

70-767 Dumps

Implementing a SQL Data Warehouse (beta)

<https://www.certleader.com/70-767-dumps.html>



NEW QUESTION 1

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 chose.
- ▶ 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 2

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 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 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 4

You are building a server to host a data warehouse.

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

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

*Minimize hardware costs.

Which RAID configuration should you recommend?

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

Answer: C

Explanation:

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

NEW QUESTION 5

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

You are 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 6

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

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

You 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 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 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 8

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.dimension_city table to Dimension.city. Data loss is not permissible, and you must not leave traces of the old table in the data warehouse. The fact.Transaction table has measures named RawCost and Totalsale that calculate the wholesale cost of materials. You plan to create a measure that calculates the profit margin based on the two 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 create the ProfitMargin measure for the fact. Transaction table.

How should you complete the MDX statement? To answer, select the appropriate MDX segments in the answer area.

Answer area

```

CREATE MEMBER CURRENTCUBE.Measures.ProfitMargin
CREATE SET
CREATE SUBCUBE
AS 'Measures. [TotalSale] [RawCost] /Measures. [TotalSale] [RawCost]'
    
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer area

```

CREATE MEMBER CURRENTCUBE.Measures.ProfitMargin
CREATE SET
CREATE SUBCUBE
AS 'Measures. [TotalSale] [RawCost] /Measures. [TotalSale] [RawCost]'
    
```

NEW QUESTION 9

You have a data warehouse.

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

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

Actions

Add a file to the LowCost filegroup.

Rename the Fact.ErrorLog table to Fact.ErrorLogBak.

Drop the Fact.ErrorLog table.

Create a new Fact.ErrorLog table on the LowCost filegroup.

Add a filegroup named LowCost to the database.

Reorganize the clustered index on the Fact.ErrorLog table in the new filegroup.

Rebuild the clustered index on the Fact.ErrorLog table in the new filegroup.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: Add a filegroup named LowCost to the database. First create a new filegroup.

Step 2:

The next stage is to go to the 'Files' page in the same Properties window and add a file to the filegroup (a filegroup always contains one or more files)

Step 3:

To move a table to a different filegroup involves moving the table's clustered index to the new filegroup. While this may seem strange at first this is not that surprising when you remember that the leaf level of the clustered index actually contains the table data. Moving the clustered index can be done in a single statement using the DROP_EXISTING clause as follows (using one of the AdventureWorks2008R2 tables as an example) :

```
CREATE UNIQUE CLUSTERED INDEX PK_Department_DepartmentID ON HumanResources.Department(DepartmentID) WITH (DROP_EXISTING=ON,ONLINE=ON) ON SECONDARY
```

This recreates the same index but on the SECONDARY filegroup.

References:

<http://www.sqlmatters.com/Articles/Moving%20a%20Table%20to%20a%20Different%20Filegroup.aspx>

NEW QUESTION 10

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

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

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

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

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

You need to optimize performance.

Solution: You run the following Transact-SQL statement:

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

References:

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

NEW QUESTION 10

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 13

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 16

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

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

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

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

The query fails to return results.

You need to determine why the query fails.

Solution: You run the following Transact-SQL statement:

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

To use submit_time we must use sys.dm_pdw_exec_requests table. References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec>

NEW QUESTION 20

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the

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

- A. Yes
- B. No

Answer: B

Explanation:

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

NEW QUESTION 23

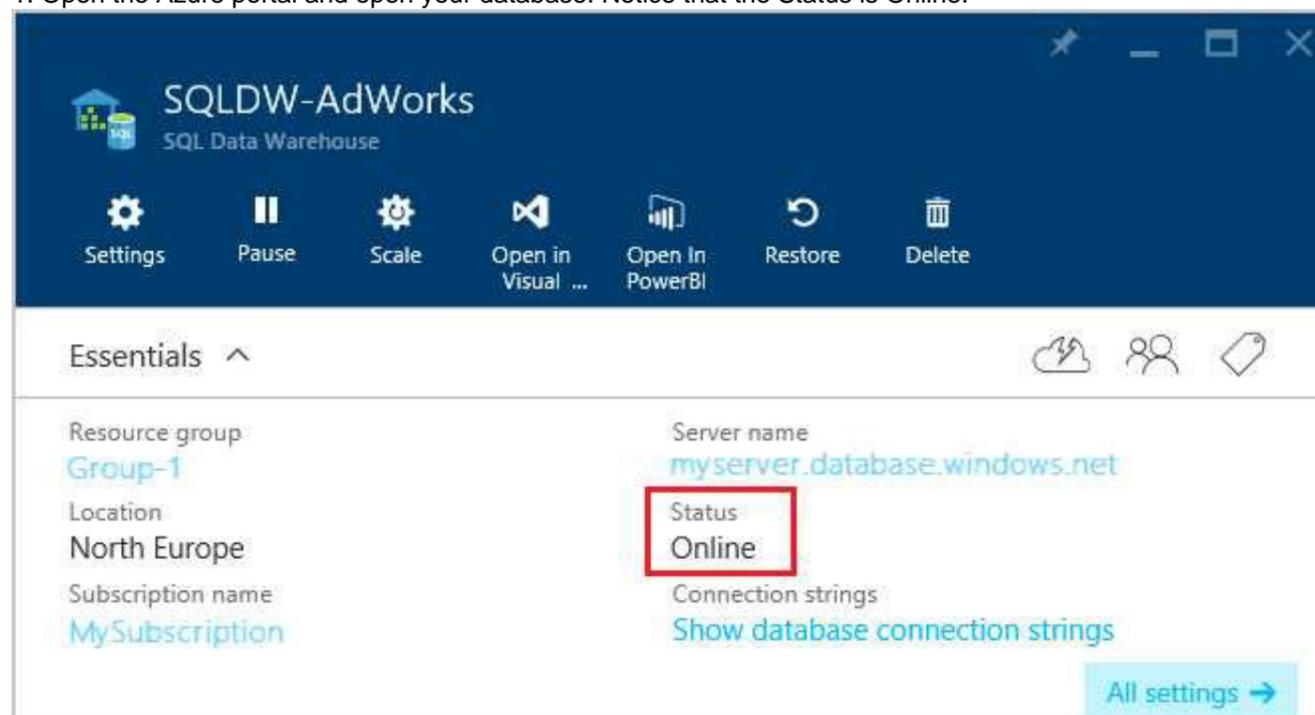
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have a Microsoft Azure SQL Data Warehouse instance that must be available six months a day for reporting. You need to pause the compute resources when the instance is not being used. Solution: You use the Azure portal. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

To pause a SQL Data Warehouse database, use any of these individual methods. Pause compute with Azure portal
Pause compute with PowerShell
Pause compute with REST APIs
Note: To pause a database:
1. Open the Azure portal and open your database. Notice that the Status is Online.



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

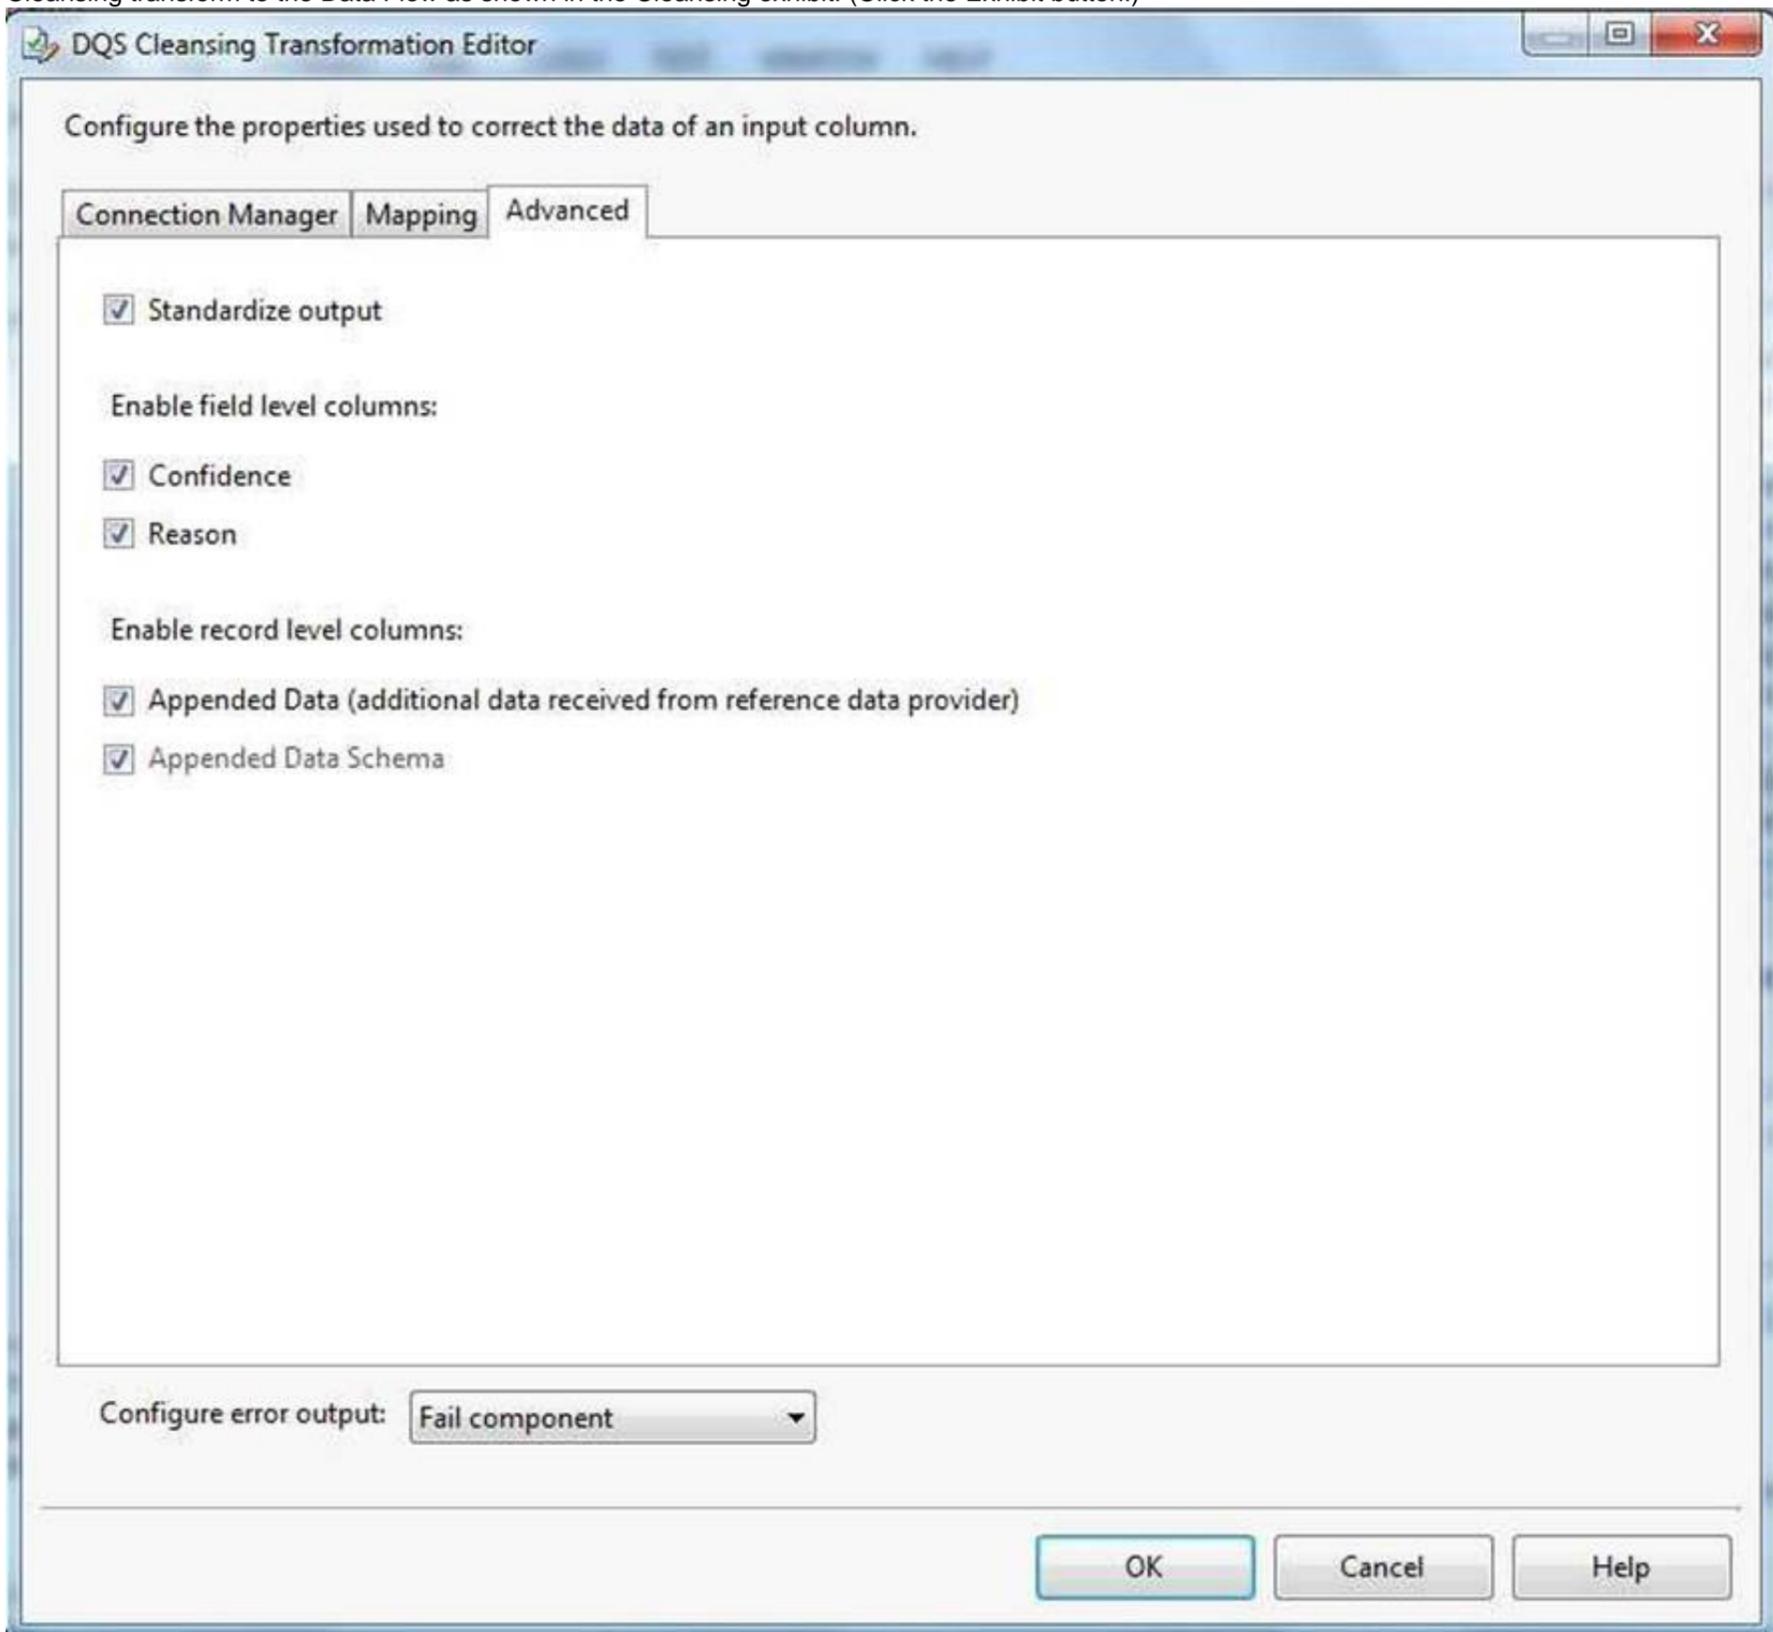
References:
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-compute-overview> <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-compute-portal#pause-c>

NEW QUESTION 25

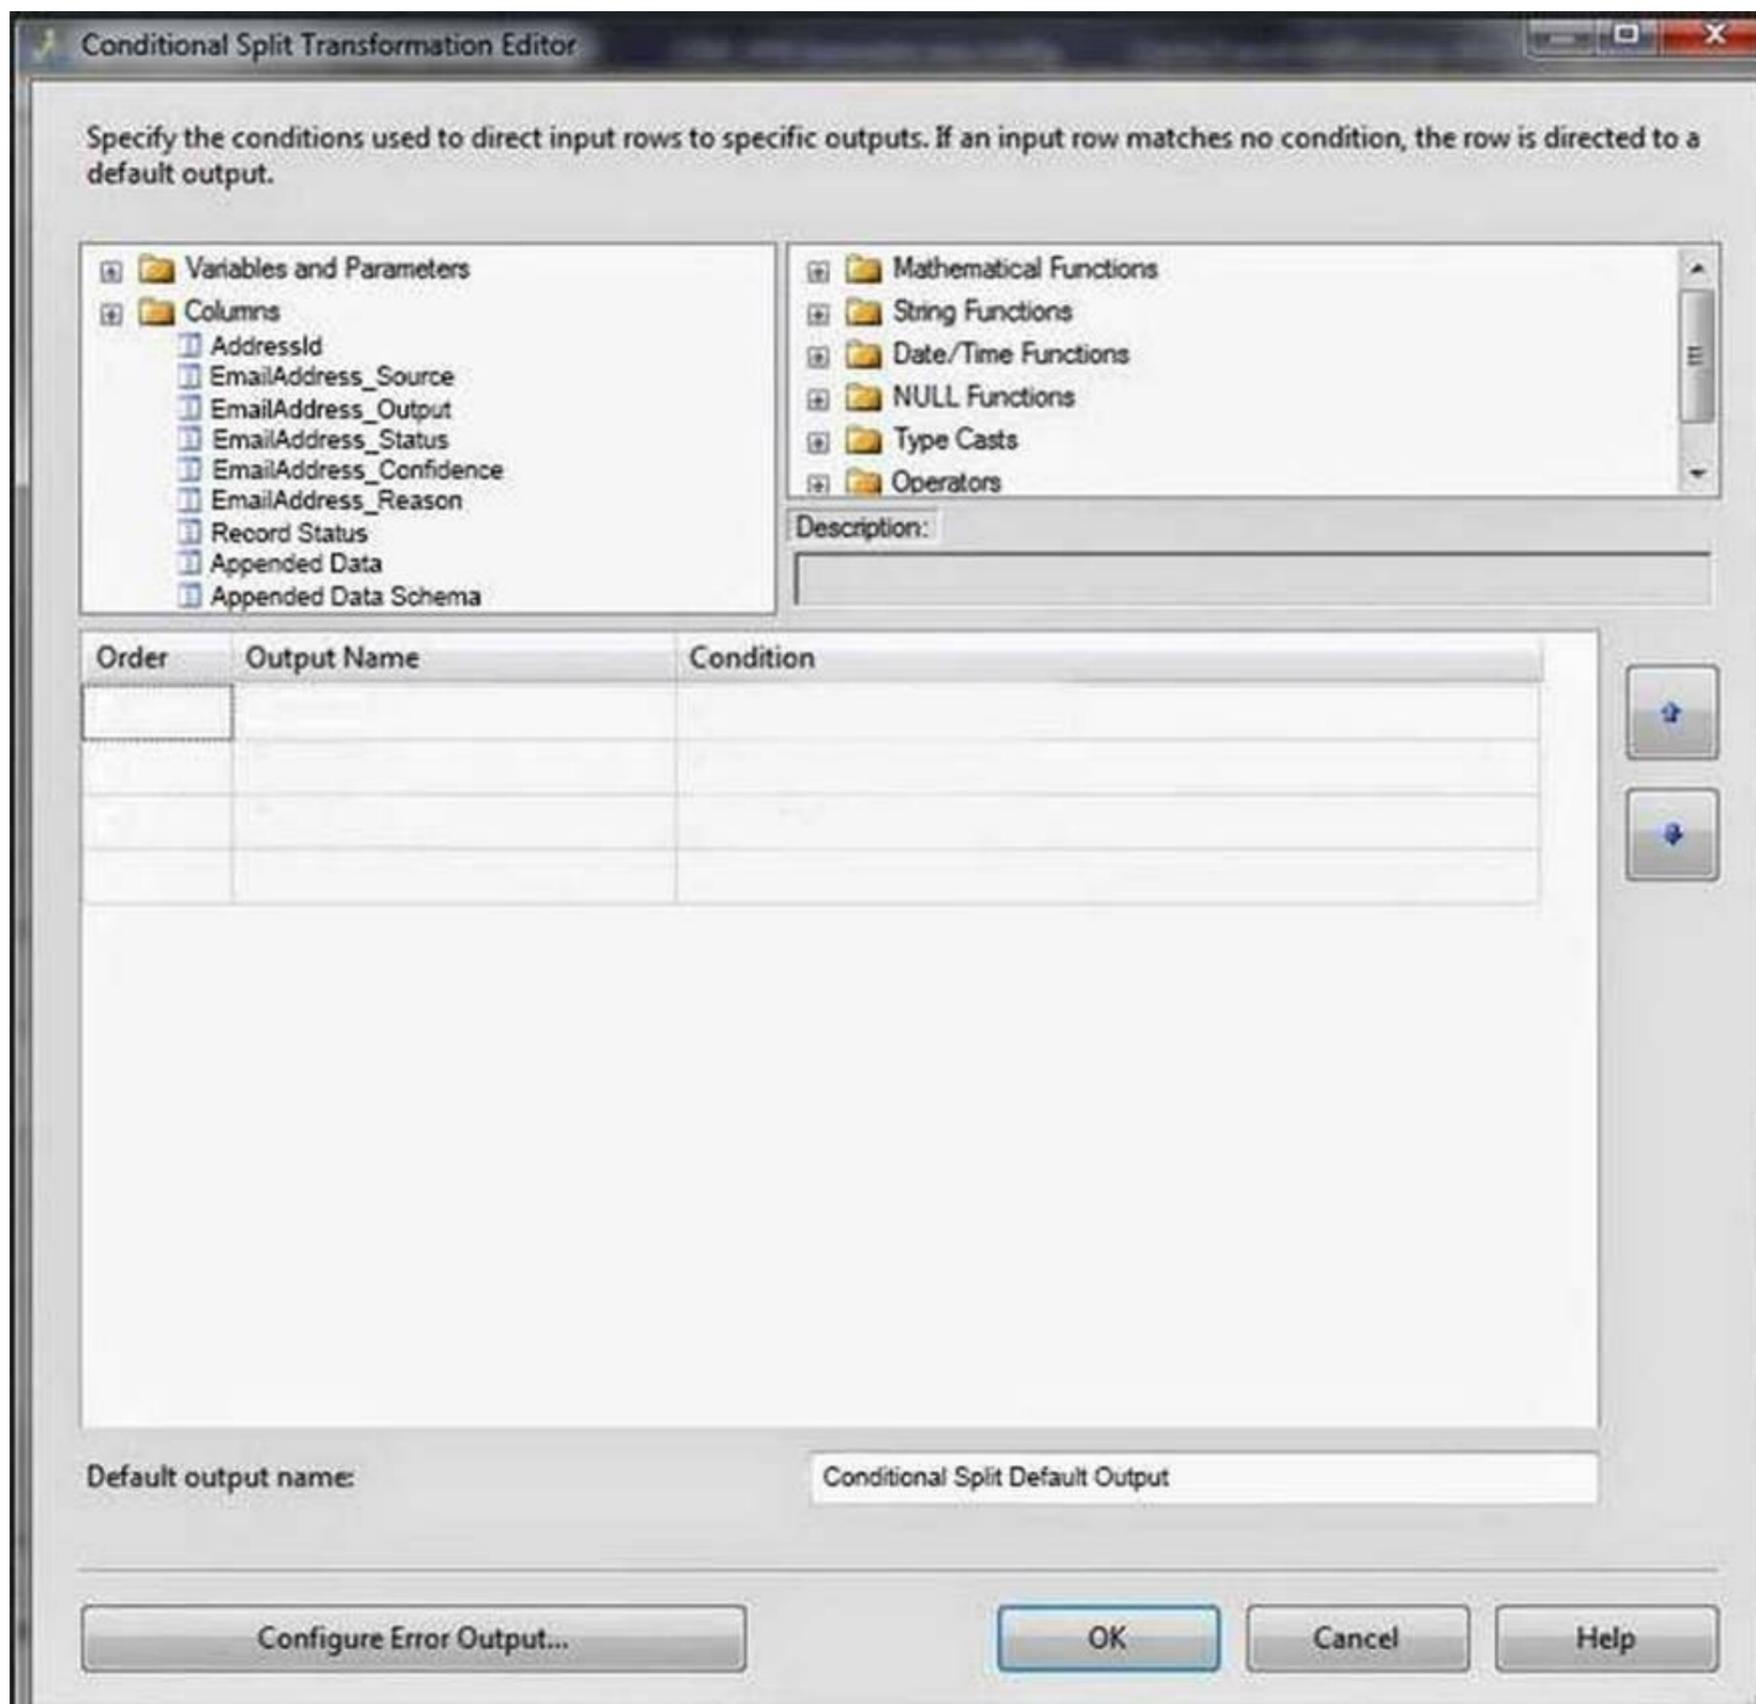
You have a Microsoft SQL Server Integration Services (SSIS) package that contains a Data Flow task as shown in the Data Flow exhibit. (Click the Exhibit button.)



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 30

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 34

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 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 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 39

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 41

You manage an inventory system that has a table named Products. The Products table has several hundred columns.

You generate a report that relates two columns named ProductReference and ProductName from the Products table. The result is sorted by a column named QuantityInStock from largest to smallest.

You need to create an index that the report can use.

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

Answer Area

CREATE INDEX lx_product

ON dbo.Products

INCLUDE

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

CREATE INDEX lx_product

CLUSTERED
 NONCLUSTERED

ON dbo.Products

(ProductReference)
 (QuantityInStock)
 (ProductName)

INCLUDE

(Products)
 (ProductReference)
 (ProductName, ProductReference)

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 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 45

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 50

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 temporarily. The solution must ensure that tracking changes can be restored quickly in a few weeks.

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: C

Explanation:

catalog.stop_operation stops a validation or instance of execution in the Integration Services catalog.

References:

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

NEW QUESTION 55

You are developing a Microsoft SQL Server Integration Services (SSIS) package to incrementally load new and changed records from a data source.

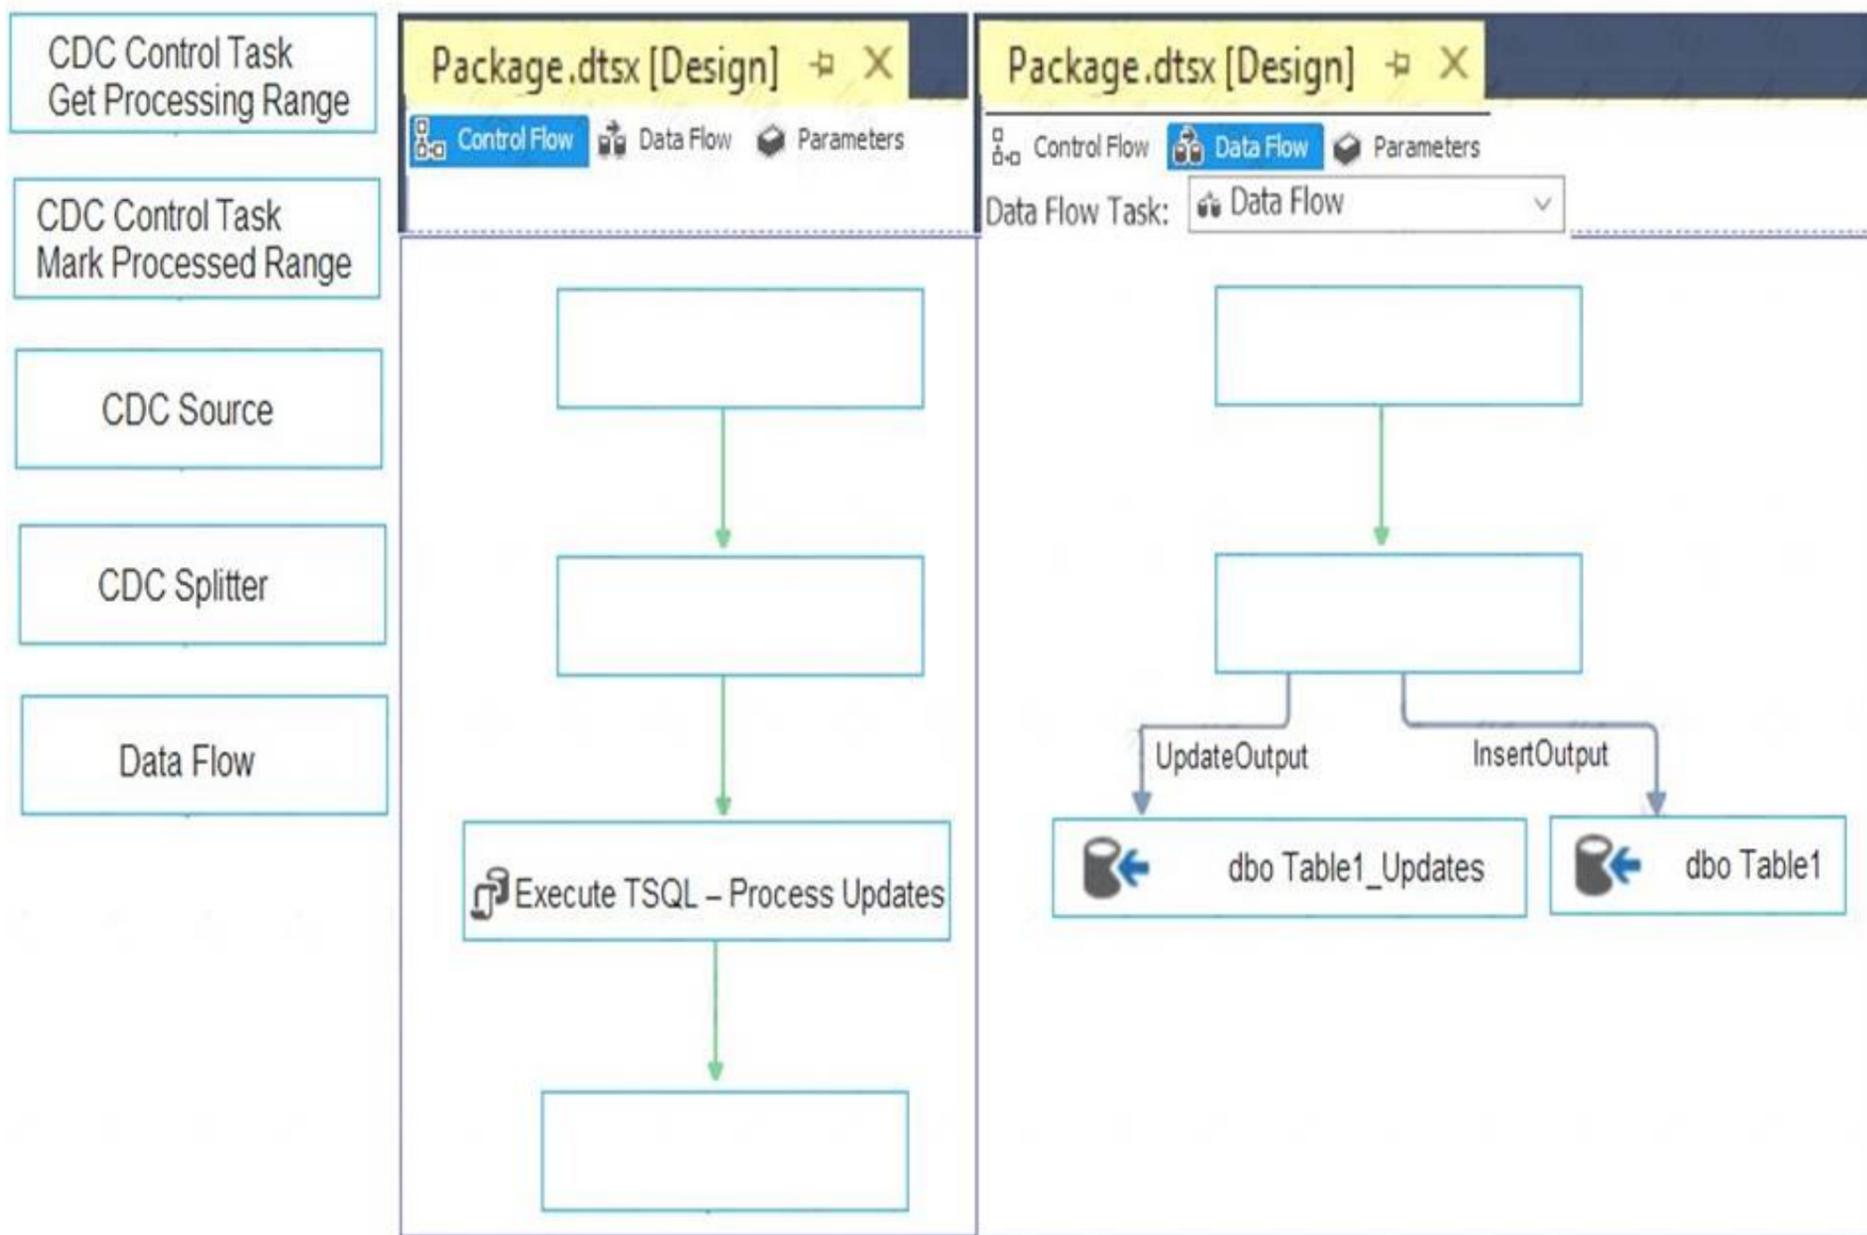
The SSIS package must load new records into Table1 and updated records into Table1_Updates. After loading records, the package must call a Transact-SQL statement to process updated rows according to existing business logic.

You need to complete the design of the SSIS package.

Which tasks should you use? To answer, drag the appropriate SSIS objects to the correct targets. Each SSIS object may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: CDC Control Task Get Processing Range Step 2: Mark Processed Range
Step 3: Data Flow

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

Step 4: CDC Source

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

Step 5: CDC Splitter

The CDC splitter splits a single flow of change rows from a CDC source data flow into different data flows for Insert, Update and Delete operations.

References:

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

NEW QUESTION 56

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 59

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 62

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

You are 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 63

You have a data warehouse named DW1. All data files are located on drive E. You expect queries that pivot hundreds of millions of rows for each report. You need to modify the data files to minimize latency.

What should you do?

- A. Add more data files to DW1 on drive E.
- B. Add more data files to tempdb on drive E.
- C. Remove data files from tempdb
- D. Remove data files from DW1.

Answer: B

Explanation:

The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.

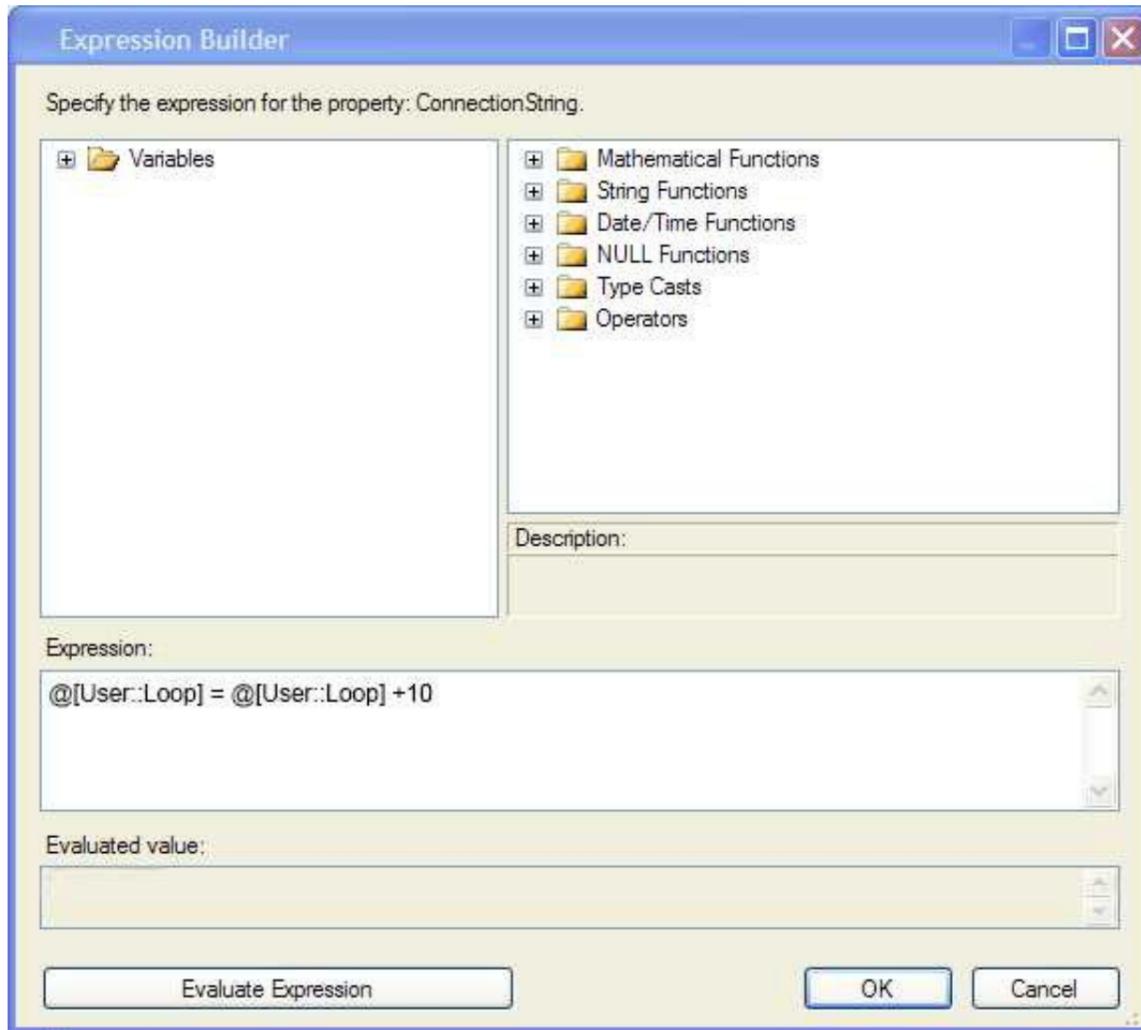
References: <https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database>

NEW QUESTION 64

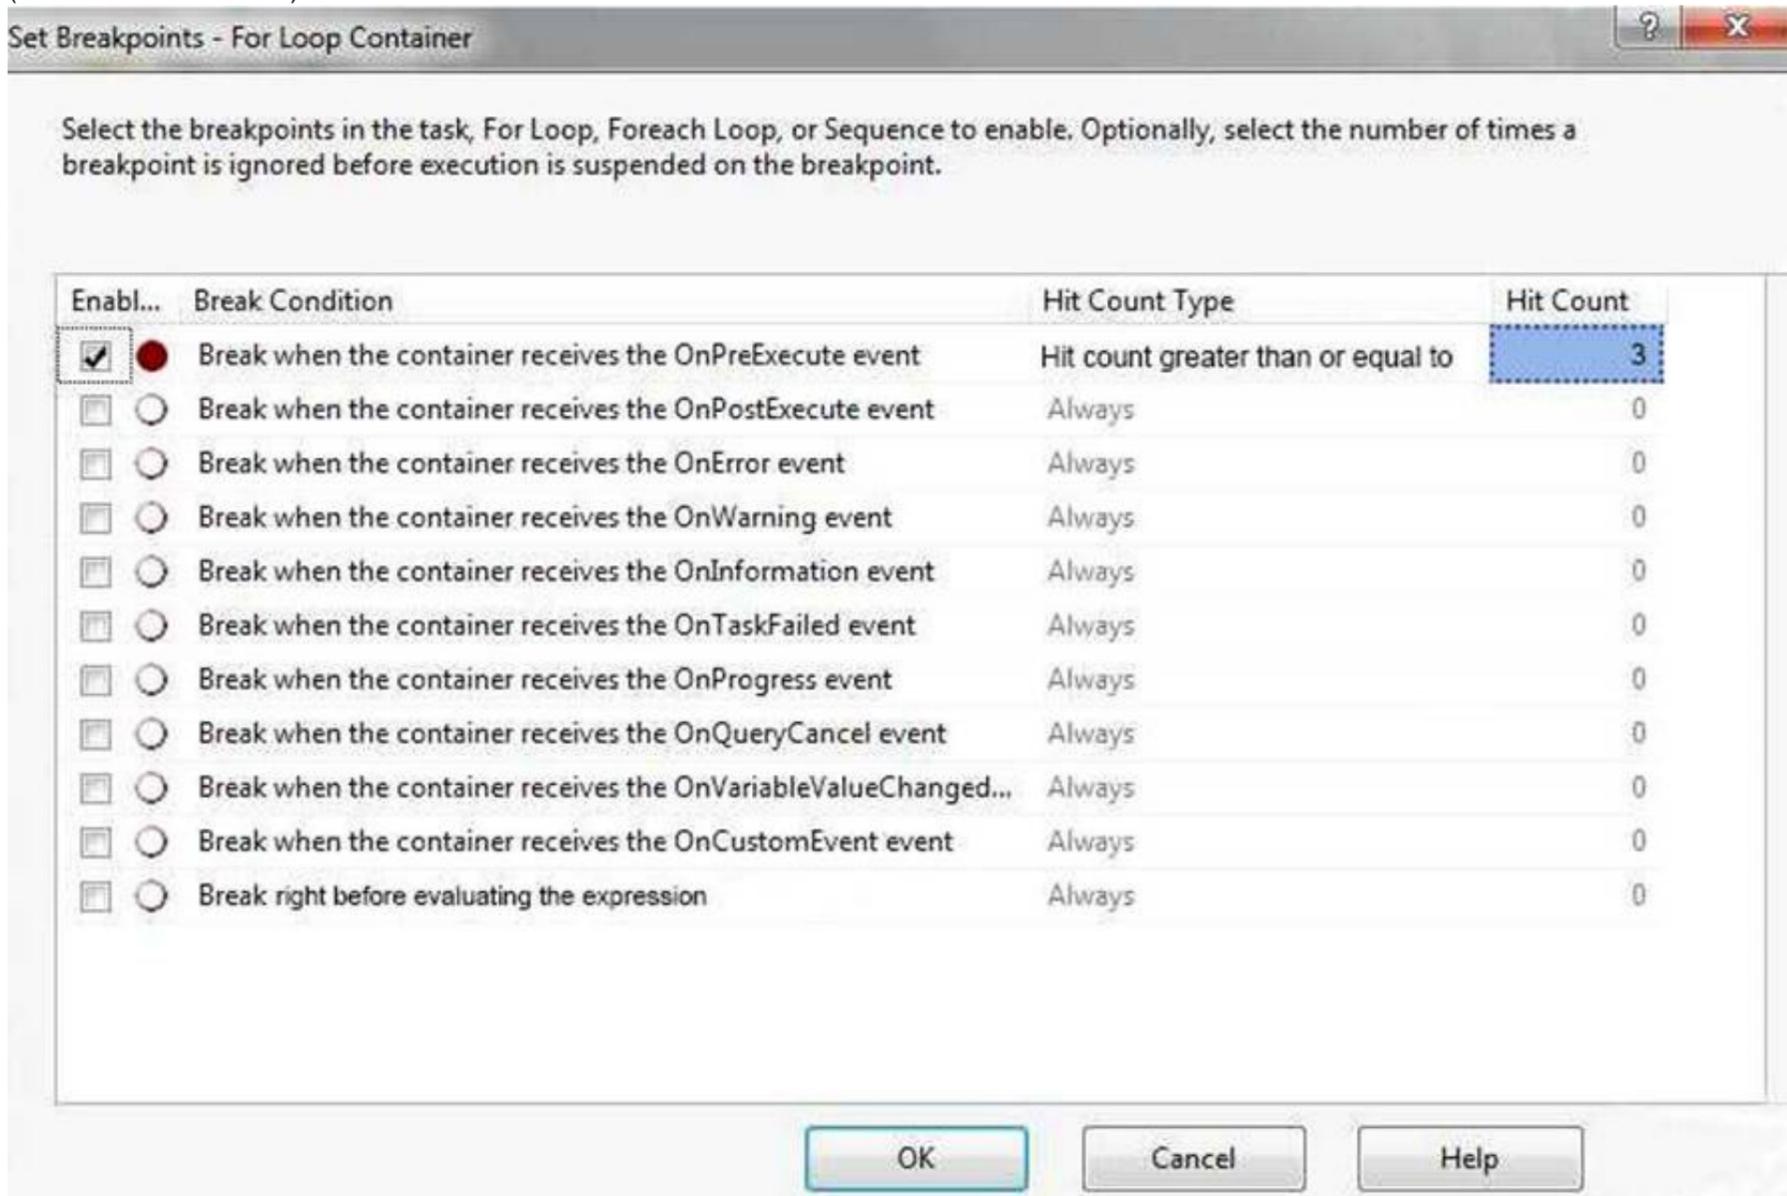
You have the Microsoft SQL Server Integration Services (SSIS) package shown in the Control flow exhibit. (Click the Exhibit button.)



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 66

You have a server that has Data Quality Services (DQS) installed. You create a matching policy that contains one matching rule. You need to configure the Similarity of Similar percentage that defines a match. Which similarity percentage will always generate a similarity score of 0?

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

Answer: A

Explanation:

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

References:

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

NEW QUESTION 67

You are developing a Microsoft SQL Server Master Data Services (MDS) solution. The model contains an entity named Product. The Product entity has three user-defined attributes named category, Subcategory, and Price, respectively. You need to ensure that combinations of values stored in the category and subcategory attributes are unique. What should you do?

- A. Create a derived hierarchy based on the category and subcategory attribute
- B. Use the category attribute as the top level for the hierarchy.
- C. Publish two business rules, one for each of the Category and Subcategory attributes.
- D. Set the value of the Attribute Type property for the Category and Subcategory attributes to Domain-based.
- E. Create a custom index that will be used by the Product entity.

Answer: D

NEW QUESTION 68

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

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

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

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

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 71

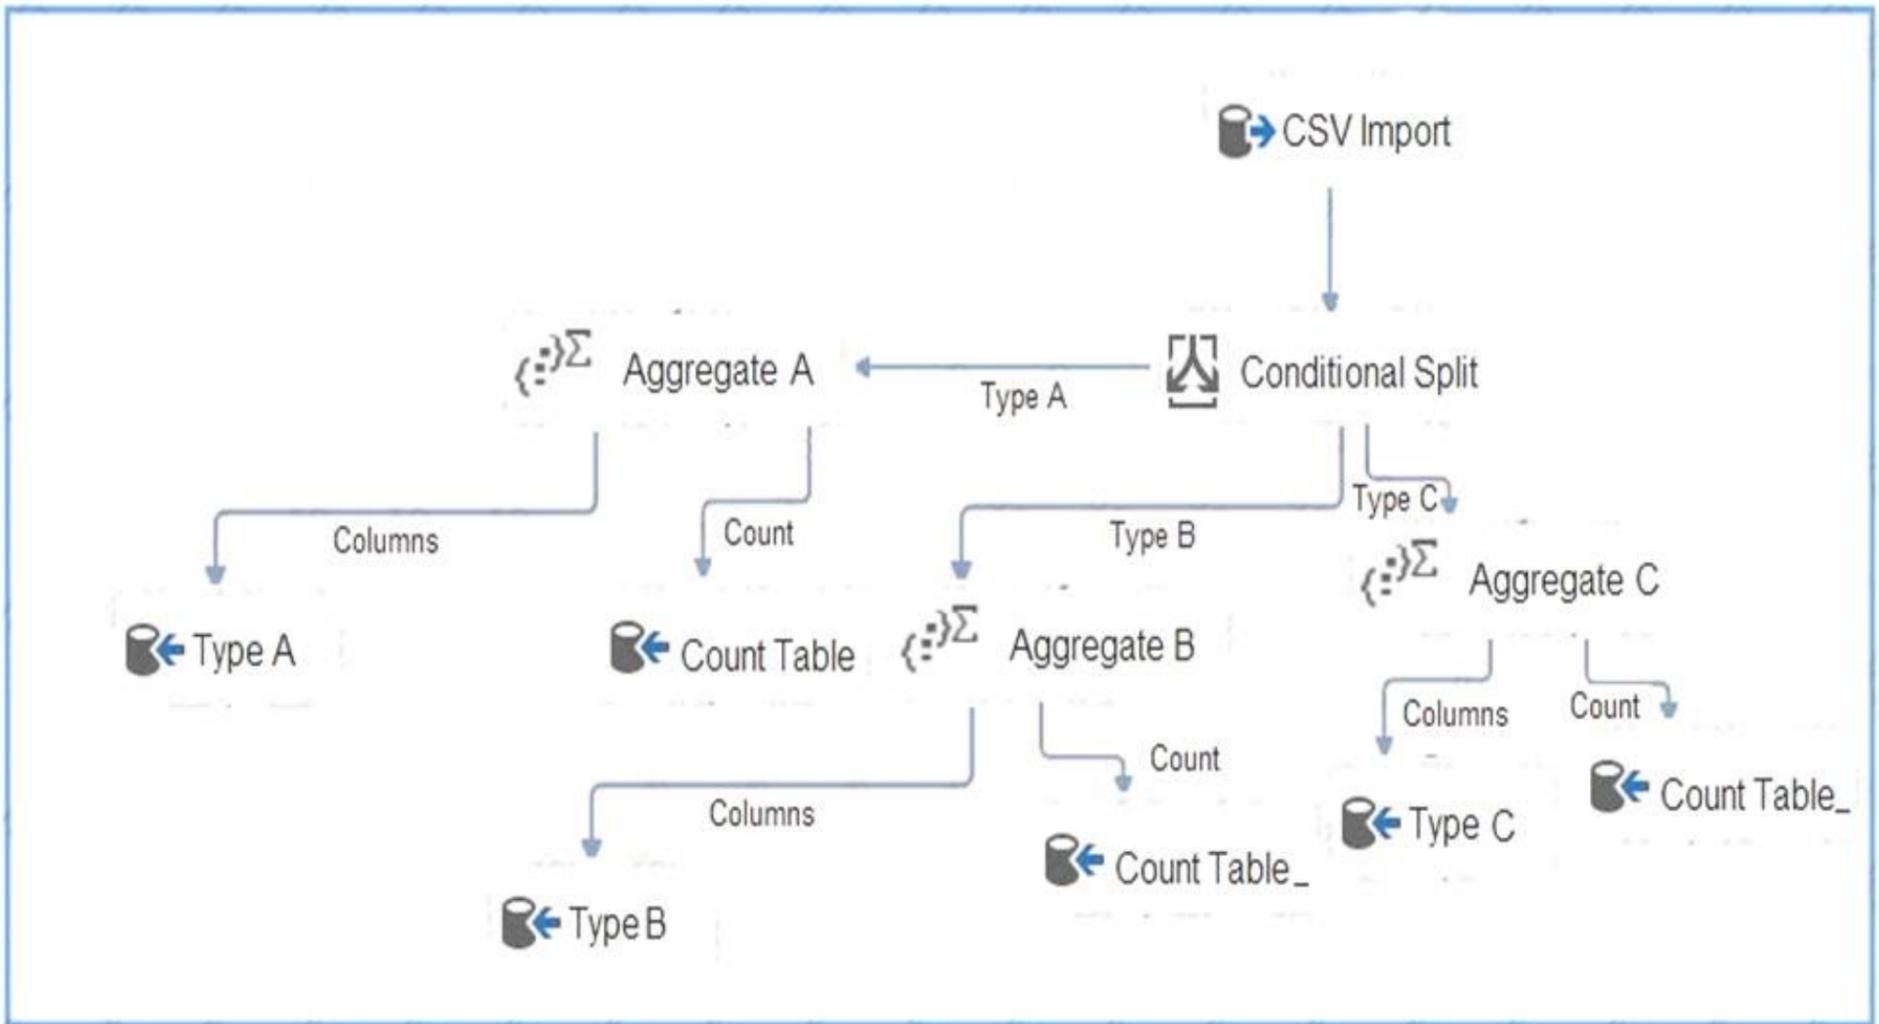
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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 75

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

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

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

 Messages

Msg 2627, Level 14, State 1, Line 13

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

The statement has been terminated.

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

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

You remove the constraint on the EmployeeSSN column. You need to ensure that values in the EmployeeSSN column are unique. For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

	Yes	No
Creating a clustered unique index on the EmployeeSSN column solves the issue.	<input type="radio"/>	<input type="radio"/>
Creating a filtered unique index on the EmployeeSSN column solves the issue.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 76

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

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

The table definitions are as follows:

```
CREATE TABLE [source].Table1
(
    ID INT NULL,
    CreateDate VARCHAR(100) NULL,
    Date1 DATETIME2(7) NULL,
    Number1 VARCHAR(100) NULL
```

```
CREATE TABLE error.Table1
(
  ID INT NULL,
  CreateDate VARCHAR(100) NULL,
  Date1 DATETIME2(7) NULL,
  Number1 VARCHAR(100) NULL,
  ErrorDescription VARCHAR(255) NULL
)
```

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

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 78

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

You are developing a Microsoft SQL Server Integration Services (SSIS) package.

You need to ensure that the packa

ge records the current Log Sequence Number (LSN) in the source database before the package begins reading source tables.

Which SSIS Toolbox item should you use?

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

Answer: A

Explanation:

The CDC Control task is used to control the life cycle of change data capture (CDC) packages. It handles CDC package synchronization with the initial load package, the management of Log Sequence Number (LSN) ranges that are processed in a run of a CDC package.

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

NEW QUESTION 79

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 profile the data.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Use a matching rule. References:

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

NEW QUESTION 84

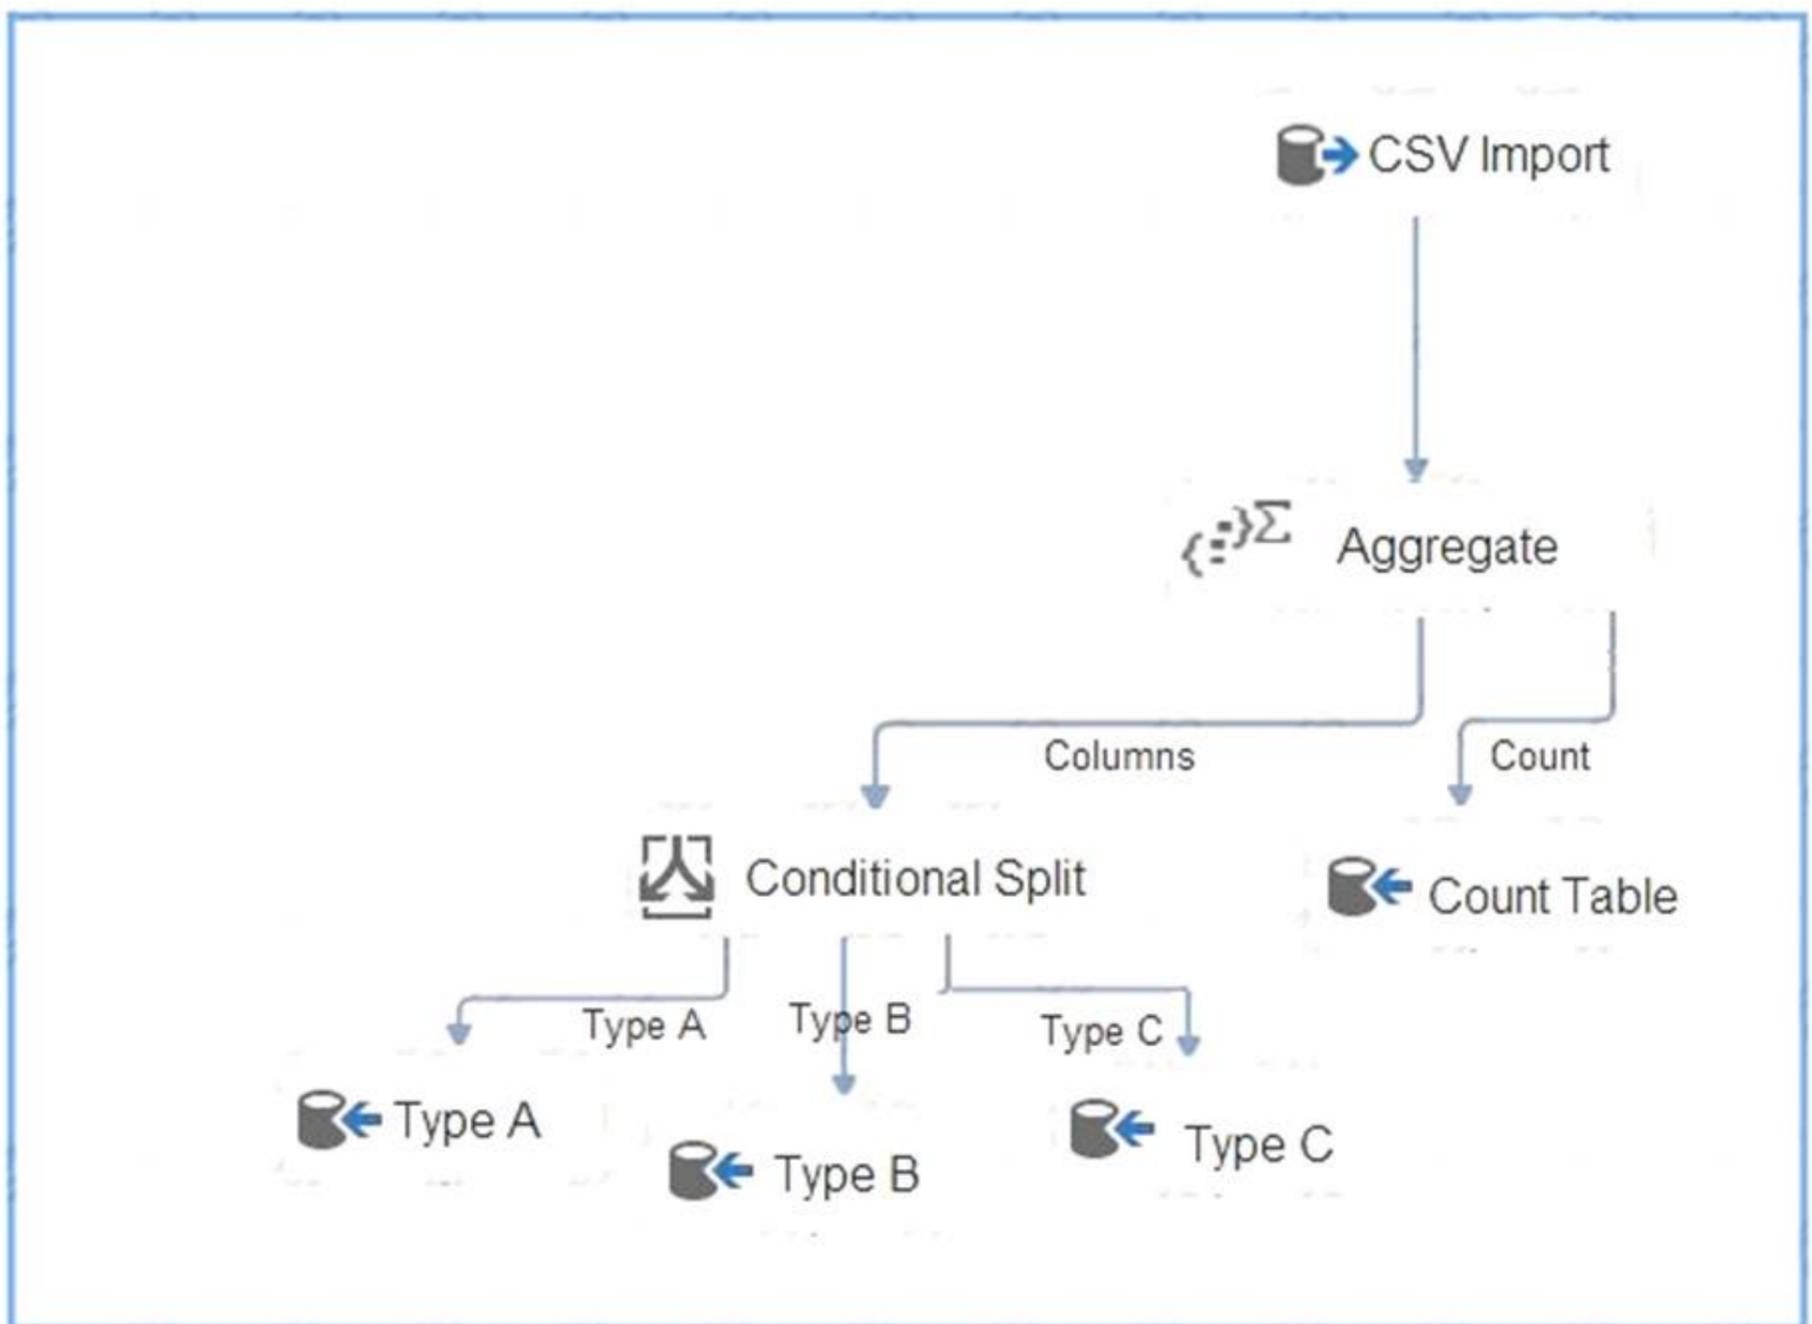
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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 86

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

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

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

A new version of that integration Services package is released that introduces several errors in the loading process.

You need to roll back the Integration Services package to the previous version. Which stored procedure should you execute?

- A. catalog.deploy_project
- B. catalog.restore_project
- C. catalog.stop.operation
- D. sys.sp_cdc.addJob
- E. sys.sp.cdc.changejob

Answer: B

Explanation:

catalog.restore_project restores a project in the Integration Services catalog to a previous version. References:

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

NEW QUESTION 91

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 94

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	
Dimension.Customer	
Dimension.Date	

Table	Technology
Dimension.SalesTerritory	<ul style="list-style-type: none"> Change Data Capture (CDC) Change Tracking Temporal table Microsoft SQL Server snapshot replication
Dimension.Customer	<ul style="list-style-type: none"> Change Data Capture (CDC) Change Tracking Temporal table Microsoft SQL Server snapshot replication
Dimension.Date	<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 97

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

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

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

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

You need to modify Dim1 to ensure that the table can be used for both fact tables.

Which two actions should you perform? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

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

Answer: AD

NEW QUESTION 100

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 70-767 Exam with Our Prep Materials Via below:

<https://www.certleader.com/70-767-dumps.html>