

70-767 Dumps

Implementing a SQL Data Warehouse (beta)

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



NEW QUESTION 1

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

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

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

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

The query fails to return results.

You need to determine why the query fails.

Solution: You run the following Transact-SQL statements:

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

We must use Label, not QueryID in the WHERE clause. References:

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

NEW QUESTION 2

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

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

You are not permitted to make changes to the client applications. You need to optimize the storage for the data warehouse.

What change should you make?

- A. Partition the Fact.Order table, and move historical data to new filegroups on lower-cost storage.
- B. Create new tables on lower-cost storage, move the historical data to the new tables, and then shrink the database.
- C. Remove the historical data from the database to leave available space for new data.
- D. Move historical data to new tables on lower-cost storage.

Answer: A

Explanation:

Create the load staging table in the same filegroup as the partition you are loading. Create the unload staging table in the same filegroup as the partition you are deleting.

From scenario: Data older than one year is accessed infrequently and is considered historical.

References:

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

NEW QUESTION 3

You need to ensure that a downstream system can consume data in a Master Data Services (MDS) system. What should you configure?

- A. a Data Collector
- B. a knowledgebase
- C. a matching policy
- D. a subscription view

Answer: D

Explanation:

Subscription views to consume your master data. References:

<https://docs.microsoft.com/en-us/sql/master-data-services/master-data-services-overview-mds?view=sql-server->

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 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.

Domain	Similarity	Weight	Prerequisite
Gender	Exact	10	<input checked="" type="checkbox"/>
City	Exact	10	<input checked="" type="checkbox"/>
State	Exact	10	<input checked="" type="checkbox"/>
Country	Exact	10	<input checked="" type="checkbox"/>
Zip	Exact	10	<input checked="" type="checkbox"/>
Birth Date	Similar	34	<input type="checkbox"/>
Address Line 1	Similar	10	<input type="checkbox"/>
First Name	Similar	33	<input type="checkbox"/>
Last Name	Similar	23	<input type="checkbox"/>

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

You verify that the data contains multiple records that have similar address values, and you expect some of the records to match. You need to increase the likelihood that the records will match when they have similar address values.

Solution: You decrease the relative weights for Address Line 1 of the matching policy. Does this meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 6

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 7

You are implementing a Microsoft SQL Server data warehouse with a multi-dimensional data model. Orders are stored in a table named Factorder. The addresses that are associated with all orders are stored in a fact table named FactAddress. A key in the FactAddress table specifies the type of address for an order.

You need to ensure that business users can examine the address data by either of the following:

- shipping address and billing address
- shipping address or billing address type Which data model should you use?

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

Answer: H

NEW QUESTION 8

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

After you answer a question in this 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 9

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

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

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

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

You have the following requirements:

- ▶ Implement table partitioning to improve the manageability of the data warehouse and to avoid the need to repopulate all transactional data each night. Use a partitioning strategy that is as granular as possible.
- ▶ Partition the Fact.Order table and retain a total of seven years of data.
- ▶ Partition the Fact.Ticket table and retain seven years of data. At the end of each month, the partition structure must apply a sliding window strategy to ensure that a new partition is available for the upcoming month, and that the oldest month of data is archived and removed.
- ▶ Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.
- ▶ 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 one 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 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 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 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 12

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?

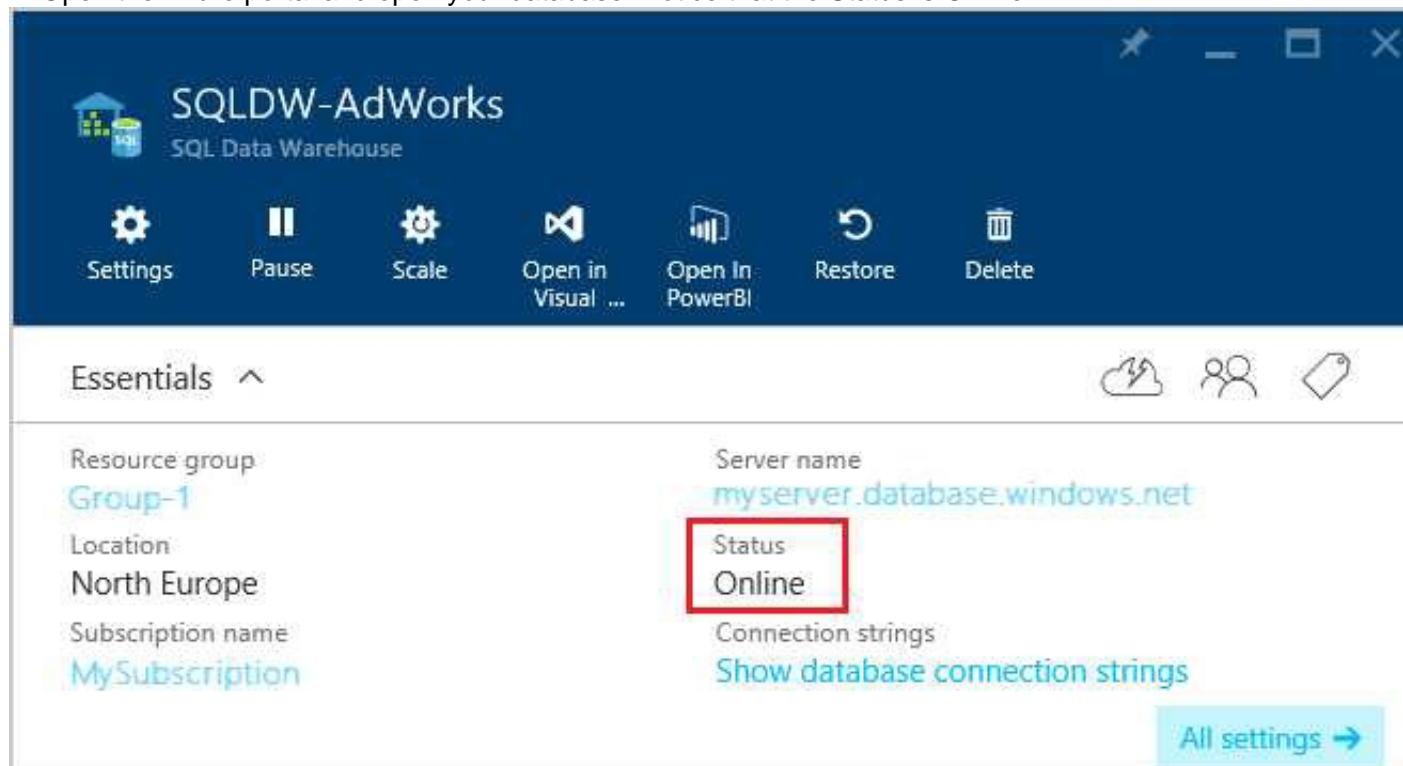
- A. Yes
B. No

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 15

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 5 PERCENT
```

Does the solution meet the goal?

- A. Yes
B. No

Answer: A

Explanation:

You can specify the sample size as a percent. A 5% statistics sample size would be helpful.

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

NEW QUESTION 16

You manage Master Data Services (MDS).

You need to create a new entity with the following requirements:

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

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

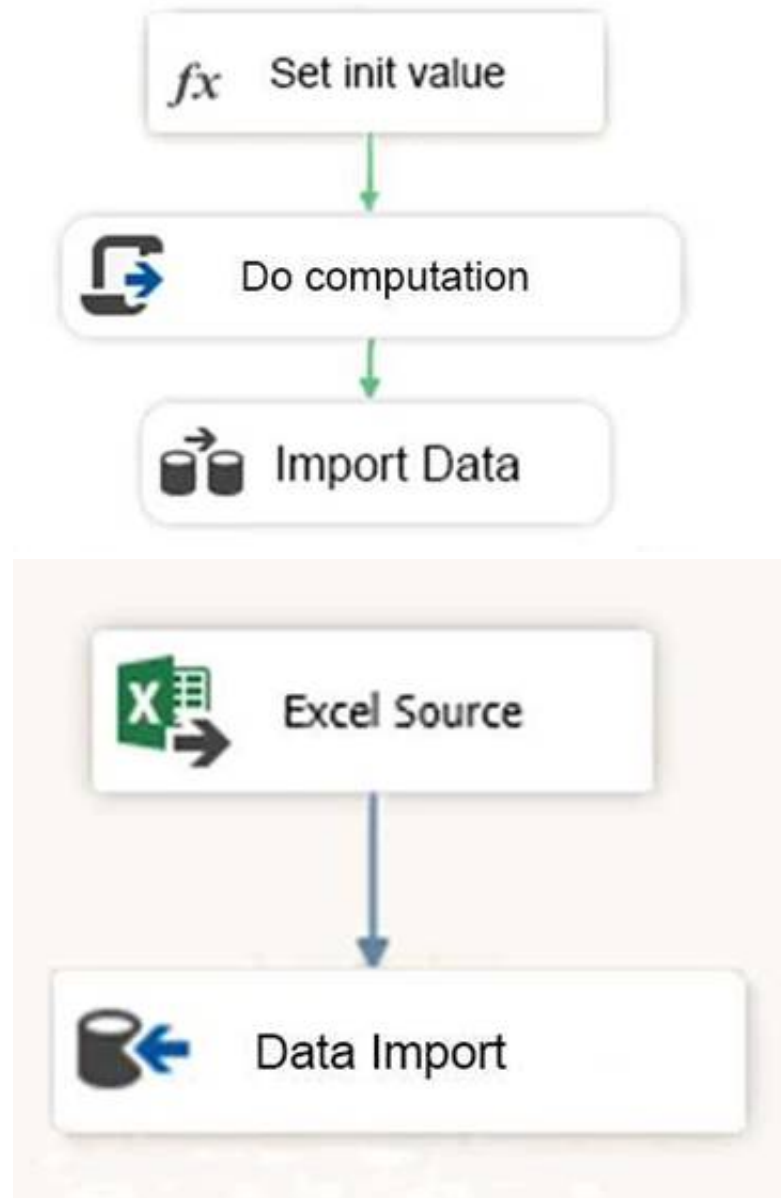
- A. Full
B. None
C. Attribute
D. Member

E. Simple

Answer: D

NEW QUESTION 21

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



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

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

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

Answer Area

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

References:

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

NEW QUESTION 26

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.)

The screenshot shows the 'DQS Cleansing Transformation Editor' dialog box with the 'Advanced' tab selected. The dialog is titled 'Configure the properties used to correct the data of an input column.' It features three tabs: 'Connection Manager', 'Mapping', and 'Advanced'. The 'Advanced' tab contains several checked options: 'Standardize output', 'Confidence', 'Reason', 'Appended Data (additional data received from reference data provider)', and 'Appended Data Schema'. Below these options is a 'Configure error output:' section with a dropdown menu set to 'Fail component'. At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

DQS Cleansing Transformation Editor

Configure the properties used to correct the data of an input column.

Connection Manager Mapping Advanced

☒ Standardize output

Enable field level columns:

☒ Confidence

☒ Reason

Enable record level columns:

☒ Appended Data (additional data received from reference data provider)

☒ Appended Data Schema

Configure error output: Fail component

OK Cancel Help

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.	<input type="radio"/>	<input type="radio"/>
You can use the EmailAddress_Status column to split the output.	<input type="radio"/>	<input type="radio"/>
You can use the EmailAddress_Reason column to split the output.	<input type="radio"/>	<input type="radio"/>

- 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 27

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 32

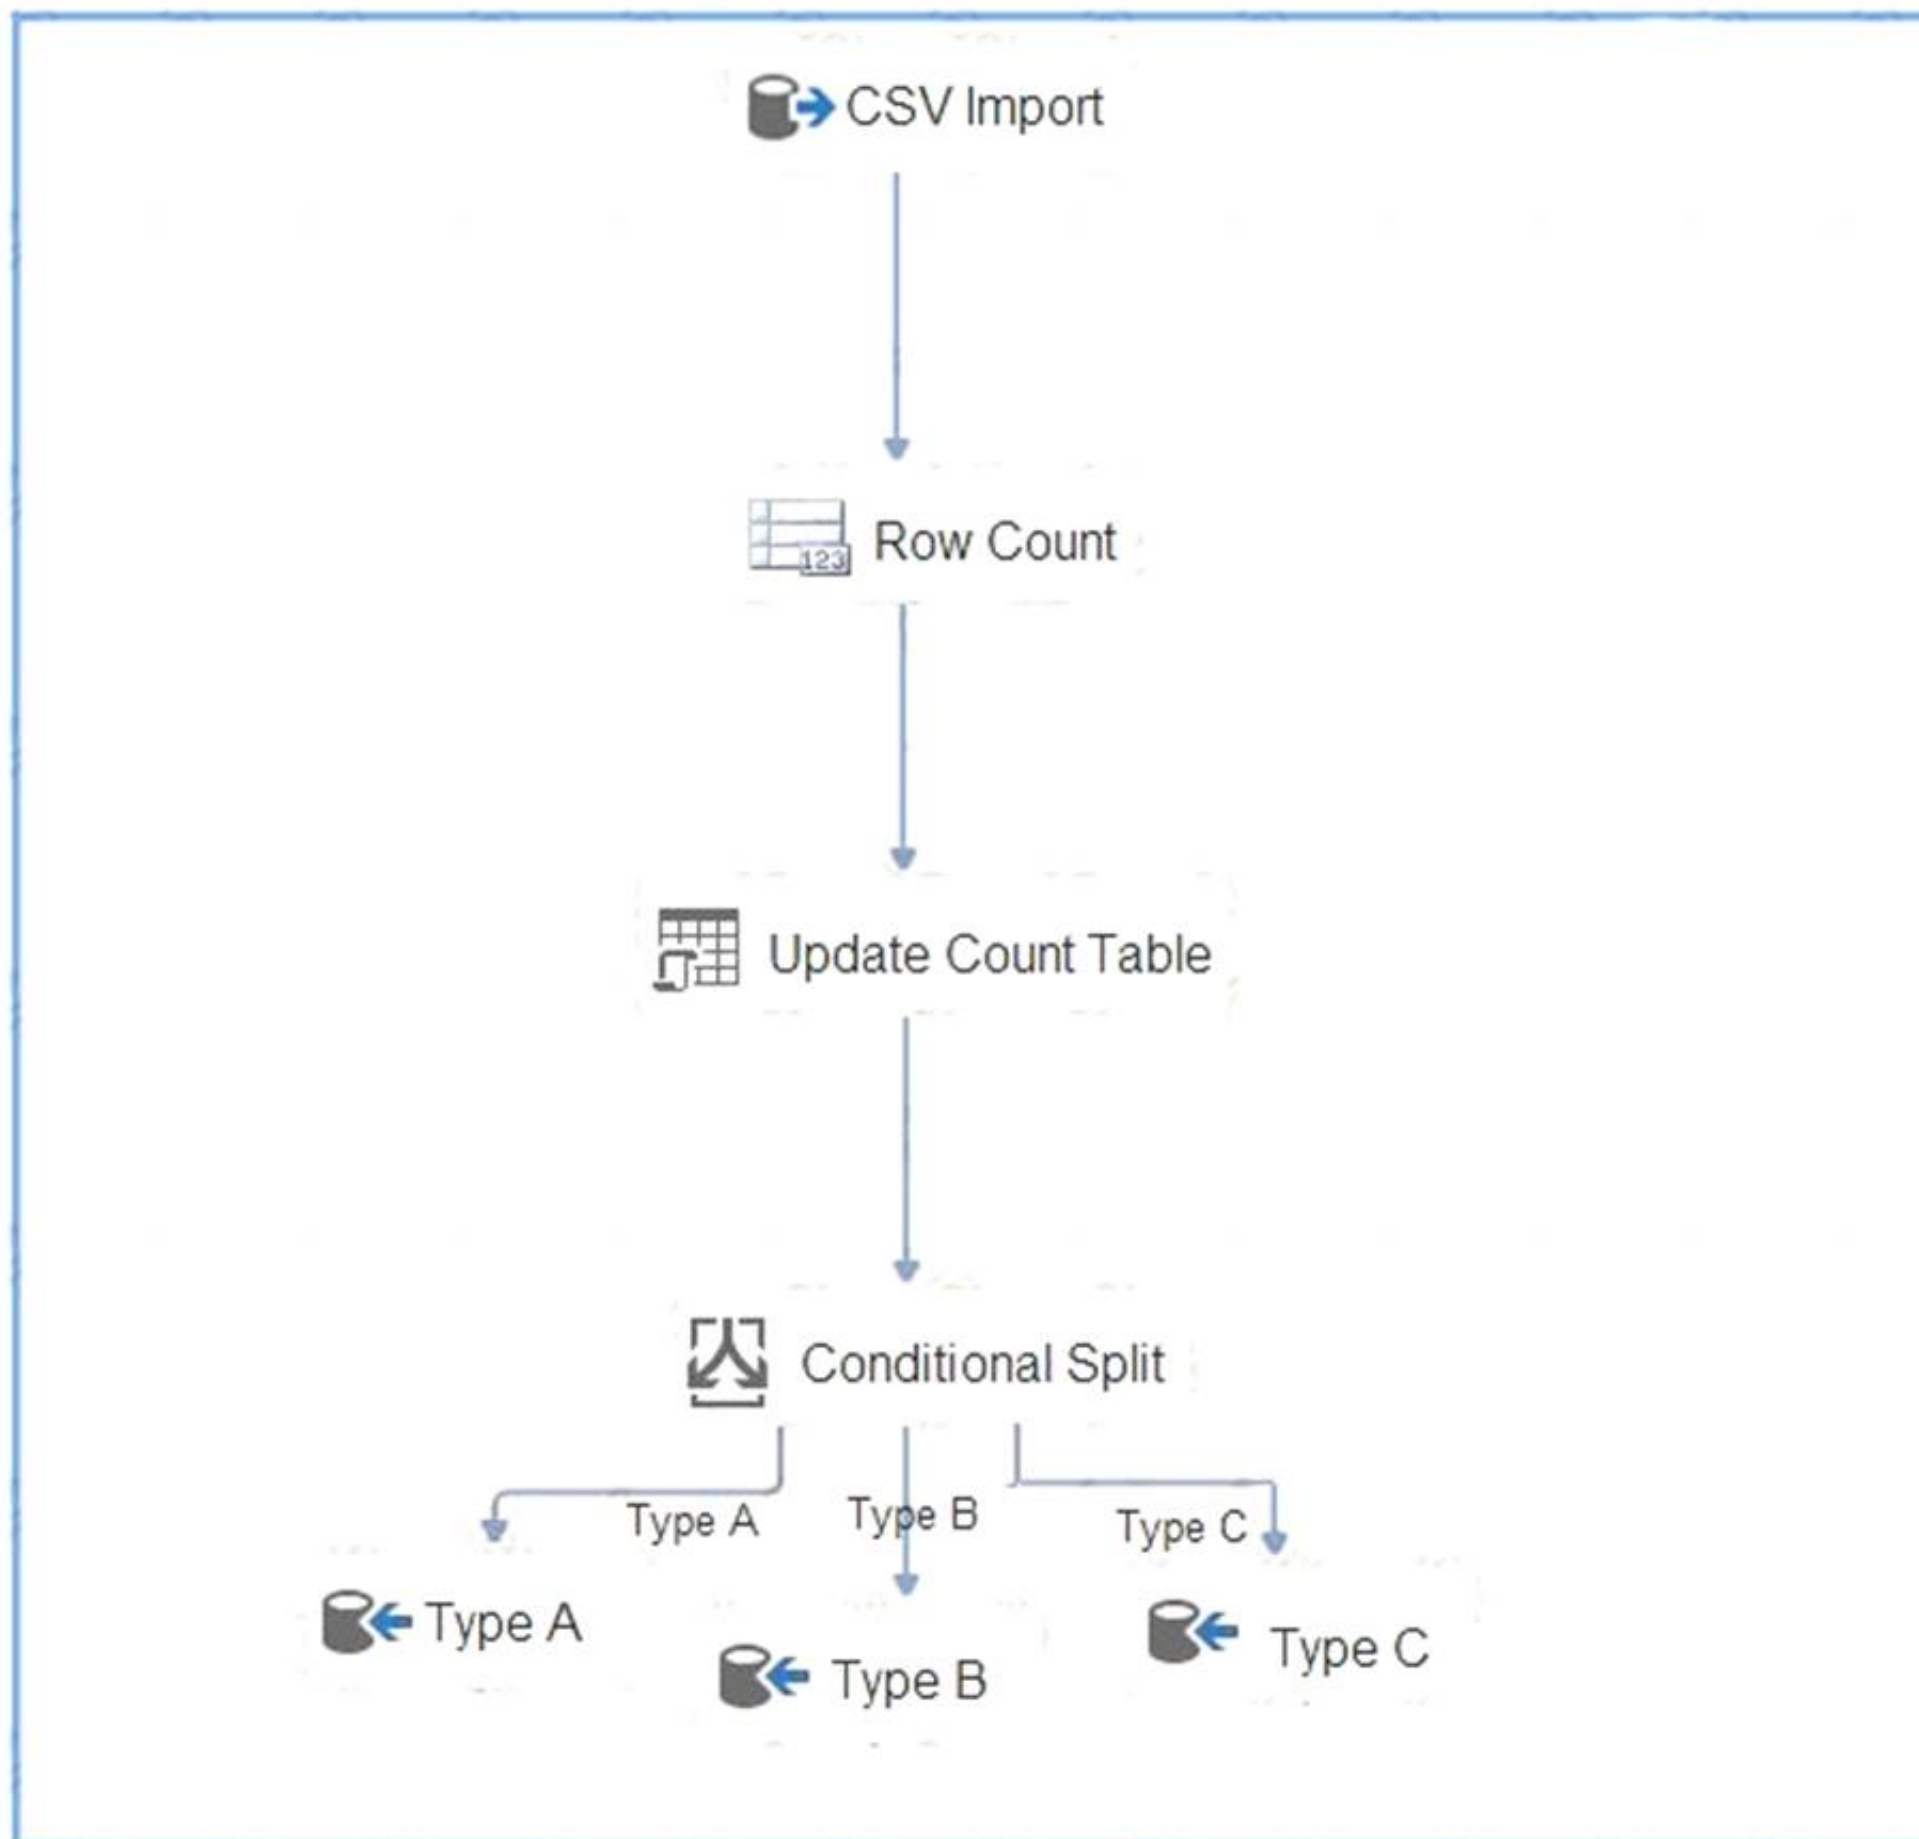
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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 37

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 41

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 42

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

	▼
CLUSTERED	
NONCLUSTERED	

 INDEX lx_product

ON dbo.Products

	▼
(ProductReference)	
(QuantityInStock)	
(ProductName)	

INCLUDE

	▼
(Products)	
(ProductReference)	
(ProductName, ProductReference)	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

CREATE

	▼
CLUSTERED	
NONCLUSTERED	

 INDEX lx_product

ON dbo.Products

	▼
(ProductReference)	
(QuantityInStock)	
(ProductName)	

INCLUDE

	▼
(Products)	
(ProductReference)	
(ProductName, ProductReference)	

NEW QUESTION 44

You have a data warehouse named DW1 that contains 20 years of data. DW1 contains a very large fact table. New data is loaded to the fact table monthly.

Many reports query DW1 for the past year of data. Users frequently report that the reports are slow.

You need to modify the fact table to minimize the amount of time it takes to run the reports. The solution must ensure that other reports can continue to be generated from DW1.

What should you do?

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

Answer: A

Explanation:

We use ALTER TABLE to partition the table.

NEW QUESTION 48

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

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

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 53

You have a database named DB1 that contains millions of rows. You plan to perform a weekly audit of the changes to the rows.

You need to ensure that you can view which rows were modified and the hour that the modification occurred. What should you do?

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

Answer: D

Explanation:

SQL Server 2017 provides two features that track changes to data in a database: change data capture and change tracking.

Change data capture provides historical change information for a user table by capturing both the fact that DML changes were made and the actual data that was changed. Changes are captured by using an asynchronous process that reads the transaction log and has a low impact on the system.

References:

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

NEW QUESTION 55

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 Management Studio (SSMS).

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 57

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

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

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

You remove the Integration Services job.

You need to stop tracking changes to the database. The solution must remove all the change data capture configurations from DB1.

Which stored procedure should you execute?

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

Answer: F

Explanation:

sys.sp_cdc_disable_db disables change data capture for all tables in the database currently enabled. All system objects related to change data capture, such as change tables, jobs, stored procedures and functions, are dropped.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-disable-db-transa>

NEW QUESTION 60

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

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

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

What should you add to the model?

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

Answer: B

Explanation:

A Master Data Services derived hierarchy is derived from the domain-based attribute relationships that already exist between entities in a model.

You can create a derived hierarchy to highlight any of the existing domain-based attribute relationships in the model.

NEW QUESTION 62

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>▼</div> <div> Turn off semi-additive behavior. Enable the First Child semi-additive function. Enable the FirstNonEmpty semi-additive function. Enable the LastNoneEmpty semi-additive function. Enable the Count semi-additive function. Enable the None semi-additive function. </div>
Measure2	<div>▼</div> <div> Turn off semi-additive behavior. Enable the First Child semi-additive function. Enable the FirstNonEmpty semi-additive function. Enable the LastNoneEmpty semi-additive function. Enable the Count semi-additive function. Enable the None semi-additive function. </div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

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

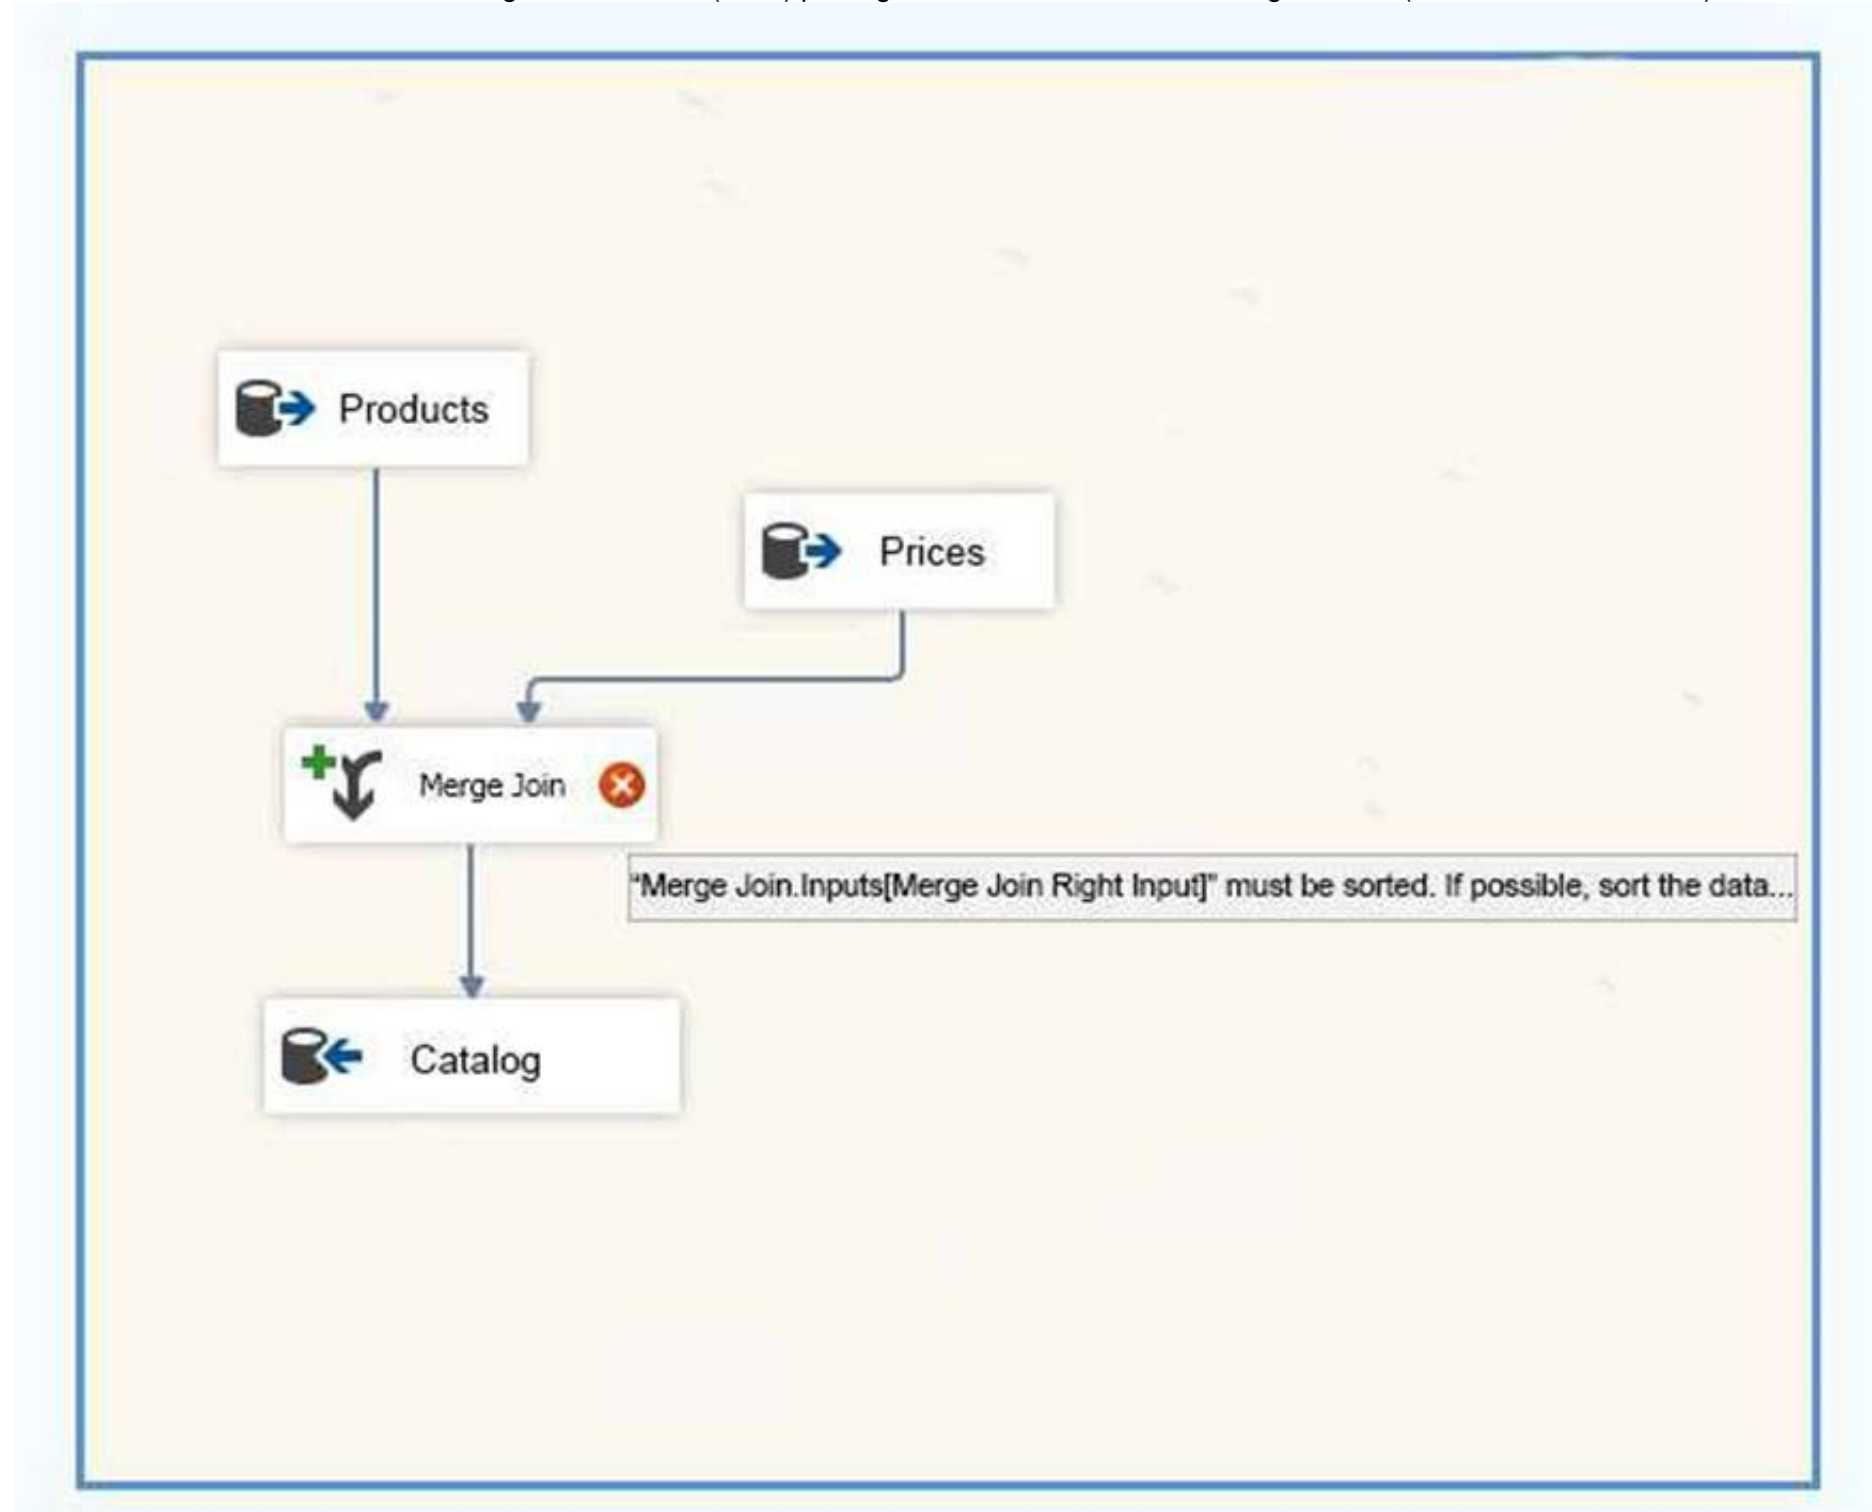
FirstNonEmpty: The member value is evaluated as the value of its first child along the time dimension that contains data.

References:

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

NEW QUESTION 64

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.)

OLE DB Source Editor

Configure the properties used by a data flow to obtain data from any OLE DB provider.

Connection Manager
Columns
Error Output

Specify an OLE DB connection manager, a data source, or a data source view, and select the data access mode. If using the SQL command access mode, specify the SQL command either by typing the query or by using Query Builder.

OLE DB connection manager:
Catalog

Data access mode:
SQL command

SQL command text:
SELECT ReferenceNr, Price
FROM dbo.Prices
ORDER BY ReferenceNr

Parameters...
Build Query...
Browse...
Parse Query

Preview...

OK Cancel Help

Advanced Editor for OLE DB Source

The advanced editor provides access to the low-level properties of data flow components. Additionally, the advanced editor can be used to configure components that do not have a custom user interface.

Connection Managers | Component Properties | Column Mappings | **Input and Output Properties**

Specify properties for the inputs and outputs of the data flow component.

Inputs and outputs:

- OLE DB Source Output
 - External Columns
 - Output Columns
 - ReferenceNr
 - Price
- OLE DB Source Error Output

Common Properties

ComparisonFlags	
Description	
ErrorOrTruncationOperation	Conversion
ErrorRowDisposition	RD_FailComponent
ExternalMetadataColumnID	111
ID	112
IdentificationString	Prices.Outputs[OLE DB Source
LineageID	112
MappedColumnID	0
Name	ReferenceNr
SortKeyPosition	0
SpecialFlags	0
TruncationRowDisposition	RD_FailComponent

Data Type Properties

CodePage	1252
DataType	string [DT_STR]
Length	100
Precision	0

ID

Buttons: Add Output, Add Column, Remove Output, Remove Column, Refresh, OK, Cancel, Help

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 65

You have a database named OnlineSales that contains a table named Customers. You plan to copy incremental changes from the Customers table to a data warehouse every hour.
You need to enable change tracking for the Customers table.
How should you complete the Transact-SQL statements? 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

DATABASE [OnlineSales]

CHANGE_TRACKING = ON

SYSTEM_VERSIONING = ON

RECOVERY FULL

TABLE [dbo].[Customers]

ENABLE CHANGE_TRACKING

QUERY_STORE = ON

ENABLE_BROKER

Answer Area

ALTER

SET

(CHANGE_RETENTION = 2 DAYS, AUTO_CLEANUP = ON)

ALTER

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:
Box 1: DATABASE [OnlineSales]
Before you can use change tracking, you must enable change tracking at the database level. The following example shows how to enable change tracking by using ALTER DATABASE.
ALTER DATABASE AdventureWorks2012 SET CHANGE_TRACKING = ON
(CHANGE_RETENTION = 2 DAYS, AUTO_CLEANUP = ON) Box 2: CHANGE_TRACKING = ON
ALTER SET CHANGE_RETENTION
Box 3: ALTER TABLE [dbo].[Customers]
Change tracking must be enabled for each table that you want tracked. When change tracking is enabled, change tracking information is maintained for all rows in

the table that are affected by a DML operation.

The following example shows how to enable change tracking for a table by using ALTER TABLE. ALTER TABLE Person.Contact ENABLE CHANGE_TRACKING

WITH (TRACK_COLUMNS_UPDATED = ON) Box 4: ENABLE CHANGE_TRACKING

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/track-changes/enable-and-disable-change-tracking-sql->

NEW QUESTION 67

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 69

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 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 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 75

You have a data warehouse named DW1.

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

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

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

How many partitions should you recommend?

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

Answer: B

NEW QUESTION 76

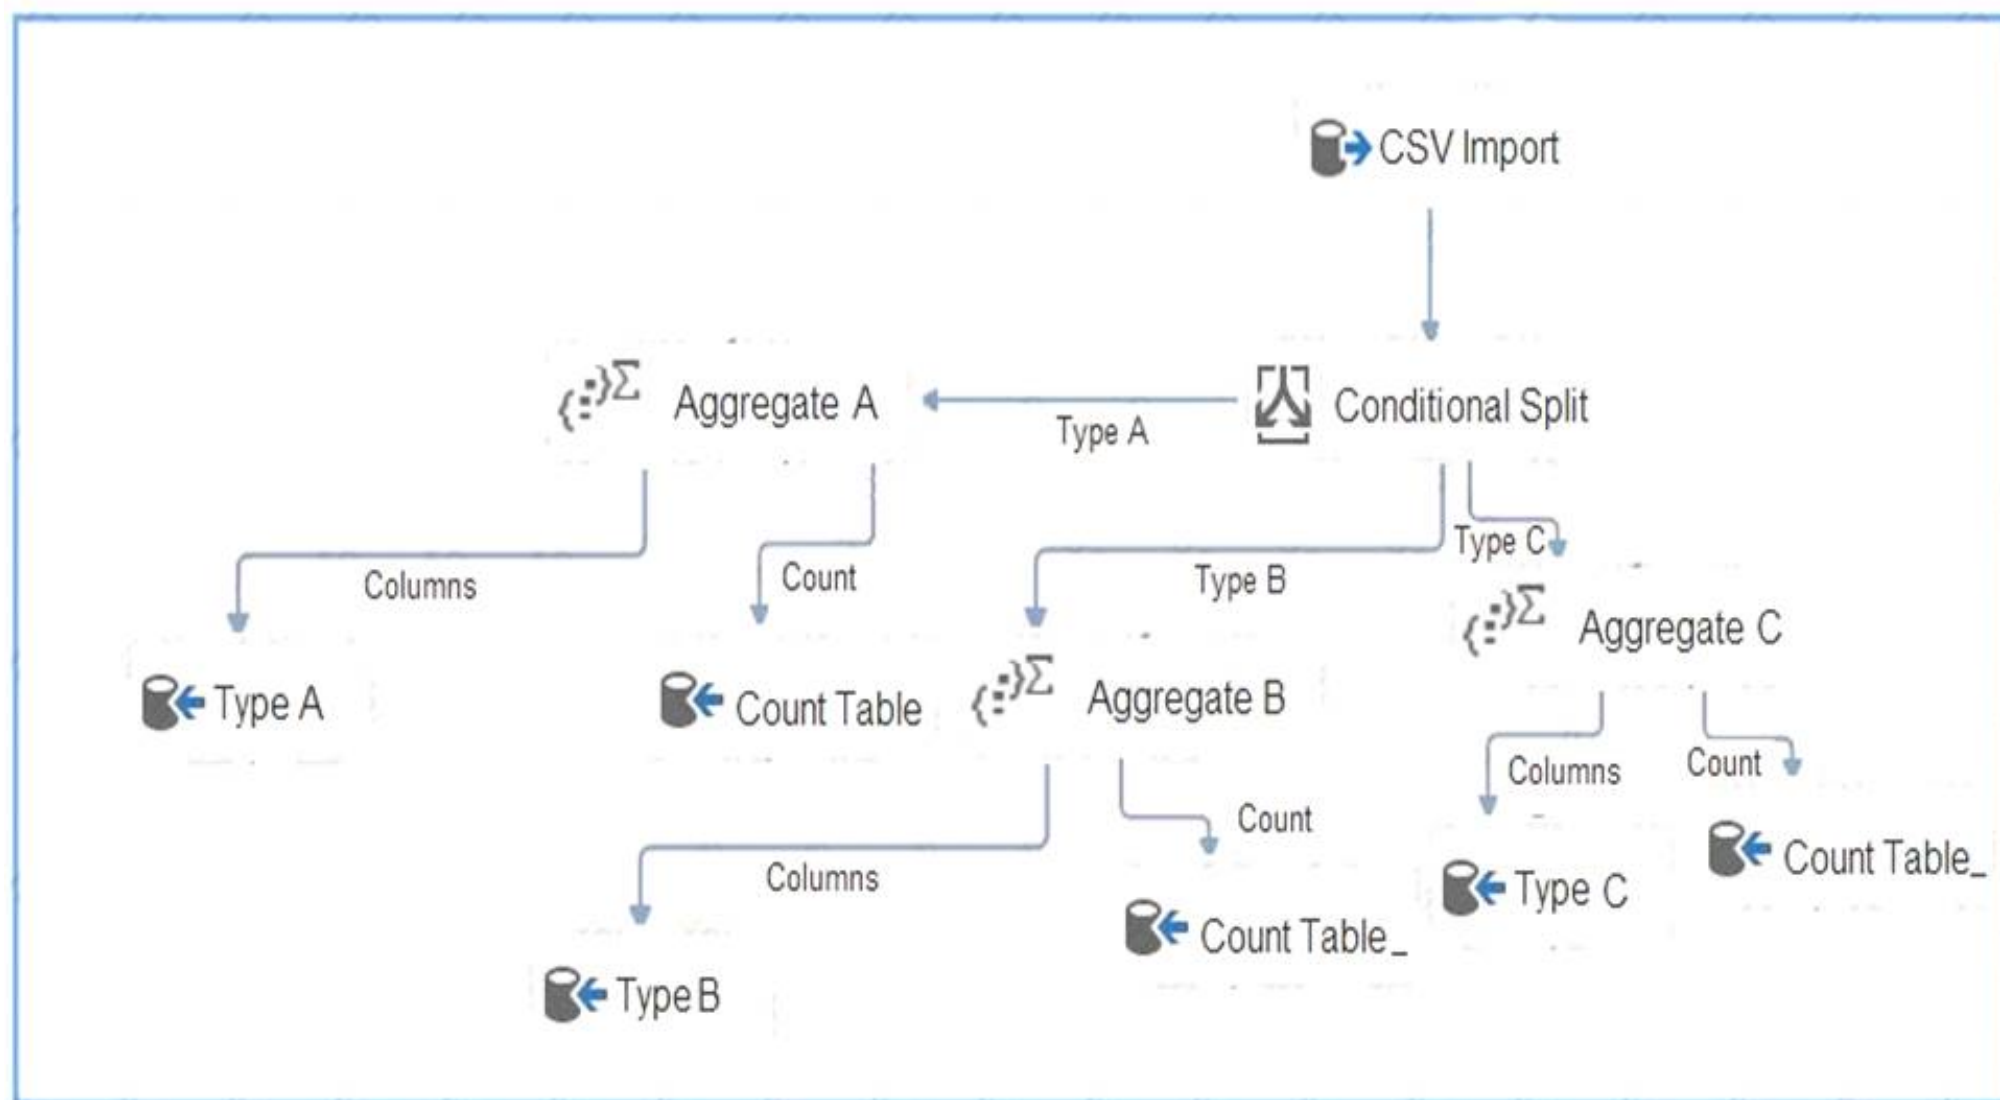
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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 79

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 in the cause of the error. (Click the Exhibit button.)

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

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

Answer Area

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

- 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 80

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 84

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 87

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 92

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 95

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

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

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

Each day, data from the table OnlineOrder in DB2 must be exported by partition. The tables must not be locked during the process.

You need to write a Microsoft SQL Server Integration Services (SSIS) package that performs the data export. What should you use?

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

Answer: E

Explanation:

The Union All transformation combines multiple inputs into one output. For example, the outputs from five different Flat File sources can be inputs to the Union All

transformation and combined into one output.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/union-all-transformation>

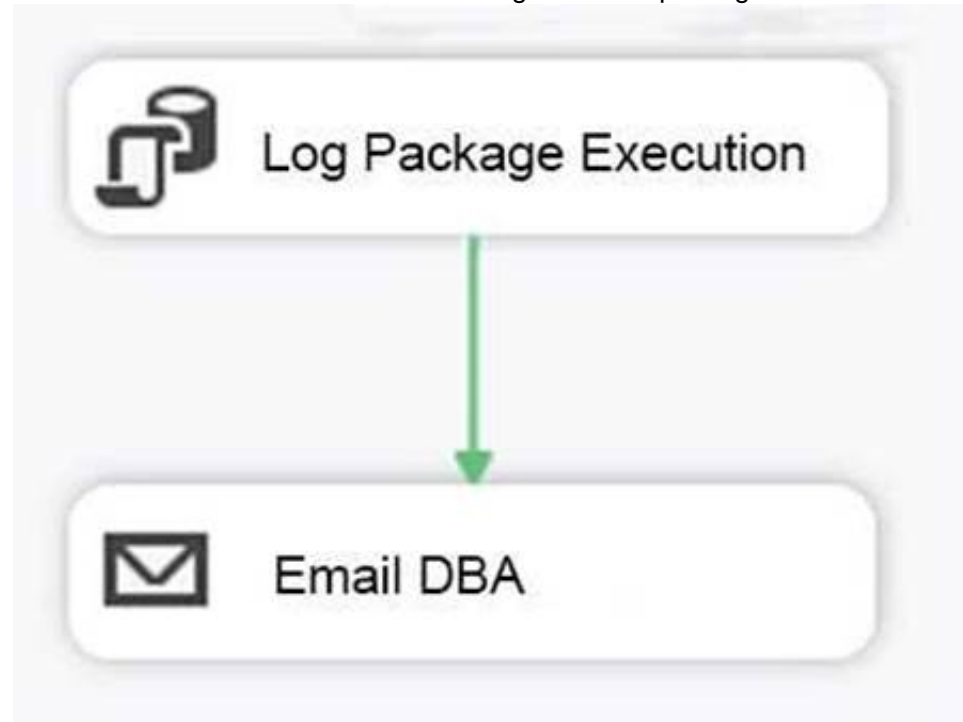
NEW QUESTION 97

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

After you answer a question in this 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 an ASP.NET assembly. You add a script task that references this assembly to each data warehouse load package.

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 100

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 104

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"/> <div> Change Data Capture (CDC) Change Tracking Temporal table Microsoft SQL Server snapshot replication </div>
Dimension.Customer	<input type="text"/> <div> Change Data Capture (CDC) Change Tracking Temporal table Microsoft SQL Server snapshot replication </div>
Dimension.Date	<input type="text"/> <div> Change Data Capture (CDC) Change Tracking Temporal table Microsoft SQL Server snapshot replication </div>

- 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 106

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 109

.....

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>