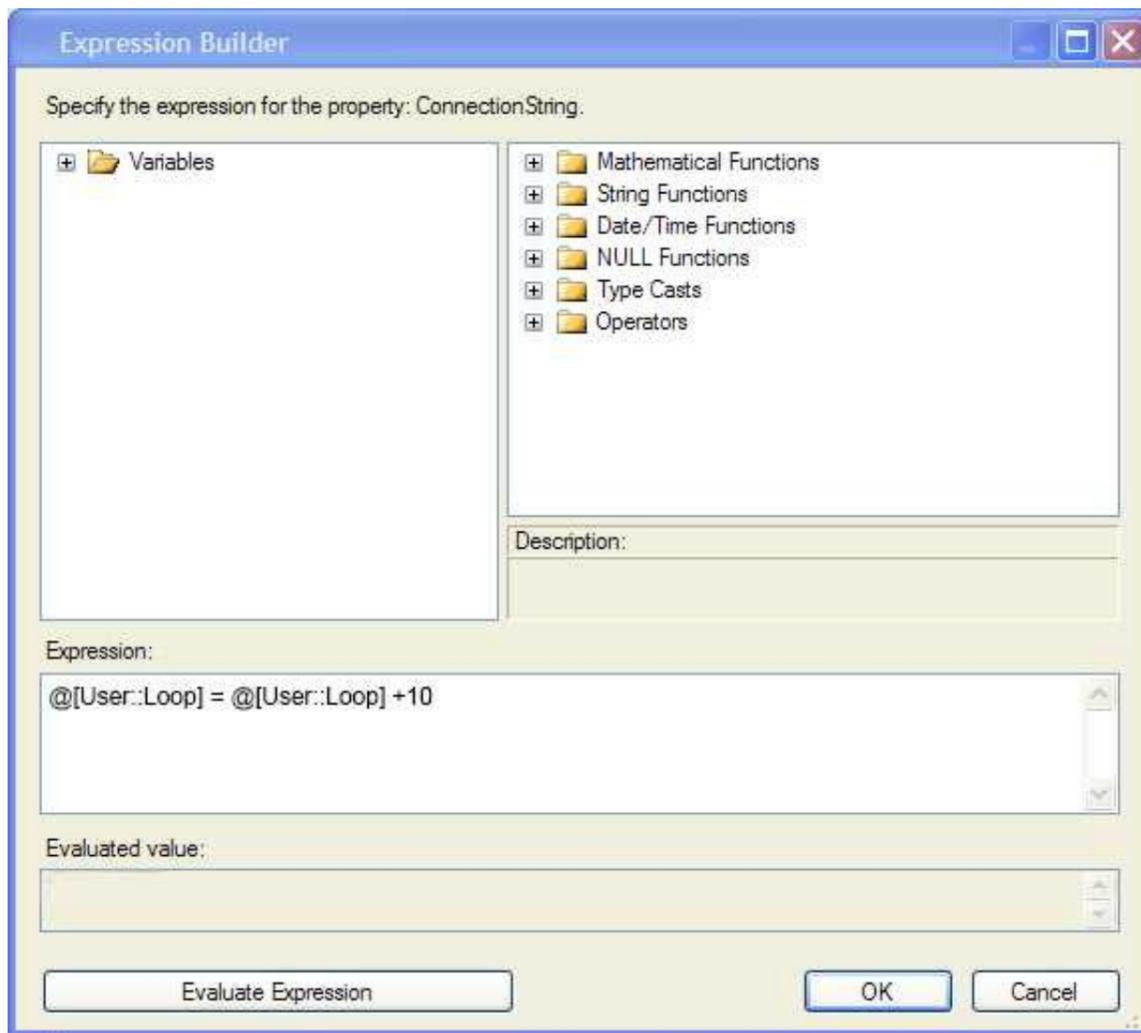
References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-statistics>

NEW QUESTION 117

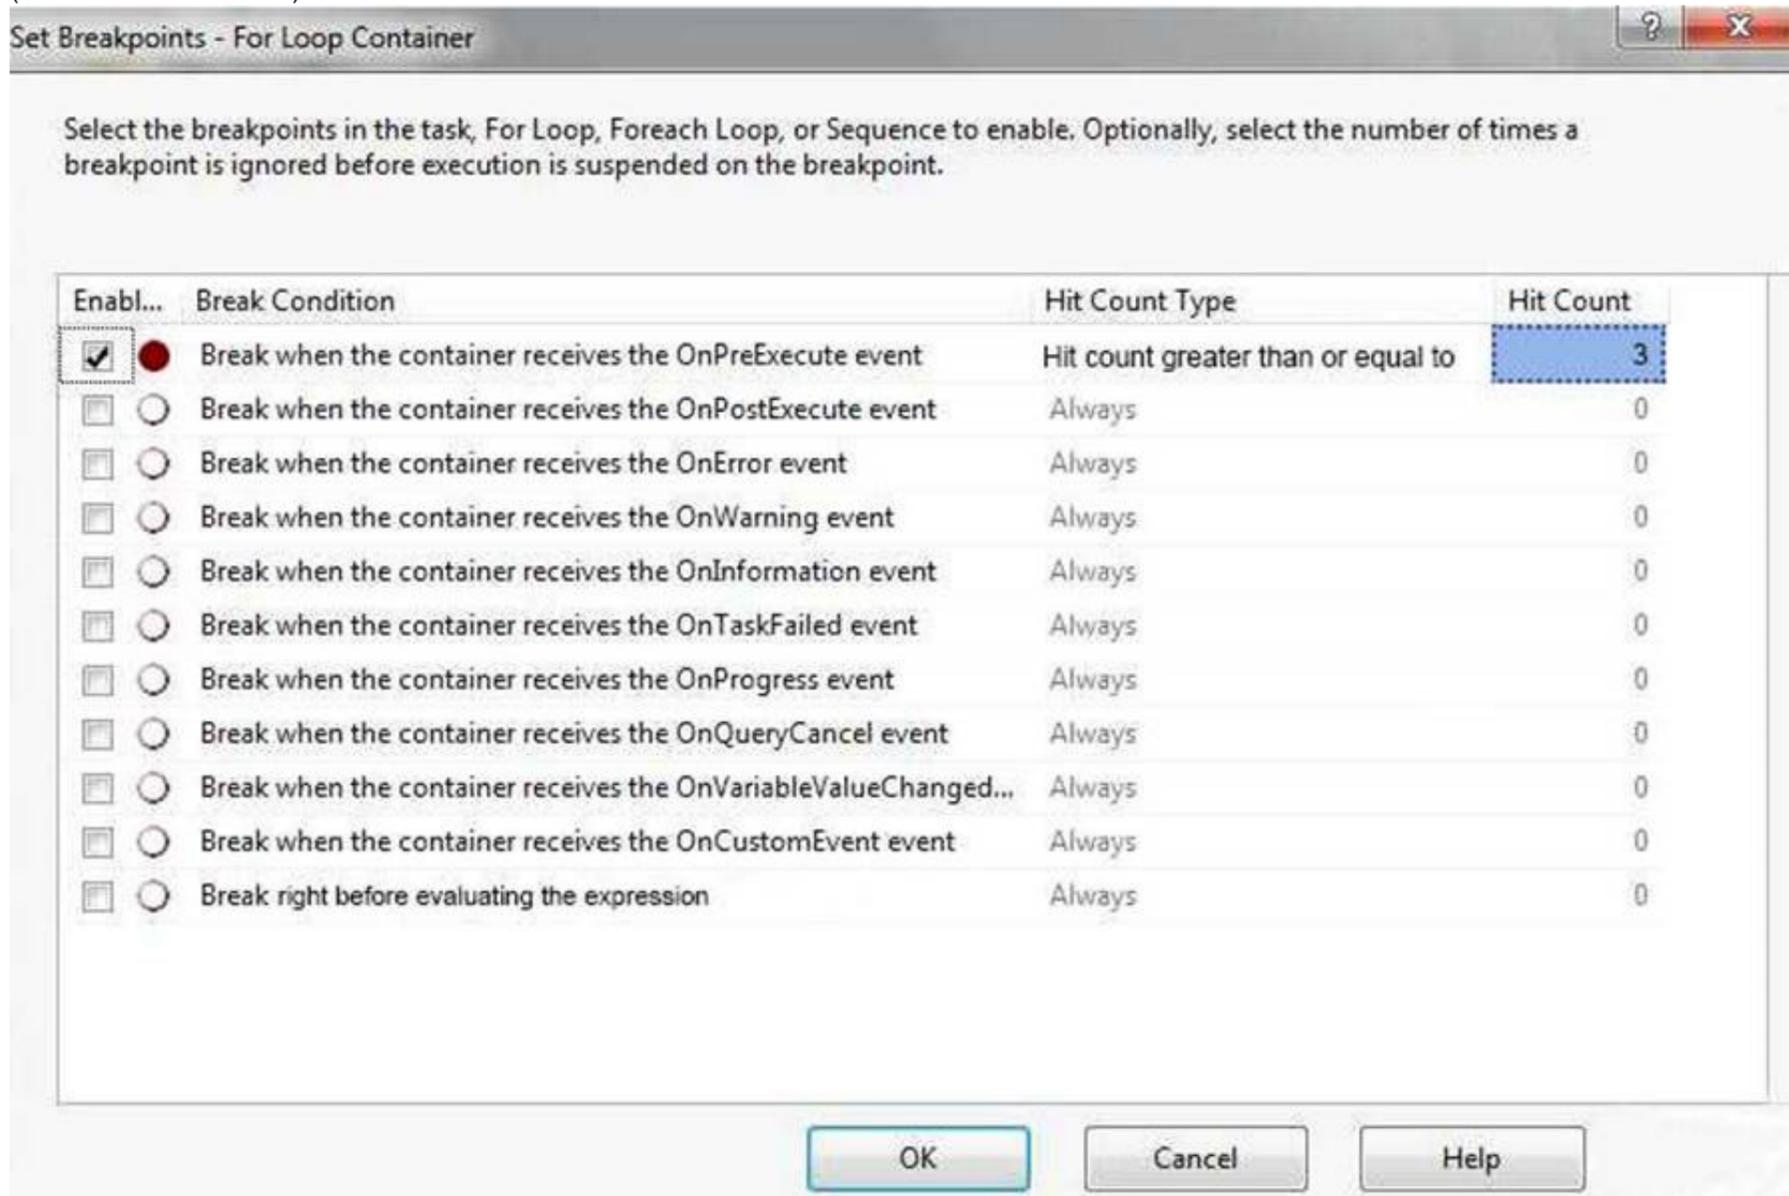
You have the Microsoft SQL Server Integration Services (SSIS) package shown in the Control flow exhibit. (Click the Exhibit button.)



The package iterates over 100 files in a local folder. For each iteration, the package increments a variable named loop as shown in the Expression task exhibit. (Click the Exhibit button) and then imports a file. The initial value of the variable loop is 0.



You suspect that there may be an issue with the variable value during the loop. You define a breakpoint on the Expression task as shown in the BreakPoint exhibit. (Click the Exhibit button.)



You need to check the value of the loop variable value.

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

Answer Area

| | Yes | No |
|---|-----------------------|-----------------------|
| The value of the loop variable is 20 after the breakpoint is reached for the first time. | <input type="radio"/> | <input type="radio"/> |
| The loop variable resets to 0 when the breakpoint is reached. | <input type="radio"/> | <input type="radio"/> |
| When the code stops at a breakpoint, you can change the value of the loop variable. | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 121

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 126

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 131

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. Your company uses Microsoft SQL Server to deploy a data warehouse to an environment that has a SQL Server Analysis Services (SSAS) instance. The data warehouse includes the Fact.Order table as shown in the following table definition. The table has no indexes.

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

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

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

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

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

Solution: You create one columnstore index that includes the [Order Date Key], [Tax Amount], and [Total Excluding Tax] columns.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

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 133

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

| DateID | Stock-ID | OpeningPrice | Closing-Price | QuantityTraded | BrokerID | NumberOfTrades | Market-ID |
|----------|----------|--------------|---------------|----------------|----------|----------------|-----------|
| 20170301 | 22 | 30.20 | 34.23 | 100 | 10 | 1 | 1 |
| 20170301 | 31 | 10.05 | 12.23 | 110 | 10 | 2 | 2 |
| 20170302 | 22 | 30.89 | 34.76 | 899 | 5 | 1 | 1 |

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

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

Answer: C

Explanation:

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

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

NEW QUESTION 134

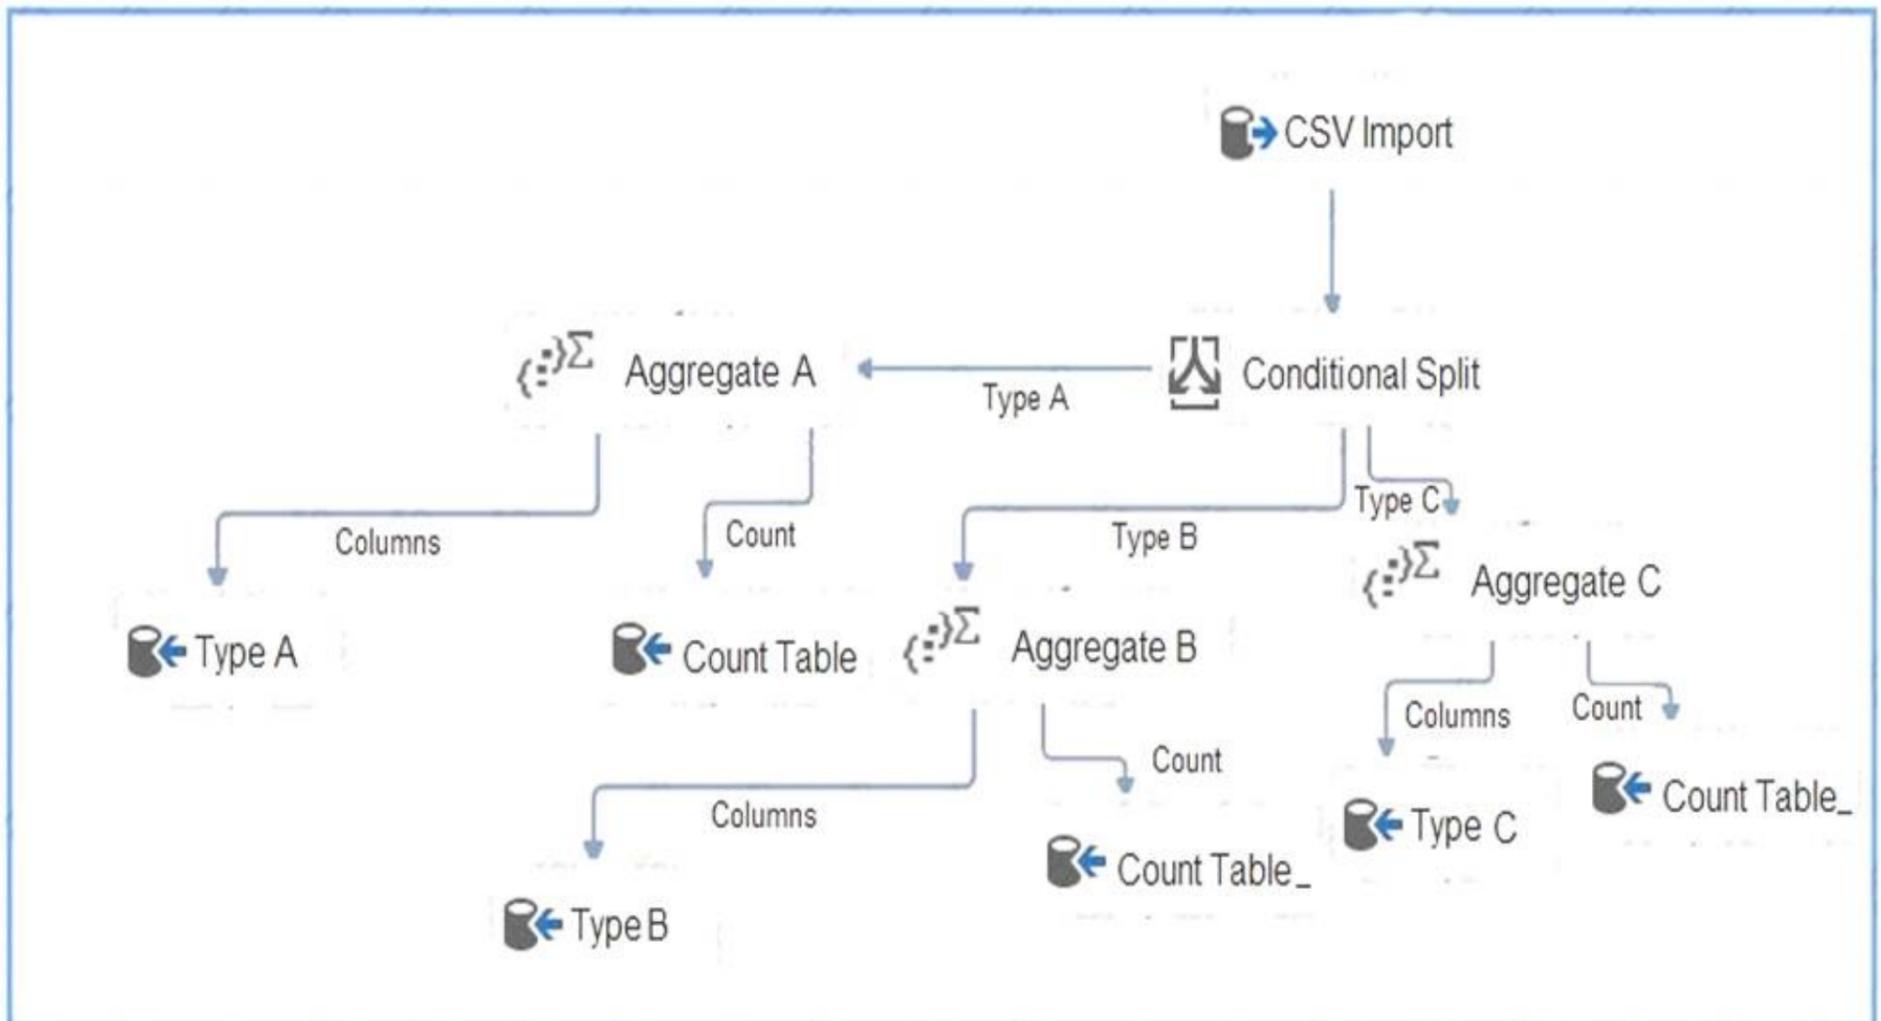
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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 137

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 142

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 145

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 148

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 |
| 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 statements

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

CREATE FUNCTION dbo.fn_stretch_by_date(@date
DATETIME2)
RETURNS TABLE
WITH SCHEMABINDING
AS
RETURN SELECT 1 AS (s_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:

```
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 (s_eligible WHERE @date <
CONVERT(datetime2, '1/1/2015', 101)
GO

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

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

Answer area

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

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

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

NEW QUESTION 153

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

You lock the current version. You need to manage the model versions.

Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area, and arrange them in the correct order.

Actions

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

Answer Area



- 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 155

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

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

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

How should you complete the Transact-SQL statement? To answer, drag the appropriate Transact-SQL segments to the correct locations. Each Transact-SQL segment may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Transact-SQL segments

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

Answer area

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

- A. Mastered
- B. Not Mastered

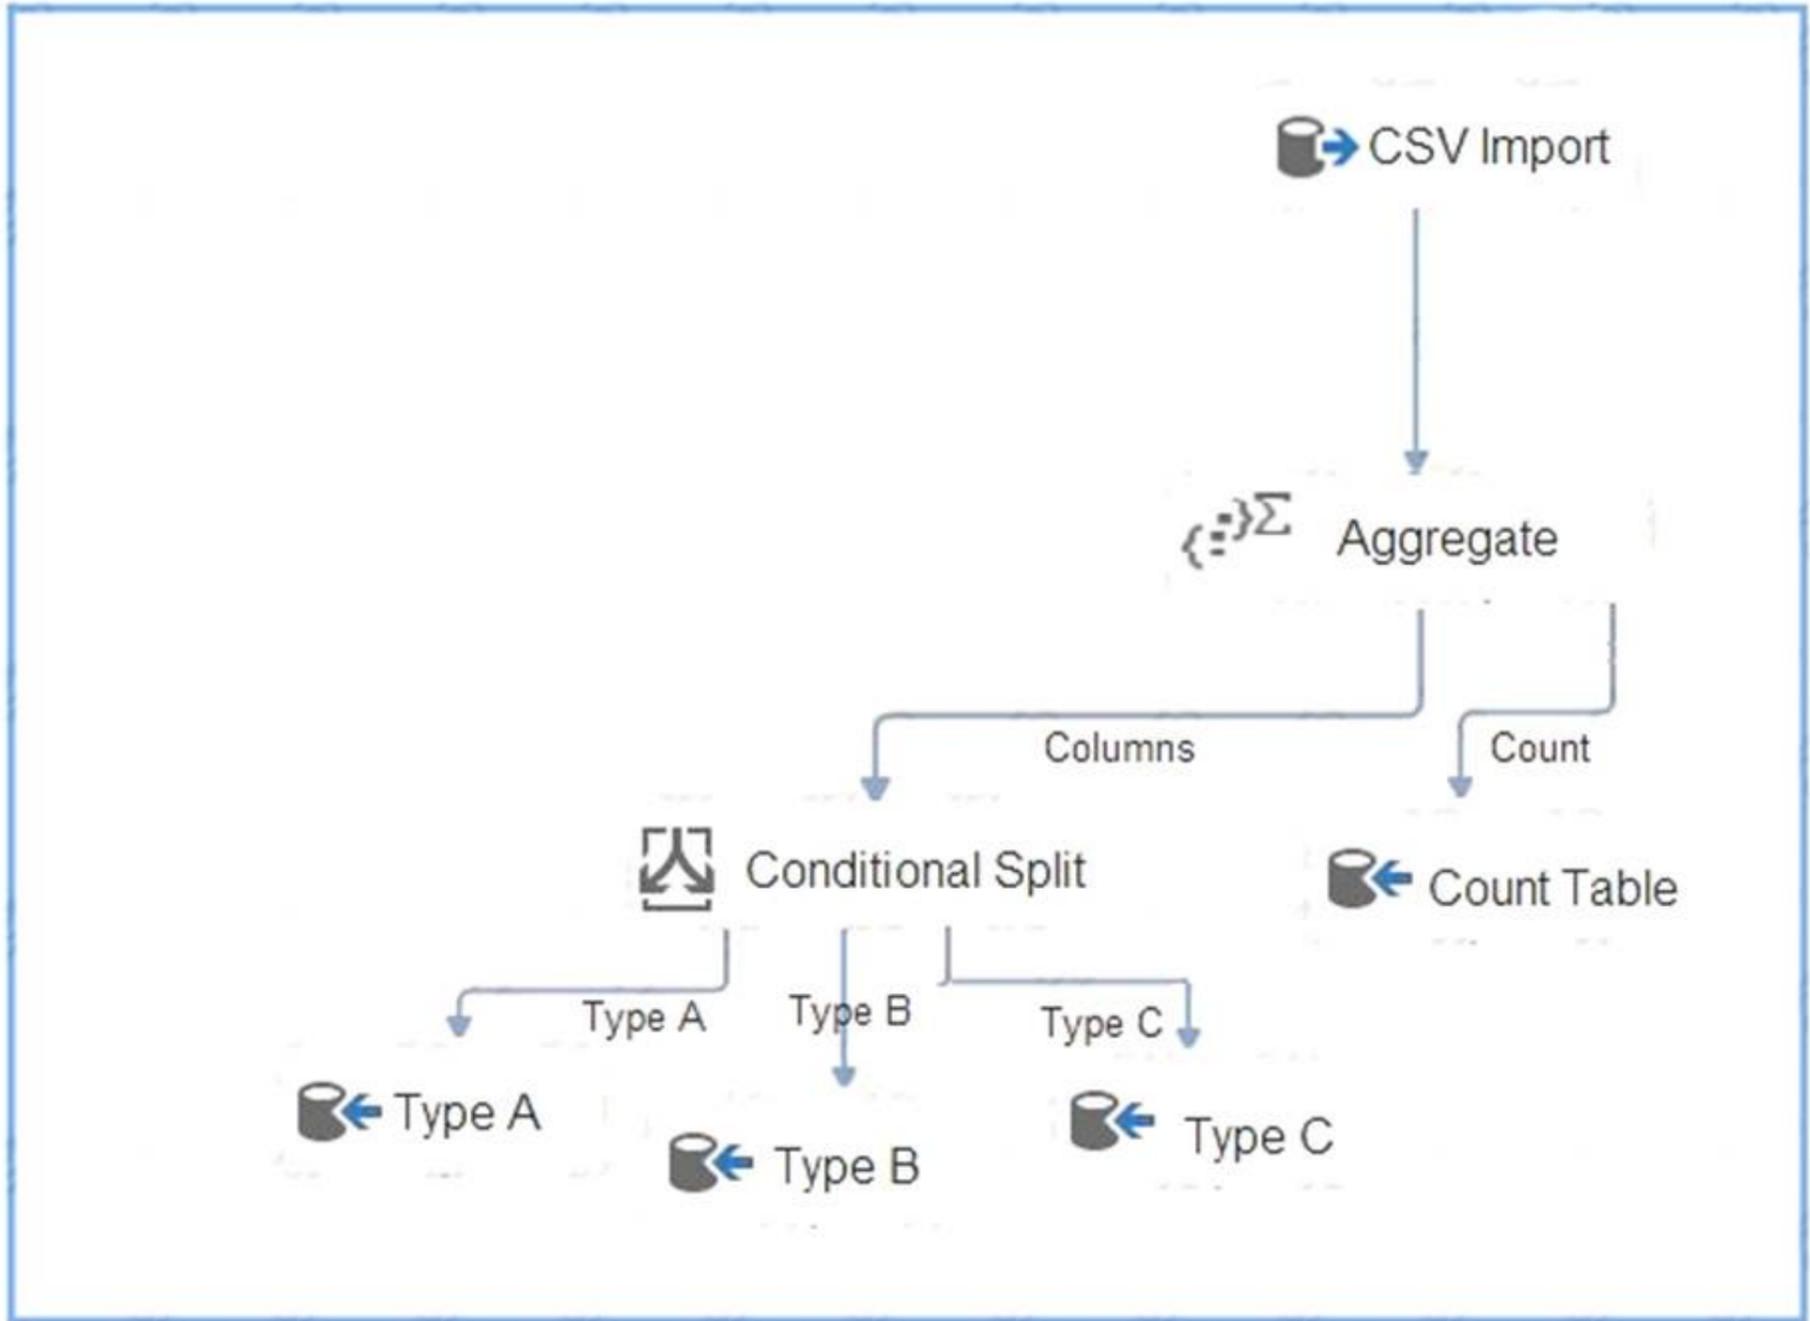
Answer: A

Explanation:

Box 1: EmailAddress
 The EmailAddress table is the target. Box 2: EmailAddress_import
 The EmailAddress_import table is the source. Box 3: NOT MATCHED BY TARGET
 Box 4: NOT MATCHED BY SOURCE
 References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/merge-transact-sql>

NEW QUESTION 160

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this 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 163

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

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

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

You plan to load at least one million rows of data each night from DB1 into the OnlineOrder table. You must load data into the correct partitions using a parallel process.

You create 24 Data Flow tasks. You must place the tasks into a component to allow parallel load. After all of the load processes compete, the process must proceed to the next task.

You need to load the data for the OnlineOrder table. What should you use?

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

H. Foreach Loop container

Answer: H

Explanation:

The Parallel Loop Task is an SSIS Control Flow task, which can execute multiple iterations of the standard Foreach Loop Container concurrently.

References:

<http://www.cozyroc.com/ssis/parallel-loop-task>

NEW QUESTION 164

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

Start of repeated scenario

You have a Microsoft SQL Server data warehouse instance that supports several client applications. The data warehouse includes the following tables:

Dimension.SalesTerritory, Dimension.Customer,

Dimension.Date, Fact.Ticket and Fact.Order. The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. The Fact.Order table is

optimized for weekly reporting, but the company wants to change it to daily. The FactOrder table is loaded by using an ETL process. Indexes have been added to the table over time, but the presence of these indexes slows data loading.

All data in the data warehouse is stored on a shared SAN. All tables are in a database named DB1. You have a second database named DB2 that contains copies of production data for a development environment. The data warehouse has grown and the cost of storage has increased. Data older than one year is accessed infrequently

and is considered historical.

- Implement table partitioning to improve the manageability of the data warehouse and to avoid the need to repopulate all transactional data each night Use a partitioning strategy that is as granular as possible.
- Partition the FactOrder table and retain a total of seven years of data.
- Partition the Fact.Ticket table and retain seven years of data. At the end of each month, the partition structure must apply a sliding window strategy to ensure that a new partition is available for the upcoming month, and that the oldest month of data is archived and removed.
- Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.
- Incrementally load all tables in the database and ensure that all incremental changes are processed.
- Maximize the performance during the data loading process for the Fact.Order partition.
- Ensure "that historical data remains online and available for querying.
- Reduce ongoing storage costs while maintaining query performance for current data. You are not permitted to make changes to the client applications.

End of repeated scenario

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

Which technology should you use for each table?

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

Answer area

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

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

- 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 167

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 cleanse a data flow source by removing duplicate records based on approximate matches. 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: E

Explanation:

The Fuzzy Grouping transformation performs data cleaning tasks by identifying rows of data that are likely to be duplicates and selecting a canonical row of data to use in standardizing the data.

NEW QUESTION 170

You are designing a data transformation process using Microsoft SQL Server Integration Services (SSIS). You need to ensure that every row is compared with every other row during transformation.

What should you configure? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

| SSIS package option | SSIS package configuration |
|-------------------------|--|
| Transformation type | <div style="border: 1px solid black; padding: 5px;"> <div style="text-align: right; margin-bottom: 5px;">▼</div> Fuzzy Grouping Merge Merge Join Multicast </div> |
| Transformation property | <div style="border: 1px solid black; padding: 5px;"> <div style="text-align: right; margin-bottom: 5px;">▼</div> Exhaustive Precision SpecialFlags </div> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

When you configure the Fuzzy Grouping transformation, you can specify the comparison algorithm that the transformation uses to compare rows in the transformation input. If you set the Exhaustive property to true, the transformation compares every row in the input to every other row in the input. This comparison algorithm may produce more accurate results, but it is likely to make the transformation perform more slowly unless the number of rows in the input is small.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/fuzzy-grouping-transformati>

NEW QUESTION 175

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions 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 First Try
- * 70-767 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 70-767 Practice Test Here](#)