

Microsoft

Exam Questions 70-767

Implementing a SQL Data Warehouse (beta)



NEW QUESTION 1

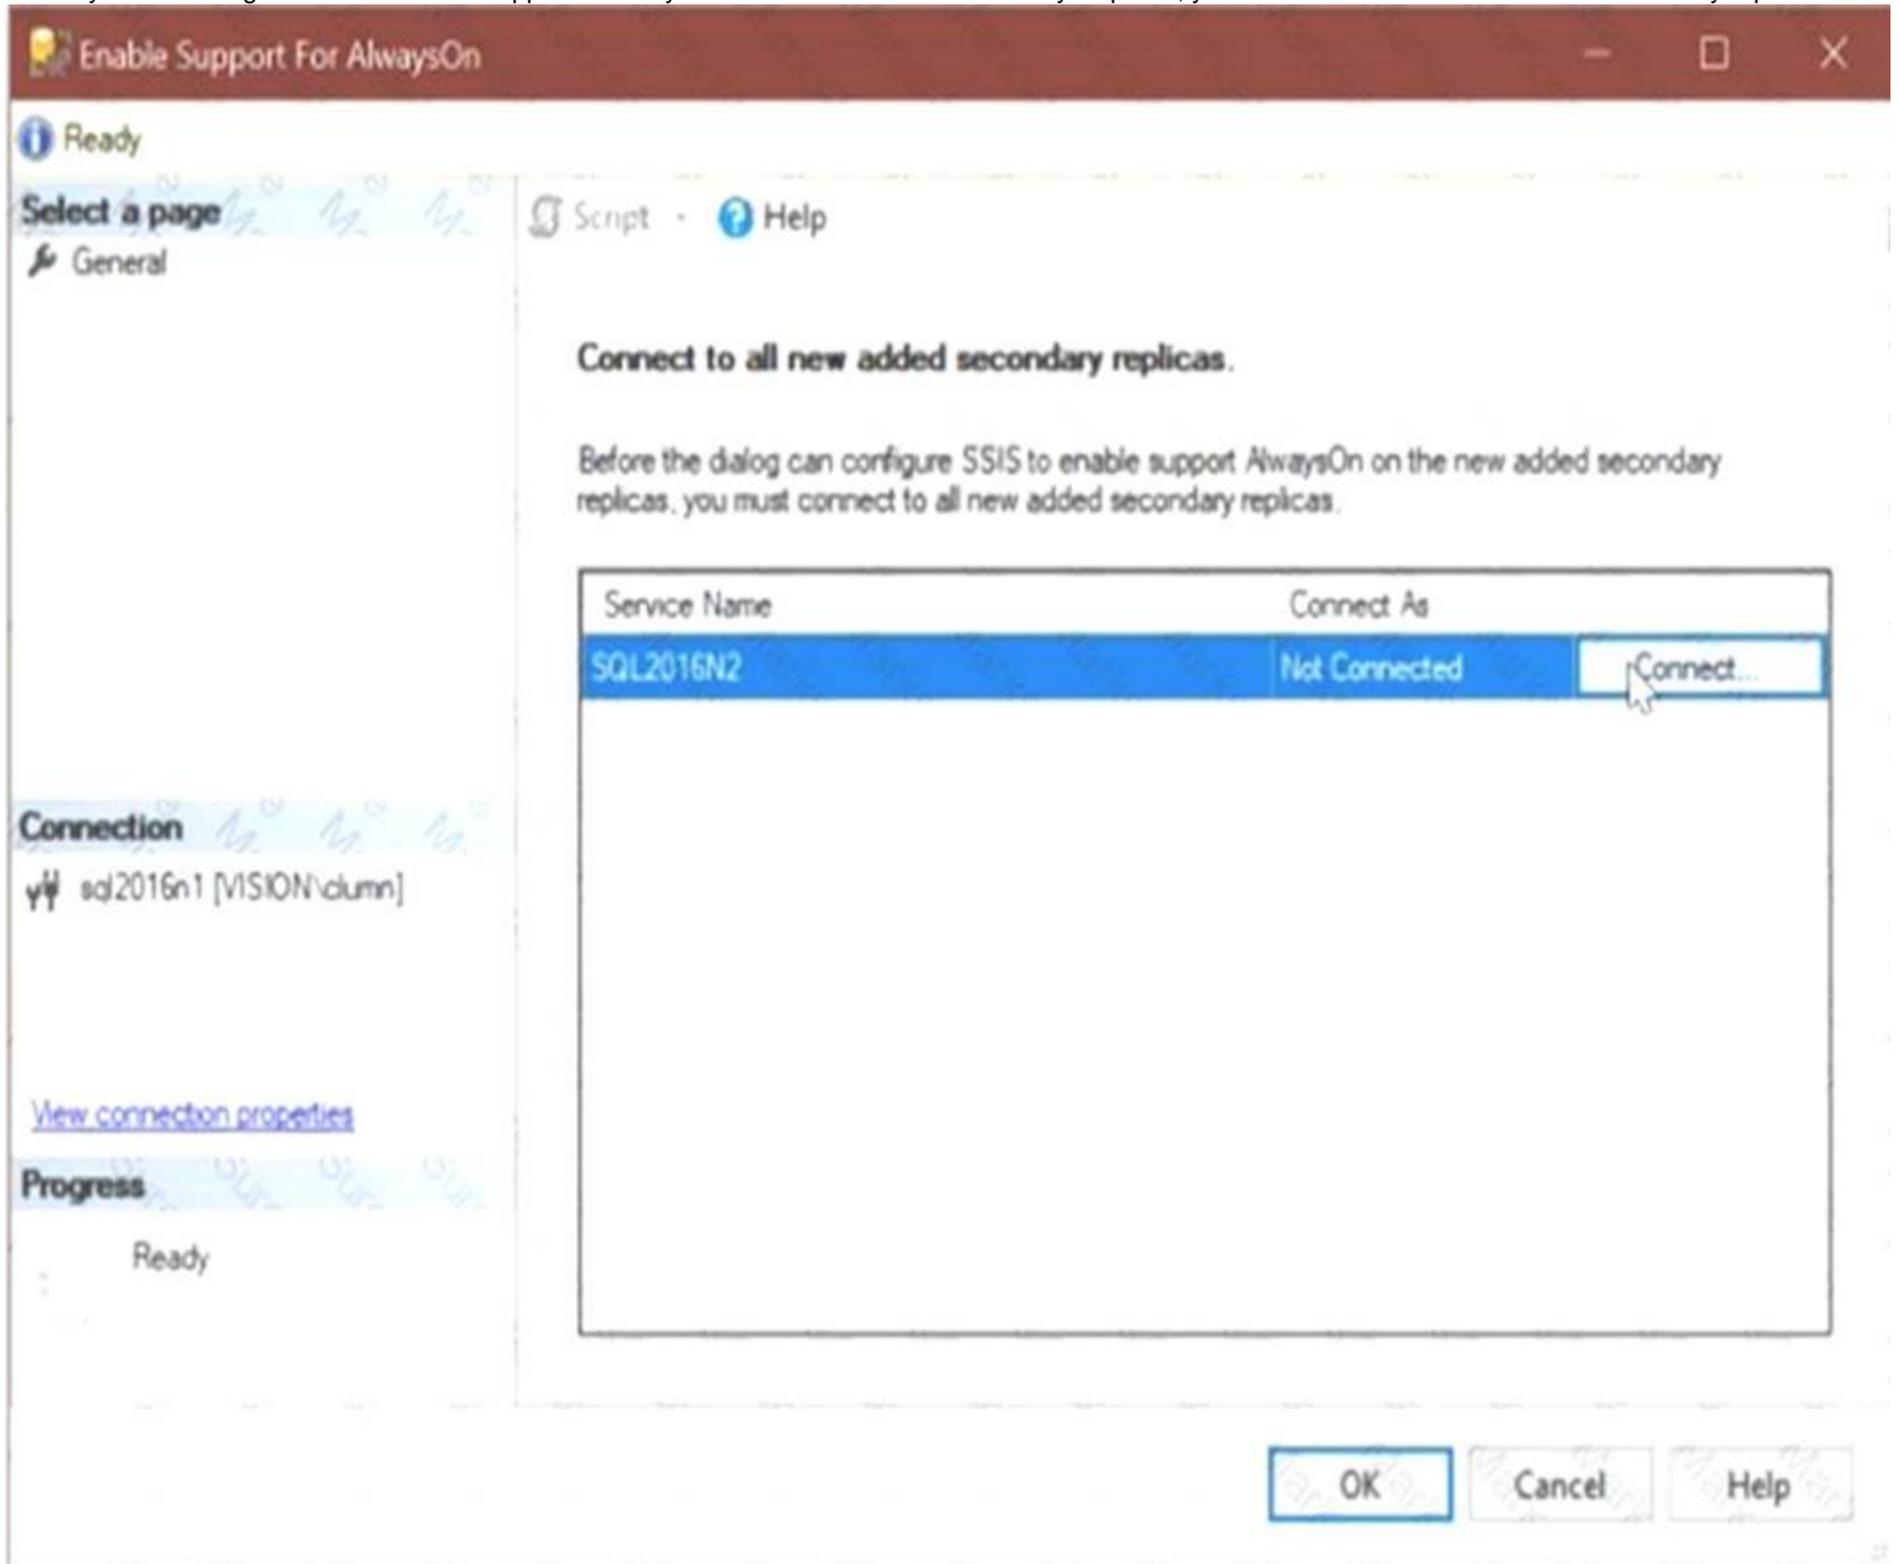
You plan to deploy several Microsoft SQL Server Integration Services (SSIS) packages to a highly available SQL Server instance. The instance is configured to use an AlwaysOn availability group that has two replicas. You need to identify which deployment method must be used to ensure that the packages are always accessible from all the nodes in the availability group. Which deployment method should you use for the packages?

- A. Deploy to the msdb database on the secondary replica.
- B. Deploy to the msdb database on the primary replica.
- C. Deploy to a file on the hard drive of the primary replica.
- D. Deploy to a shared folder on a file server.

Answer: A

Explanation:

Before you can configure SSIS to enable support of AlwaysOn on the new added secondary Replicas, you must connect to all new added secondary replicas.



Note: To use SSIS with AlwaysOn, you'll need to add the SSIS Catalog (SSISDB) into an Availability Group. You'll need to do the following steps:

- ▶ Make sure you meet the prerequisites for using AlwaysOn
- ▶ Connect to every node and create the SSISDB catalog. We need to create the catalog even on secondary nodes to create the other server-level objects (cleanup jobs, keys, accounts etc) that are used by SSIS.
- ▶ Delete the SSISDB databases on secondary nodes.
- ▶ Create an availability group, specifying SSISDB as the user database
- ▶ Specify secondary replicas.

References: <https://chrislamnah.com/2017/05/09/enabling-alwayson-for-ssisdb/>

NEW QUESTION 2

You are designing the data warehouse to import data from three different environments. The sources for the data warehouse will be loaded every hour. Scenario A includes tables in a Microsoft Azure SQL Database:

- ▶ Millions of updates and inserts occur per hour

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

Scenario B includes tables with status update changes:

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

Scenario C includes an external source database:

- ▶ Updates and inserts occur regularly.
- ▶ No changes to the database should require code changes to any reports or applications.
- ▶ Columns are added and dropped to tables in the database periodically. These schema changes should not require any interruption or reconfiguration of the change detection method 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:

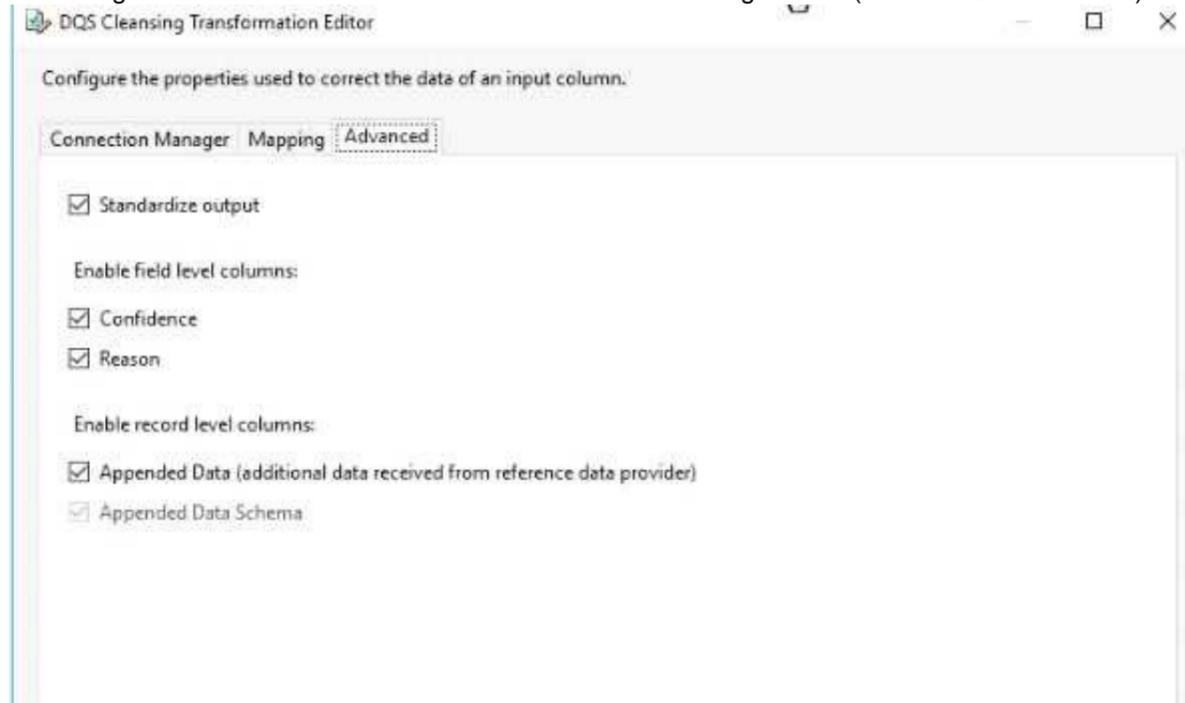
<https://docs.microsoft.com/en-us/sql/relational-databases/track-changes/track-data-changes-sql-server> <https://docs.microsoft.com/en-us/sql/relational-databases/tables/temporal-table-usage-scenarios>

NEW QUESTION 3

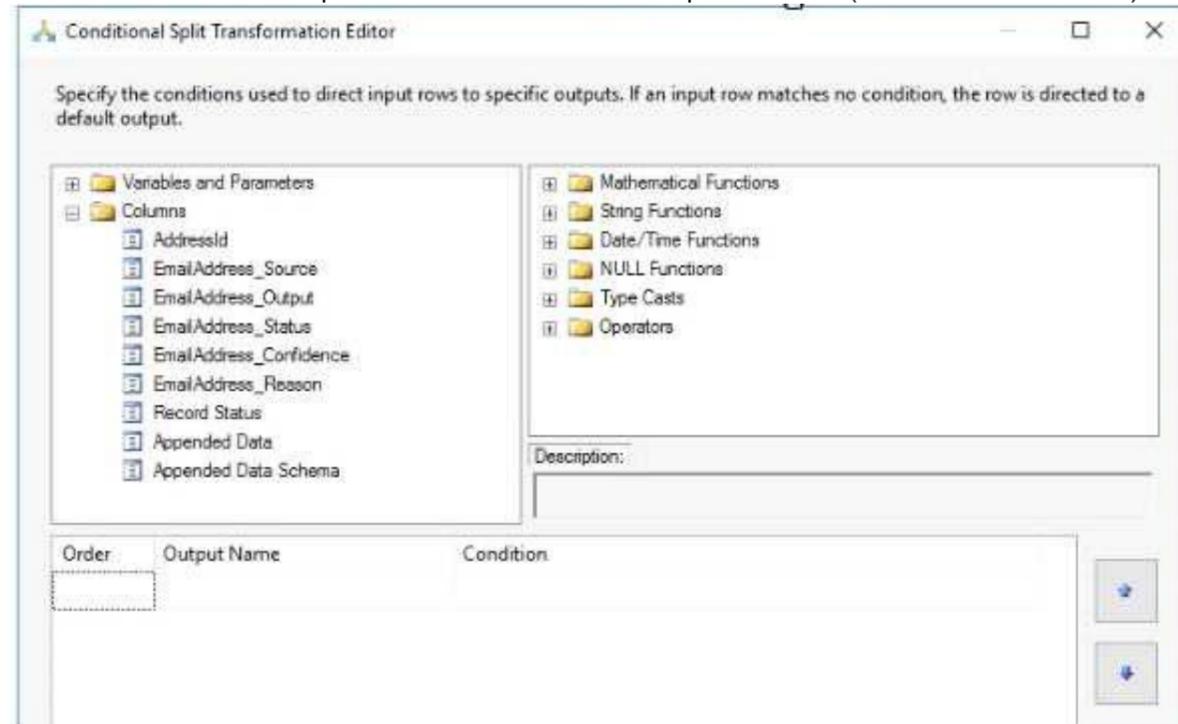
You have a Microsoft SQL Server Integration Services (SSIS) package that contains a Data Flow task as shown in the Data Flow exhibit. (Click the Exhibit button.)



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



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



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

Answer area

- | | | |
|---|-----------------------|-----------------------|
| | Yes | No |
| You can use the EmailAddress_Output column to split the output. | <input type="radio"/> | <input type="radio"/> |
| You can use the EmailAddress_Status column to split the output. | <input type="radio"/> | <input type="radio"/> |

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

You can use the EmailAddress_Output column to split the output.

Yes

No

You can use the EmailAddress_Status column to split the output.

NEW QUESTION 4

You are designing an indexing strategy for a data warehouse. The data warehouse contains a table named Table1. Data is bulk inserted into Table1. You plan to create the indexes configured as shown in the following table.

Index name	Indexing specifications
Index1	<ul style="list-style-type: none"> • Index1 contains all the data in Table1. • Queries against Index1 perform aggregation operations against hundreds of millions of rows.
Index2	<ul style="list-style-type: none"> • Index2 returns all the columns in this index. • Index2 contains 80 percent of the columns in Table1. • Index2 is used to assist with queries against other tables by performing point lookups against Table1.

Which type of index should you use to minimize the query times of each index? To answer, drag the appropriate index types to the correct indexes. Each index type may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

Index Types

Clustered	Clustered columnstore
Hash	Heap
Nonclustered	Nonclustered columnstore

Answer Area

Index1:

Index2:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Index Types

Clustered	Clustered columnstore
Hash	Heap
Nonclustered	Nonclustered columnstore

Answer Area

Index1:

Index2:

NEW QUESTION 5

Your company has a Microsoft SQL Server data warehouse instance. The human resources department assigns all employees a unique identifier. You plan to store this identifier in a new table named Employee.

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

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

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

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

Answer Area

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

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

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

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

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

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

NEW QUESTION 6

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 7

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 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 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 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 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 Totaisale 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.

Results Messages

	Transaction Key	Date Key	Customer Key	Bill To Customer Key	Supplier Key	Transaction Type Key	Payment Method Key	WWI Invoice ID
1	7	2013-01-01	375	202	0	1	0	7
2	11	2013-01-01	387	202	0	1	0	11
3	12	2013-01-01	330	202	0	1	0	12
4	13	2013-01-01	274	202	0	1	0	13
5	16	2013-01-01	215	202	0	1	0	16
6	25	2013-01-01	298	202	0	1	0	25
7	26	2013-01-01	285	202	0	1	0	26
8	30	2013-01-01	368	202	0	1	0	30
9	35	2013-01-01	232	202	0	1	0	35
10	39	2013-01-01	346	202	0	1	0	39
11	41	2013-01-01	216	202	0	1	0	41
12	63	2013-01-02	224	202	0	1	0	42
13	64	2013-01-02	264	202	0	1	0	43
14	65	2013-01-02	268	202	0	1	0	44
15	70	2013-01-02	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
 CREATE SET
 CREATE SUBCUBE

CURRENTCUBE.Measures.ProfitMargin

AS 'Measures. [TotalSale]
 [RawCost]

/Measures. [TotalSale]
 [RawCost]

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer area

CREATE MEMBER
 CREATE SET
 CREATE SUBCUBE

CURRENTCUBE.Measures.ProfitMargin

AS 'Measures. [TotalSale]
 [RawCost]

/Measures. [TotalSale]
 [RawCost]

NEW QUESTION 10

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

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 14

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 15

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 17

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

- A. dtutil.exe /FILE c:\repository\package1.dtsx /DestServer Server1 /COPY SQL;package1.dtsx
- B. dtutil.exe /I /FILE c:\repository\package1.dtsx
- C. dtutil.exe /SQL package1 /COPY OTS;c:\repository\package1.dtsx
- D. dtutil.exe /SQL package1 /DELETE

Answer: A

NEW QUESTION 20

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 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 the following line-of-business solutions:

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

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

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

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

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

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 25

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 27

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question. You are developing a Microsoft SQL Server Integration Services (SSIS) package. The package design consists of the sources shown in the following diagram:



Each source contains data that is not sorted. You need to combine data from all of the sources into a single dataset. Which SSIS Toolbox item should you use?

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

Answer: C

NEW QUESTION 32

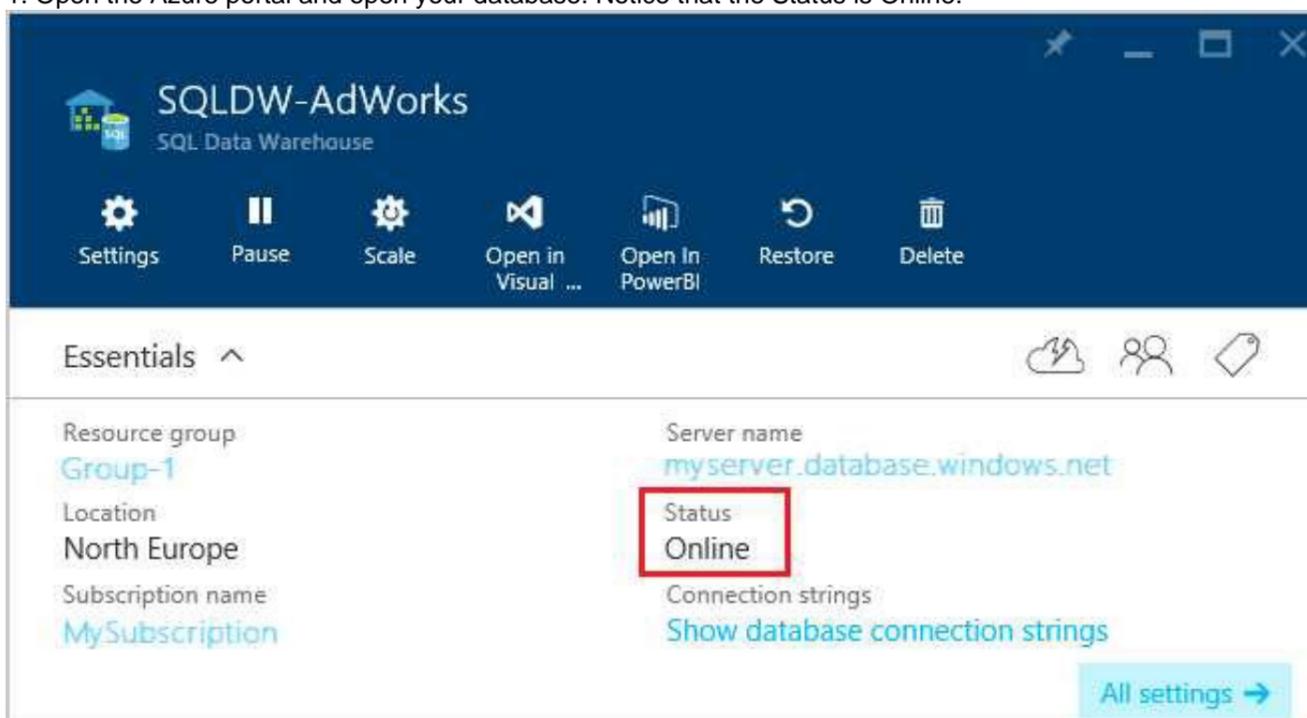
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this 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 36

You manage Master Data Services (MDS). You plan to create entities and attributes and load them with the data. You also plan to match data before loading it into Data Quality Services (DQS).

You need to recommend a solution to perform the actions.

What should you recommend?

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

Answer: A

Explanation:

In the Master Data Services Add-in for Excel, matching functionality is provided by Data Quality Services (DQS). This functionality must be enabled to be used.

- ▶ To enable Data Quality Services integration
- ▶ Open Master Data Services Configuration Manager.
- ▶ In the left pane, click Web Configuration.
- ▶ On the Web Configuration page, select the website and web application.
- ▶ In the Enable DQS Integration section, click Enable integration with Data Quality Services.
- ▶ On the confirmation dialog box, click OK.

References:

<https://docs.microsoft.com/en-us/sql/master-data-services/install-windows/enable-data-quality-services-integrati>

NEW QUESTION 40

You manage Master Data Services (MDS).

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

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

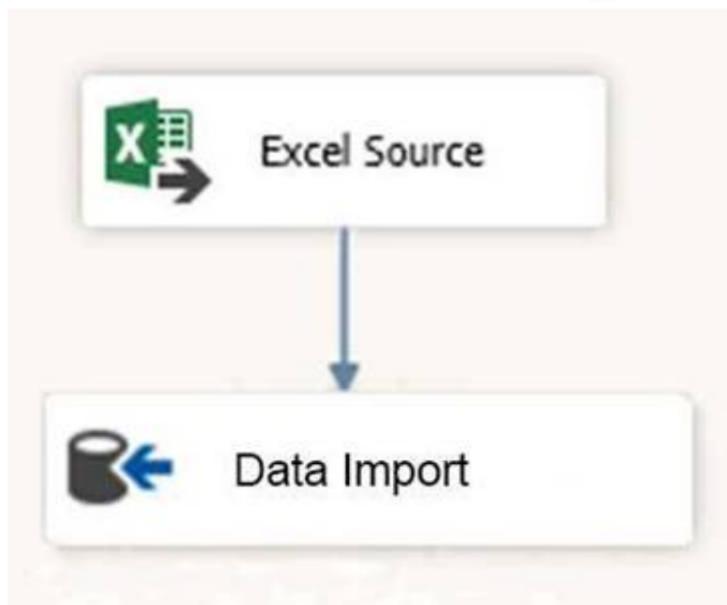
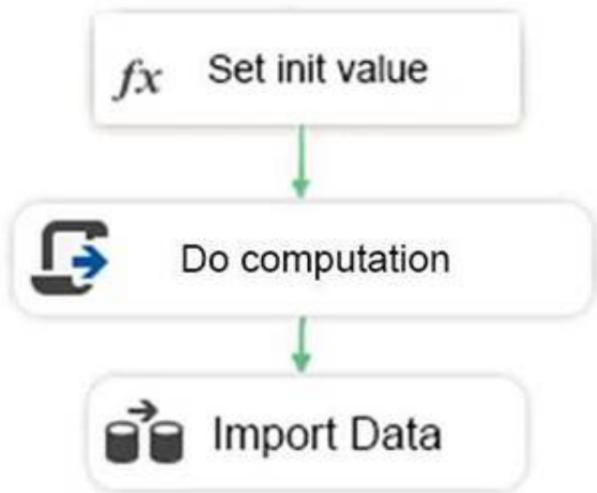
You need to configure the Transaction Log Type setting. Which type should you use?

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

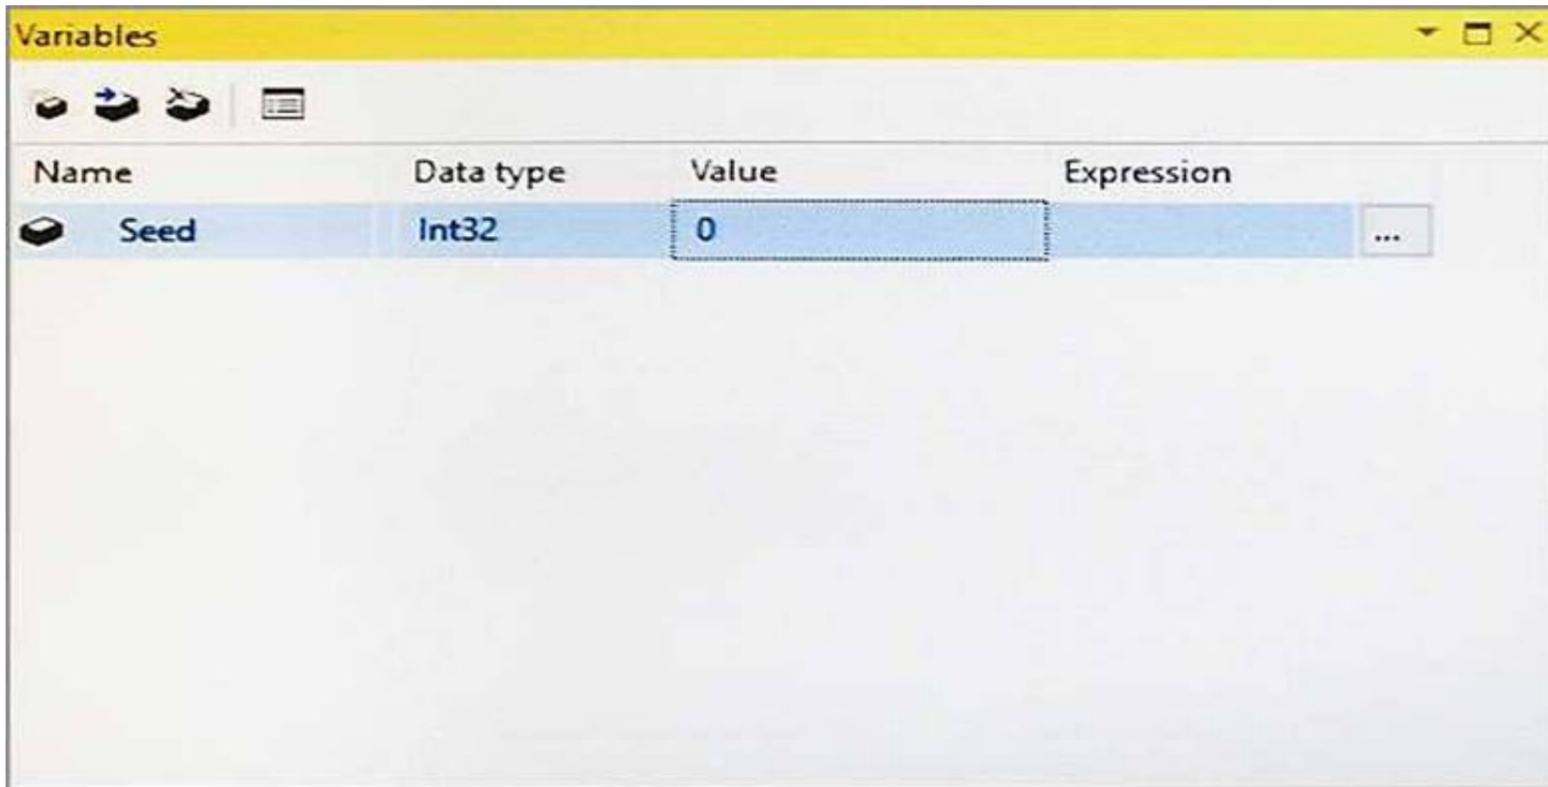
Answer: D

NEW QUESTION 45

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.



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 47

You are the administrator of a Microsoft SQL Server Master Data Services (MDS) model. The model was developed to provide consistent and validated snapshots of master data to the ETL processes by using subscription views. A new model version has been created.

You need to ensure that the ETL processes retrieve the latest snapshot of master data. What should you do?

- A. Add a version flag to the new version, and create new subscription views that use this version flag.
- B. Create new subscription views for the new version.
- C. Update the subscription views to use the new version.
- D. Update the subscription views to use the last committed version.

Answer: A

Explanation:

When a version is ready for users or for a subscribing system, you can set a flag to identify the version. You can move this flag from version to version as needed. Flags help users and subscribing systems identify which version of a model to use.

References: <https://docs.microsoft.com/en-us/sql/master-data-services/versions-master-data-services>

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 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 53

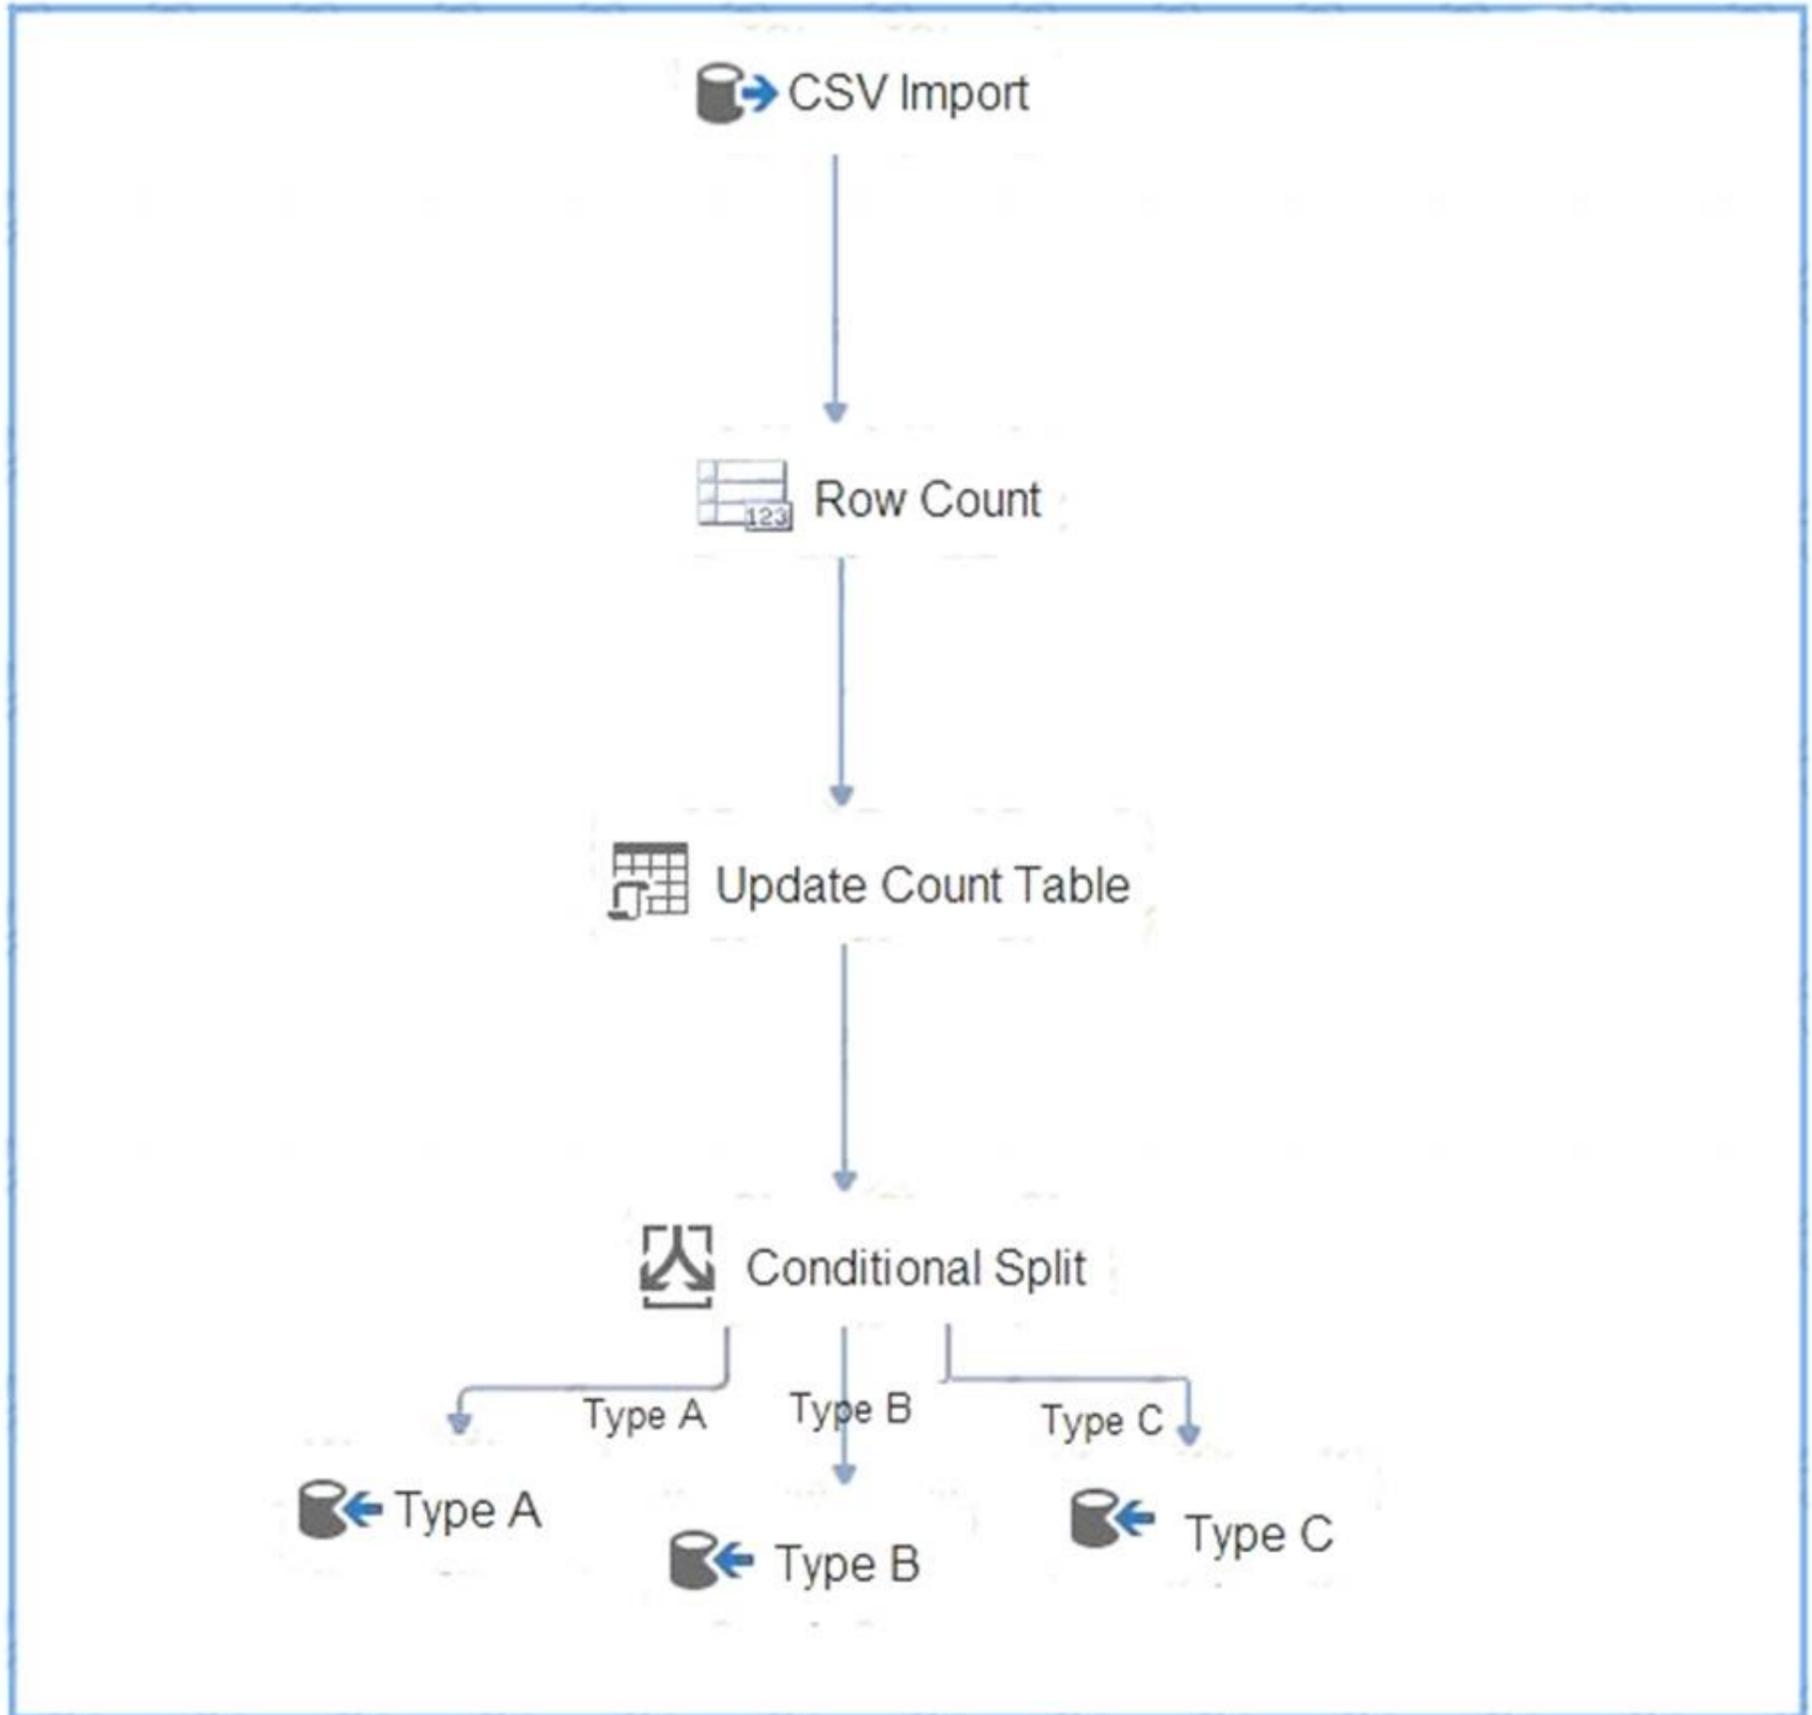
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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 56

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 60

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 64

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 68

You have a series of analytic data models and reports that provide insights into the participation rates for sports at different schools. Users enter information about sports and participants into a client application. The application stores this transactional data in a Microsoft SQL Server database. A SQL Server Integration Services (SSIS) package loads the data into the models.

When users enter data, they do not consistently apply the correct names for the sports. The following table shows examples of the data entry issues.

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

You need to create a new knowledge base to improve the quality of the sport name data.

How should you configure the knowledge base? To answer, select the appropriate options in the dialog box in the answer area.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Spot 1: Create Knowledge base from: None

Select None if you do not want to base the new knowledge base on an existing knowledge base or data file.

NEW QUESTION 71

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

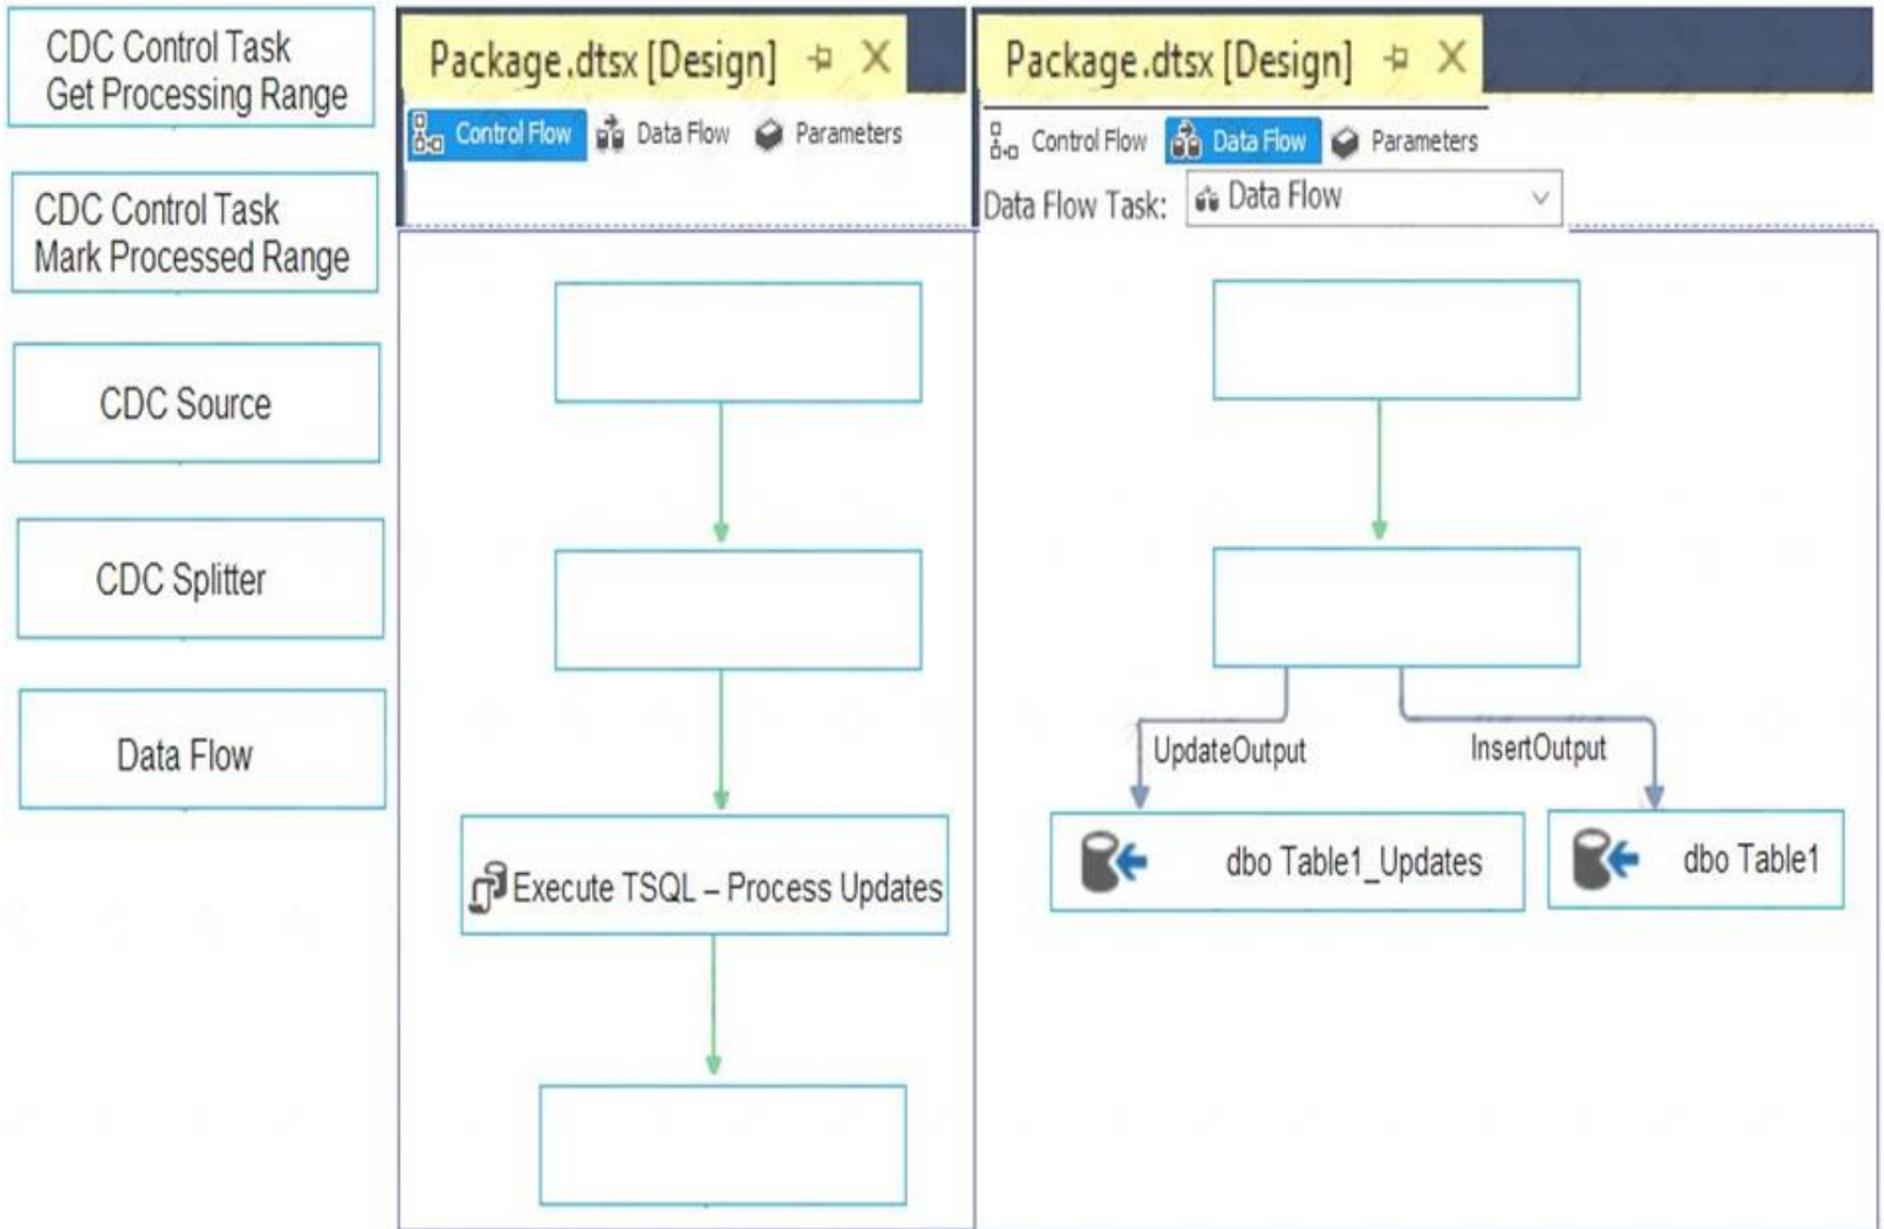
The SSIS package must load new records into Table1 and updated records into Table1_Updates. After loading records, the package must call a Transact-SQL statement to process updated rows according to existing business logic.

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 73

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 75

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

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

You need to configure the measures.

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

NOTE: Each correct selection is worth one point.

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:

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

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

References:

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

NEW QUESTION 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 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 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 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 82

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 configure the server to automatically delete the transaction logs nightly.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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 83

You have a data quality project that focuses on the Products catalog for the company. The data includes a product reference number.

The product reference should use the following format: Two letters followed by an asterisk and then four or five numbers. An example of a valid number is XX*55522. Any reference number that does not conform to the format must be rejected during the data cleansing.

You need to add a Data Quality Services (DQS) domain rule in the Products domain. Which rule should you use?

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

Answer: A

Explanation:

For a pattern matching rule:

Any letter (A...Z) can be used as a pattern for any letter; case insensitive Any digit (0...9) can be used as a pattern for any digit

Any special character, except a letter or a digit, can be used as a pattern for itself Brackets, [], define optional matching

Example: ABC:0000

This rule implies that the data will contain three parts: any three letters followed by a colon (:), which is again followed by any four digits.

NEW QUESTION 88

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 93

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.

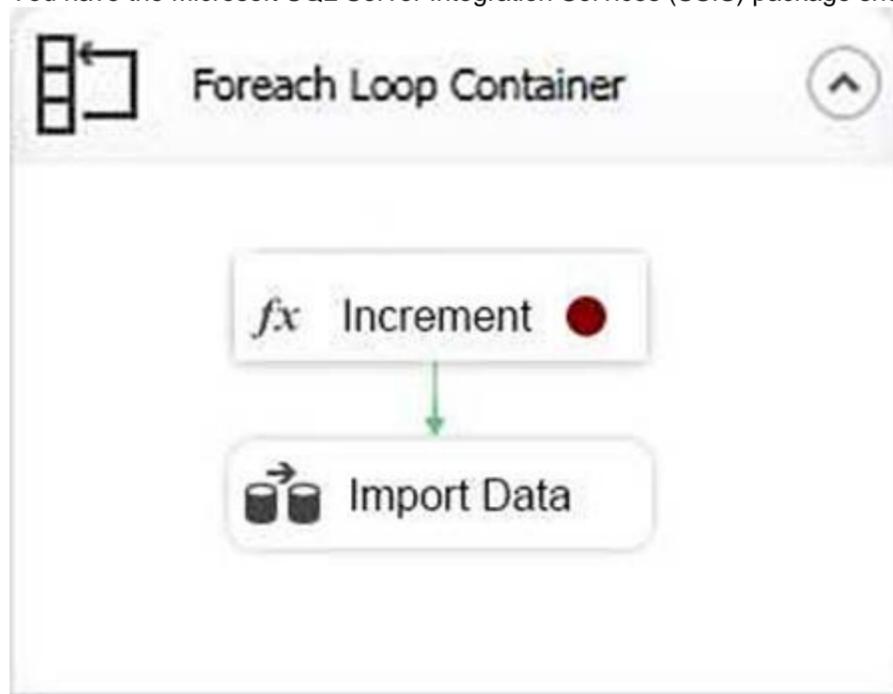
Answer: B

Explanation:

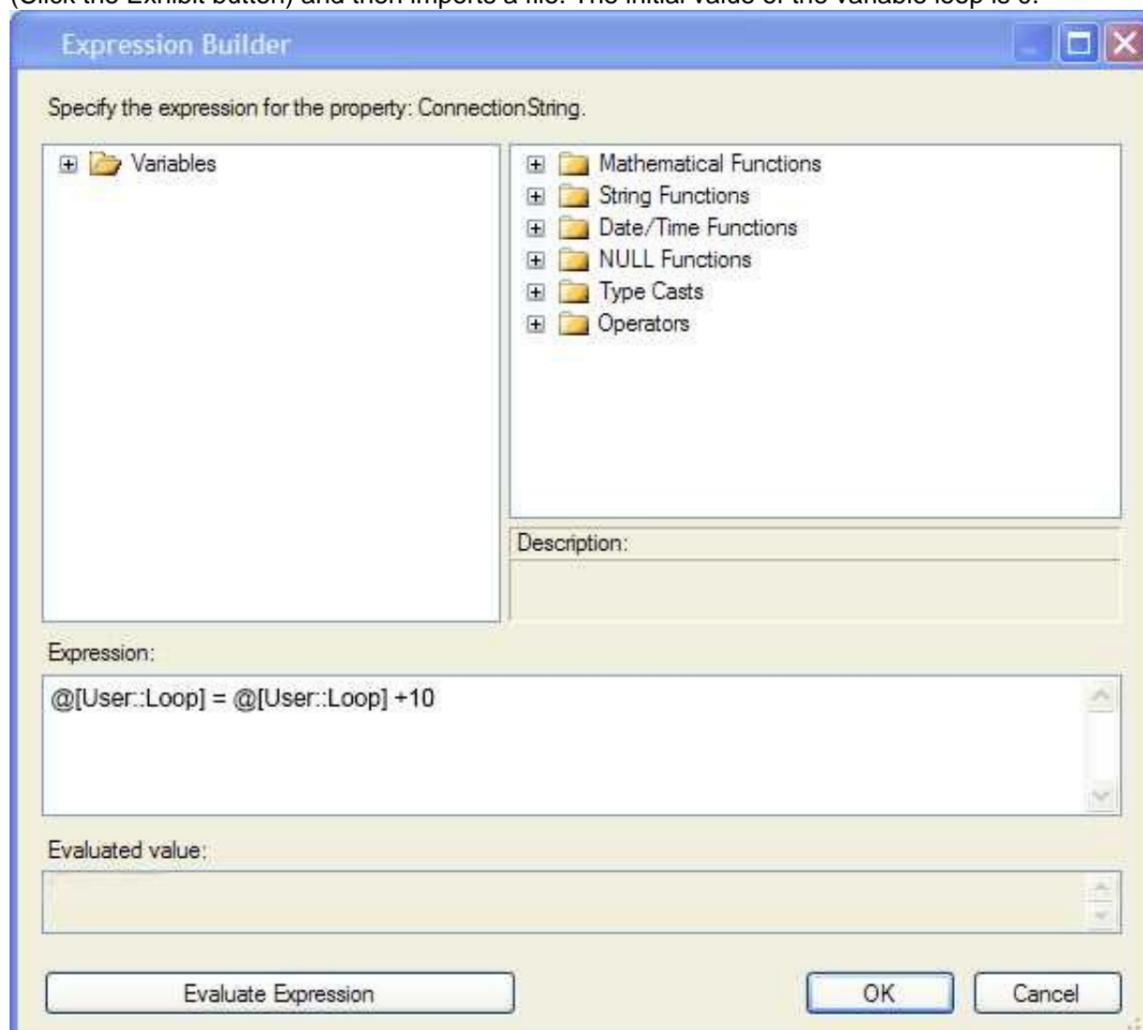
The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files 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 97

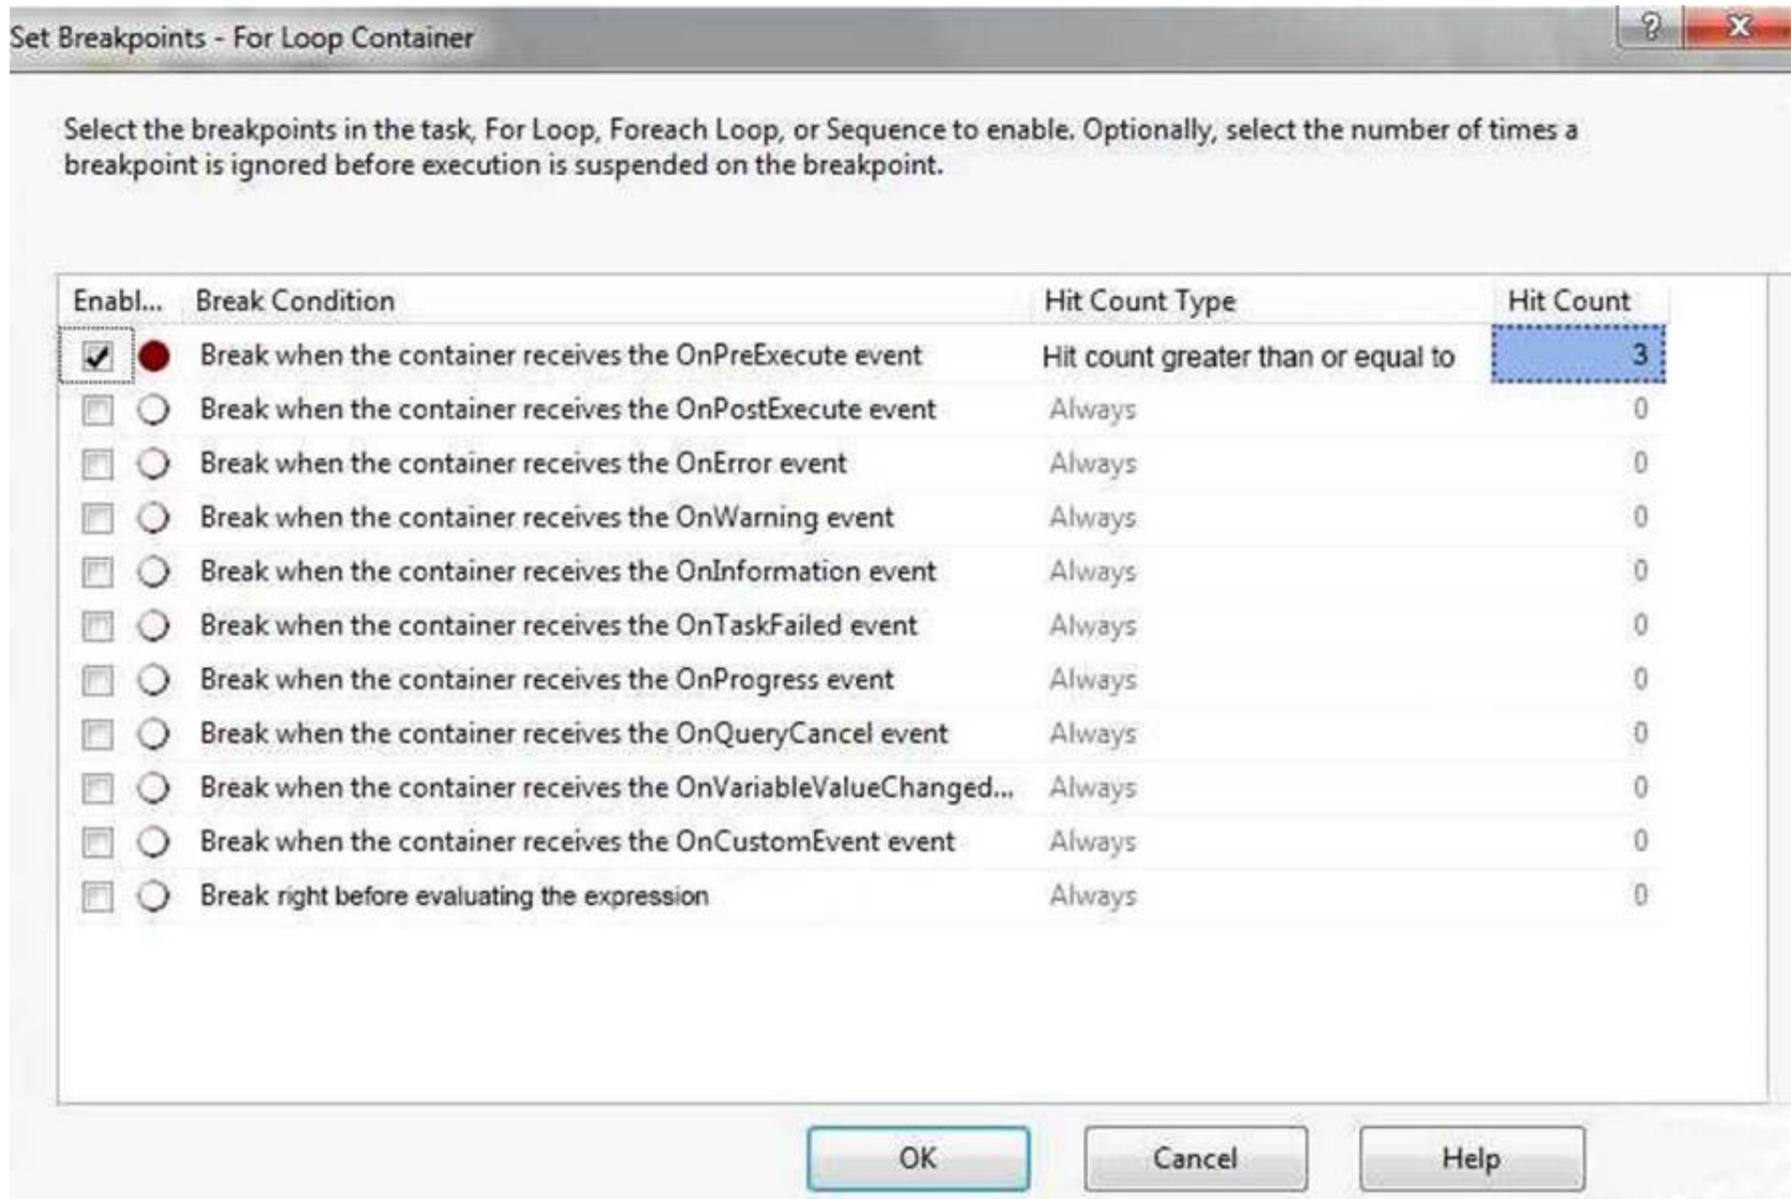
You have the Microsoft SQL Server Integration Services (SSIS) package shown in the Control flow exhibit. (Click the Exhibit button.)



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 101

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 106

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.

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

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.

End of repeated scenario

You need to perform the first step to partition the fact .Transaction table.

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

Answer area

CREATE PARTITION FUNCTION	[DateRange] (DATETIME) AS RANGE RIGHT
CREATE PARTITION SCHEME	[DateRange] (DATETIME) AS RANGE LEFT
CREATE PROCEDURE	[DateRange] (INT) AS RANGE RIGHT
CREATE RESOURCE POOL	[DateRange] (INT) AS RANGE LEFT

20160101', '20170101')

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

CREATE PROCEDURE
 [DateRange] (INT) AS RANGE LEFT

NEW QUESTION 111

You are developing a Microsoft SQL Server Data Warehouse. You use SQL Server Integration Services (SSIS) packages to import files from a Microsoft Azure blob storage to the data warehouse.

You plan to use multiple SQL Server instances and SSIS Scale Out to complete the workload faster. You must configure three SQL Server instances to run the SSIS package.

Which two actions should you perform? Each correct answer presents part of the solution. NOTE: Each correct selection is worth one point.

- A. Install The SSIS Scale Out Worker feature on two server
- B. Install the Scale Out Master role feature on one server.
- C. Deploy the SSIS project to the SSIS catalog only on the SQL Server which has the Scale Out Master role installed.
- D. Install the SSIS Scale Out Worker feature on all three server
- E. Install the Scale Out Master role on one server.
- F. Deploy the SSIS project to the SSIS catalog on all three SQL Servers in the SSIS Scale Out environment.

Answer: AD

NEW QUESTION 115

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

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

Answer: A

NEW QUESTION 119

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 122

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 123

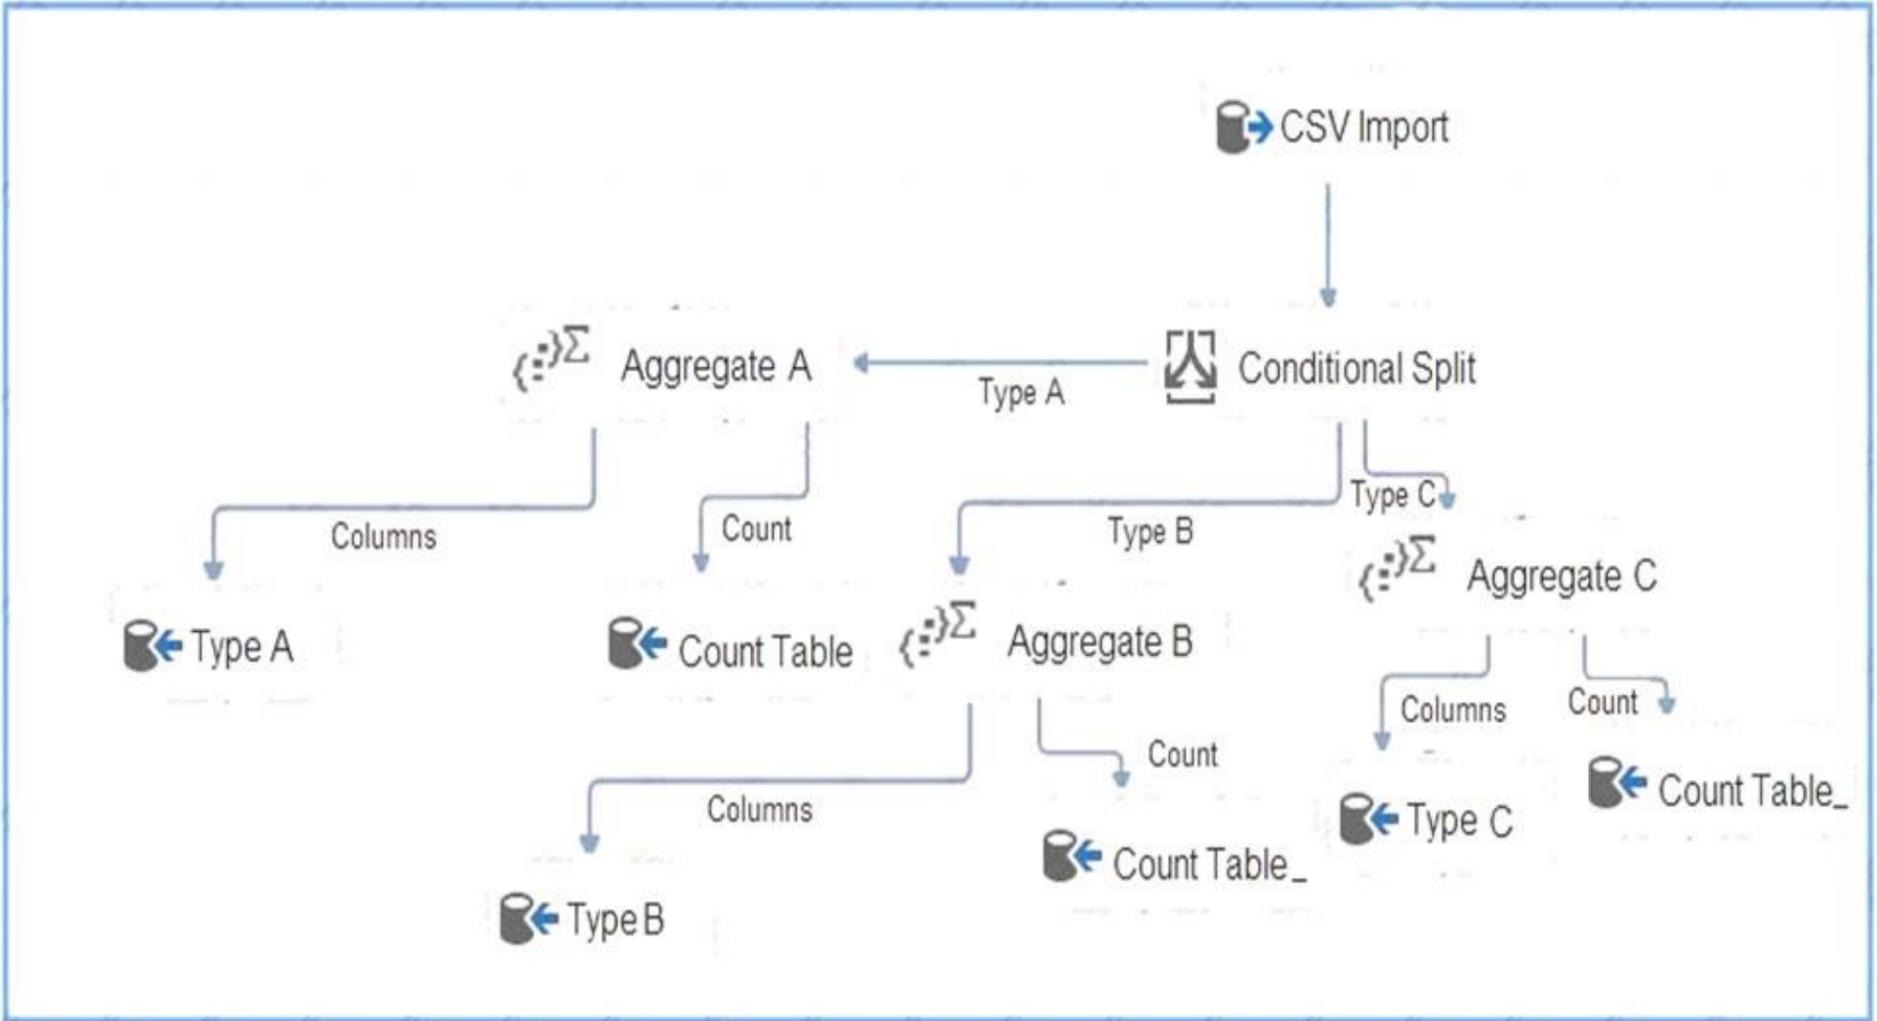
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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 127

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

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

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

You remove the Integration Services job.

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

Which stored procedure should you execute?

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

Answer: F

NEW QUESTION 131

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 134

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
 After you answer a question in this 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 137

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

Start of repeated scenario

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

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

The data warehouse includes the following tables:

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

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

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

You must implement a partitioning scheme for the fact.Transaction table to move older data to less expensive storage. Each partition will store data for a single calendar year, as shown in the exhibit (Click the Exhibit button.) You must align the partitions.

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

End of repeated scenario

You need to configure the fact. Transaction table.

Which three Transact-SQL segments should you use to develop the solution? To answer, move the appropriate Transact-SQL segments from the list of Transact-SQL segments to the answer area and arrange them in the correct order.

Results Messages

	Transaction Key	Date Key	Customer Key	Bill To Customer Key	Supplier Key	Transaction Type Key	Payment Method Key	WWI Invoice ID
1	7	2013-01-01	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

Transact-SQL segments

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

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

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

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

Answer area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

```

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

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

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

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

Answer area

```

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




```

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




```

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

NEW QUESTION 139

You administer a Microsoft SQL Server Master Data Services (MDS) model. All model entity members have passed validation. The current model version should be committed to form a record of master data that can be audited and create a new version to allow the ongoing management of the master data. You lock the current version. You need to manage the model versions. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area, and arrange them in the correct order.

Actions

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

Answer Area







B. Not Mastered

Answer: A

Explanation:

Box 1: Validate the current version.

In Master Data Services, validate a version to apply business rules to all members in the model version. You can validate a version after it has been locked.

Box 2: Commit the current version.

In Master Data Services, commit a version of a model to prevent changes to the model's members and their attributes. Committed versions cannot be unlocked.

Prerequisites:

Box 3: Create a copy of the current version.

In Master Data Services, copy a version of the model to create a new version of it. Note:

References:

NEW QUESTION 143

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 144

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 147

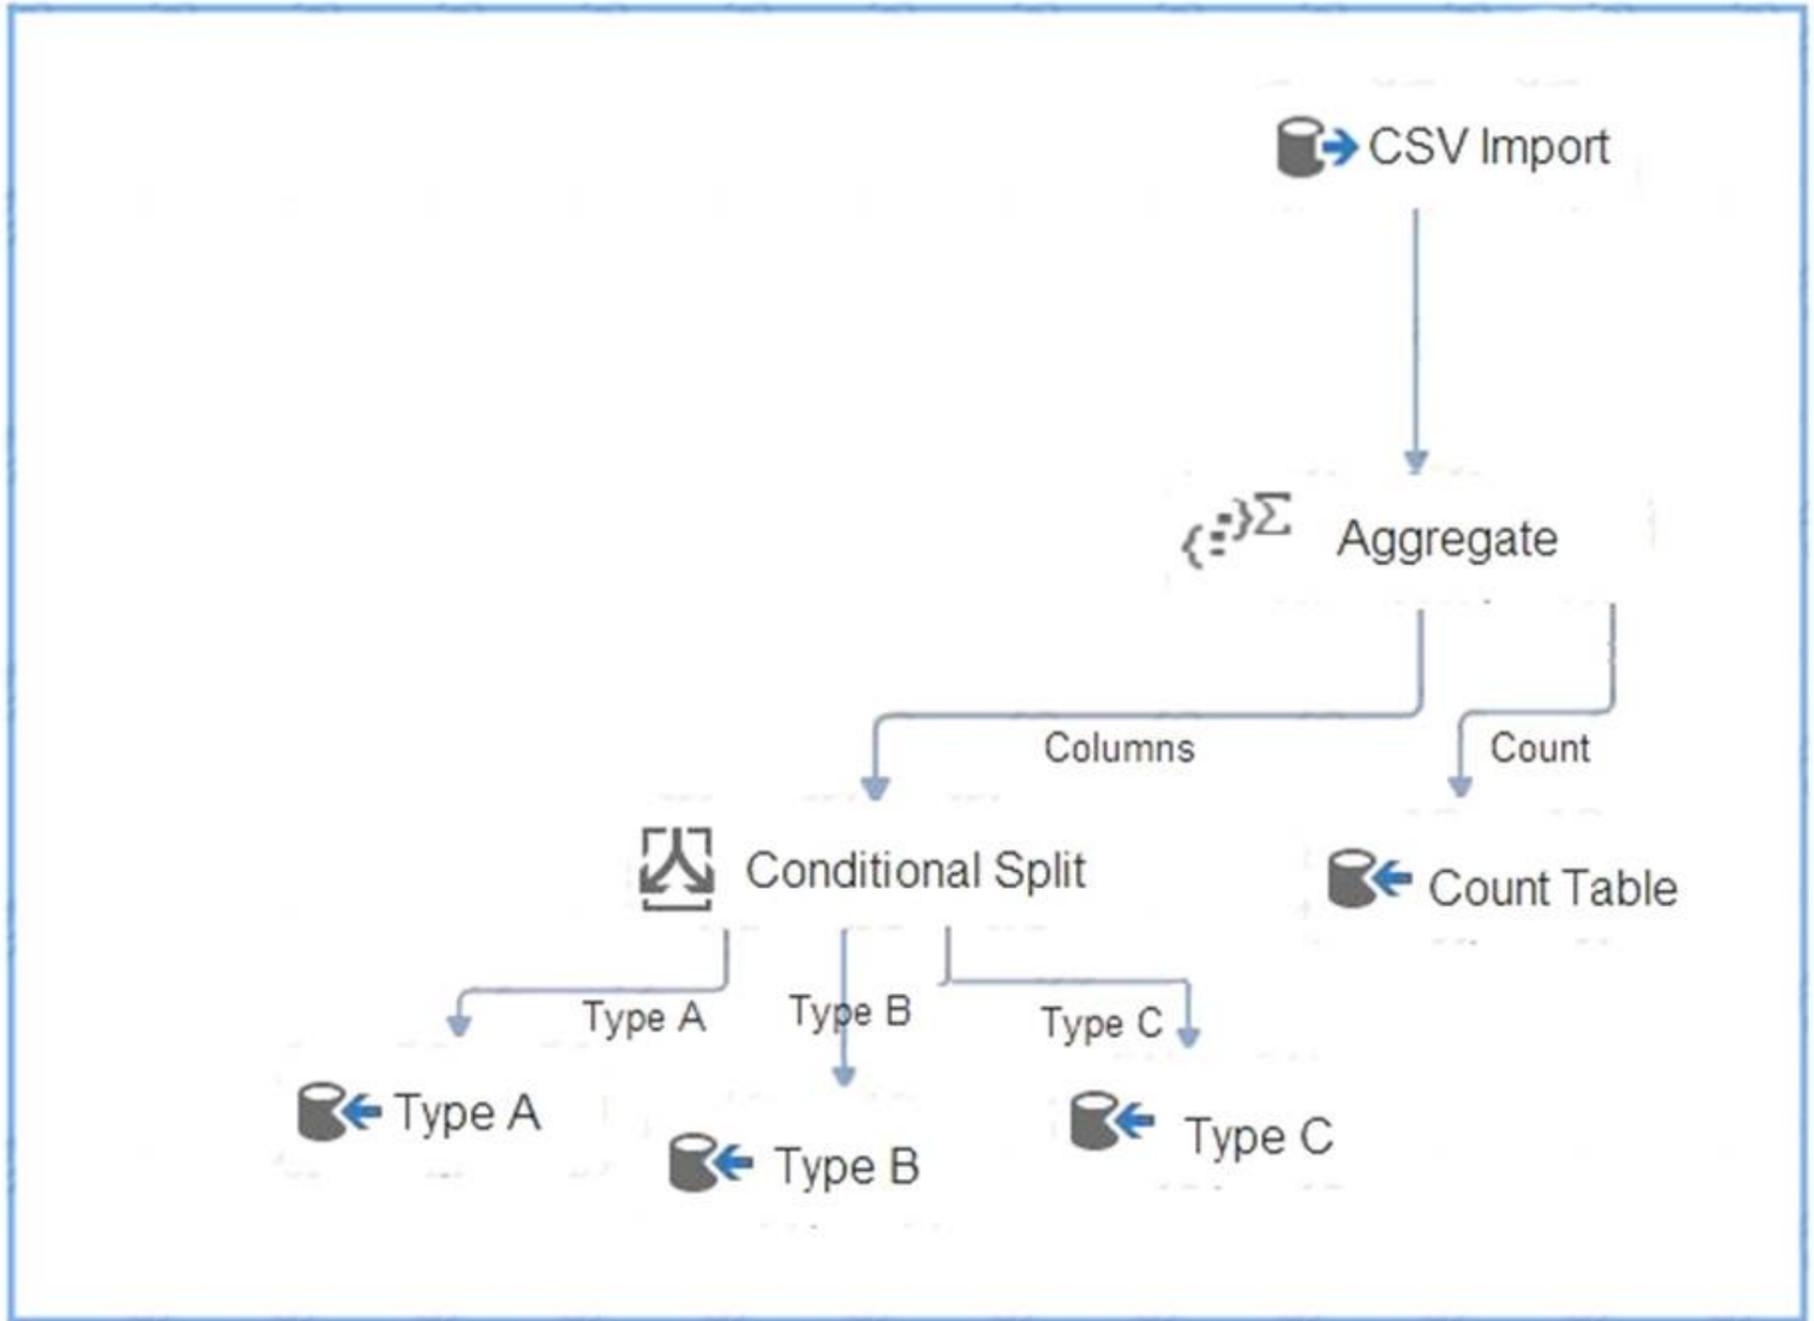
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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 149

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 151

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

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

You are the administrator of a Microsoft SOL Server Master Data Services (MDS) instance. The instance contains a model named Geography and a model named customer. The Geography model contains an entity named countryRegion.

You need to ensure that the countryRegion entity members are available in the customer model. Solution: In the Geography model, publish a business rule with a

Change Value action.
 Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 155

You are developing a Microsoft SQL Server Integration Services (SSIS) package. You enable the SSIS log provider for the Windows event log. You configure the package to use the ScriptTaskLogEntry event. You create a custom Script task.

You need to ensure that when the script completes, it writes the execution status to the event log on the server that hosts SSIS.

Which code segment should you add to the Script task?

- A. `Dts.TaskResult = (int)ScriptResults.Failure`
- B. `Dts.Events.FireWarning(0, "SSIS", "Script executed with return result " + Dts.TaskResult, String.Empty, 0)`
- C. `System.Diagnostics.EventLog.writeEntryC("SSIS", "Script executed with return result " + Dts.TaskResult, System.Diagnostics.EventLogEntryType.Information)`
- D. `Dts.TaskResult = (int)ScriptResults.Success`

Answer: D

NEW QUESTION 156

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

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

You need to cleanse a data flow source by removing duplicate records based on approximate matches. Which SSIS Toolbox item should you use?

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

Answer: E

Explanation:

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

NEW QUESTION 159

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

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

References:

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

NEW QUESTION 160

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

70-767 Practice Exam Features:

- * 70-767 Questions and Answers Updated Frequently
- * 70-767 Practice Questions Verified by Expert Senior Certified Staff
- * 70-767 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 70-767 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

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