

70-767.68q

Number: 70-767
Passing Score: 800
Time Limit: 120 min

70-767



Website: <https://vceplus.com>
VCE to PDF Converter: <https://vceplus.com/vce-to-pdf/>
Facebook: <https://www.facebook.com/VCE.For.All.VN/>
Twitter : https://twitter.com/VCE_Plus

<https://vceplus.com/>

Implementing a SQL Data Warehouse

Exam A

QUESTION 1

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

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

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

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

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

You need to optimize performance.

Solution: You run the following Transact-SQL statement:

```
CREATE STATISTICS Stat1
On Sales.SalesOrderHeader (OrderDate)
WITH SAMPLE 0 PERCENT
```



Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Microsoft recommend against specifying 0 PERCENT or 0 ROWS in a CREATE STATISTICS..WITH SAMPLE statement. When 0 PERCENT or ROWS is specified, the statistics object is created but does not contain statistics data.

References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/create-statistics-transact-sql>

QUESTION 2

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 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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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/>

QUESTION 3

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

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

You 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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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/>

QUESTION 4

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 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 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.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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/>

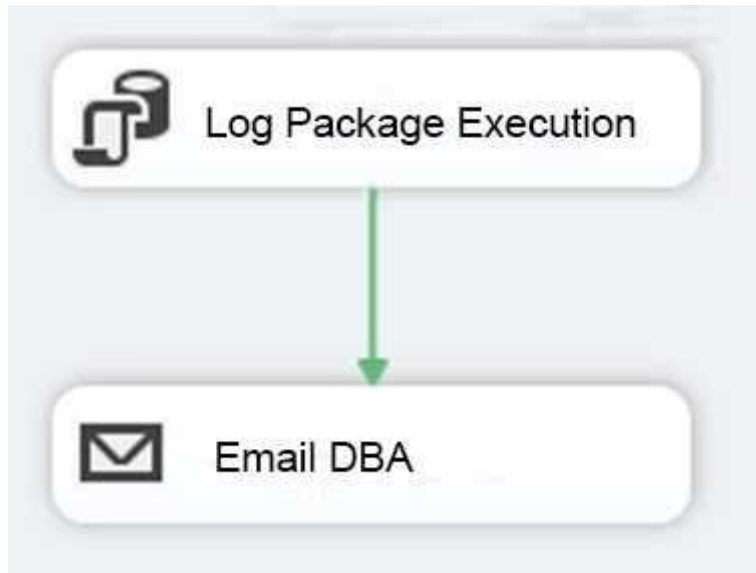
QUESTION 5

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

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

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

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



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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

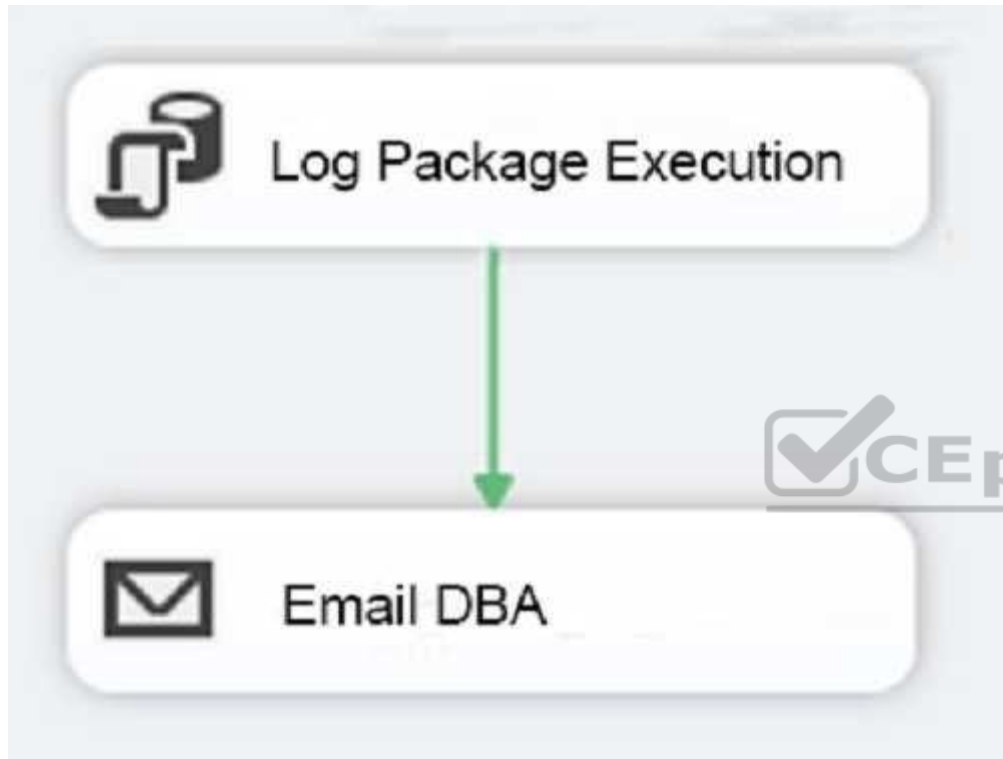
QUESTION 6

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

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

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

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



Solution: You add the control flow to a control flow package part. You add an instance of the control flow package part to each data warehouse load package.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

QUESTION 7

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 to daily. The Fact.Order table is loaded by using an ETL process. Indexes have been added to the table over time, but the presence of these indexes slows data loading.

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

You have the following requirements:

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

You are not permitted to make changes to the client applications.

You need to 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.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

We create one partition for each day, which means that a granularity of one day is used. 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.

The Fact.Order table is optimized for weekly reporting, but the company wants to change it to daily.

Maximize the performance during the data loading process for the Fact.Order partition.

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

QUESTION 8

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer choices, but the text of the scenario is exactly the same in each question in 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 to daily. The Fact.Order table is loaded by using an ETL process. Indexes have been added to the table over time, but the presence of these indexes slows data loading.

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

You have the following requirements:

- Implement table partitioning to improve the manageability of the data warehouse and to avoid the need to repopulate all transactional data each night. Use a partitioning strategy that is as granular as possible.
- Partition the Fact.Order table and retain a total of seven years of data.

- Partition the Fact.Ticket table and retain seven years of data. At the end of each month, the partition structure must apply a sliding window strategy to ensure that a new partition is available for the upcoming month, and that the oldest month of data is archived and removed.
- Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.
- Incrementally load all tables in the database and ensure that all incremental changes are processed.
- Maximize the performance during the data loading process for the Fact.Order partition.
- Ensure that historical data remains online and available for querying.
- Reduce ongoing storage costs while maintaining query performance for current data.

You are not permitted to make changes to the client applications.

You need to optimize 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.
- E. Implement row compression for the **Fact.Order** table.
- F. Move the index for the **Fact.Order** table to lower-cost storage.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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: The data warehouse has grown and the cost of storage has increased. 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-data-warehouse/>

QUESTION 9

HOTSPOT

You deploy a Microsoft Azure SQL Data Warehouse instance. The instance must be available eight hours each day.

You need to pause Azure resources when they are not in use to reduce costs.

What will be the impact of pausing resources? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

What will happen to existing queries that are running?

	▼
The data warehouse instance pauses when existing queries have completed. No new queries are permitted.	
The existing queries will be immediately terminated.	
The existing queries will be paused until the data warehouse instance is resumed.	

What will happen to the charges for the data warehouse instance?

	▼
You will stop being charged for compute resources but will continue to be charged for storage.	
You will continue to be charged for both compute resources and storage.	
You are no longer charged for storage but continue to pay for the assigned data warehouse instance units.	

Correct Answer:

Answer Area

What will happen to existing queries that are running?

<input type="radio"/> The data warehouse instance pauses when existing queries have completed. No new queries are permitted.
<input checked="" type="radio"/> The existing queries will be immediately terminated.
<input type="radio"/> The existing queries will be paused until the data warehouse instance is resumed.

What will happen to the charges for the data warehouse instance?

<input type="radio"/> You will stop being charged for compute resources but will continue to be charged for storage.
<input checked="" type="radio"/> You will continue to be charged for both compute resources and storage.
<input type="radio"/> You are no longer charged for storage but continue to pay for the assigned data warehouse instance until:

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To save costs, you can pause and resume compute resources on-demand. For example, if you won't be using the database during the night and on weekends, you can pause it during those times, and resume it during the day. You won't be charged for DWUs while the database is paused.

When you pause a database:

Compute and memory resources are returned to the pool of available resources in the data center

Data Warehouse Unit (DWU) costs are zero for the duration of the pause.

Data storage is not affected and your data stays intact.

SQL Data Warehouse cancels all running or queued operations.

When you resume a database:

SQL Data Warehouse acquires compute and memory resources for your DWU setting.

Compute charges for your DWUs resume.

Your data will be available.

You will need to restart your workload queries.

References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-compute-rest-api>

QUESTION 10

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

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

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

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

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

You need to load the data for the OnlineOrder table.

What should you use?

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

Correct Answer: H

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

QUESTION 11

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, you publish a Microsoft Excel workbook that contains a list of product names and current prices to an external website. Suppliers update pricing information in the workbook. Each supplier saves the workbook with a unique name.

Each night, the Products table is deleted and refreshed from MDS by using a Microsoft SQL Server Integration Services (SSIS) package. All files must be loaded in sequence.

You need to add a data flow in an SSIS package to perform the Excel files import in the data warehouse.

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

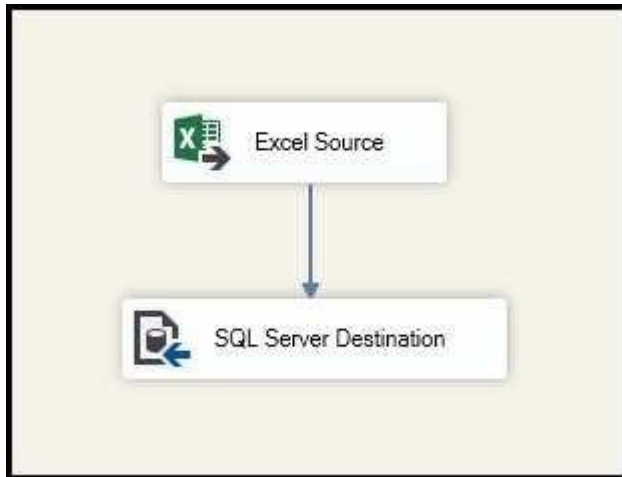
H. Foreach Loop container

Correct Answer: A **Section:**
(none) **Explanation**

Explanation/Reference:

Explanation:

If you're familiar with SSIS and don't want to run the SQL Server Import and Export Wizard, create an SSIS package that uses the Excel Source and the SQL Server Destination in the data flow.



References: <https://docs.microsoft.com/en-us/sql/integration-services/import-export-data/import-data-from-excel-to-sql>

QUESTION 12

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

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

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

Each week, you import a product catalog from a partner company to a staging table in DB2.

You need to create a stored procedure that will update the staging table by inserting new products and deleting discontinued products.

What should you use?

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



Correct Answer: G

Section: (none)

Explanation

Explanation/Reference:

Explanation:

QUESTION 13

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



Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 14

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

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

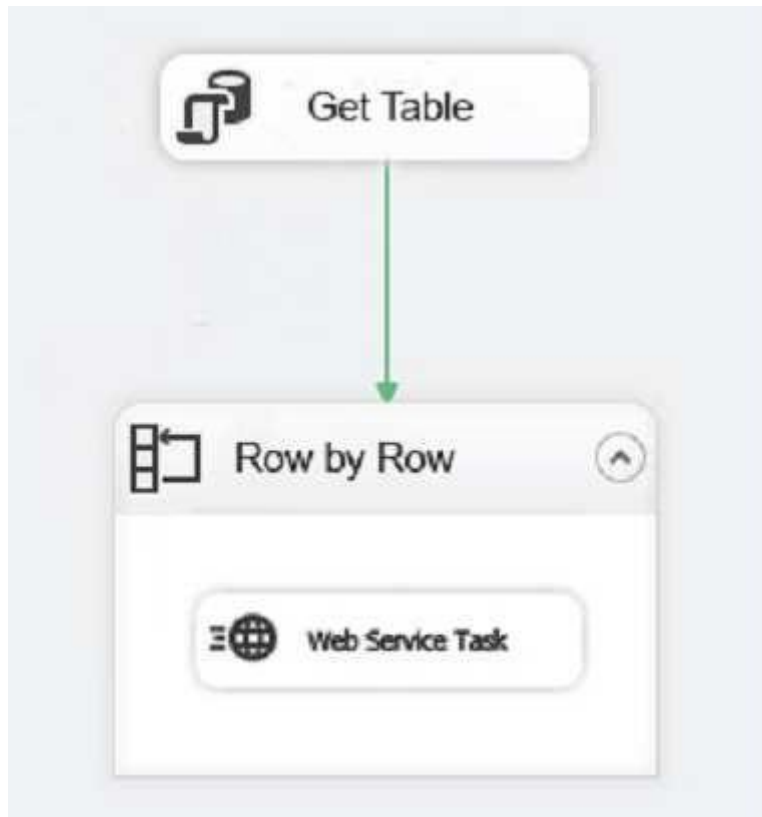
Explanation:

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>

QUESTION 15

You have a Microsoft SQL Server Integration Services (SSIS) package that includes the control flow shown in the following diagram.



You need to choose the enumerator for the Foreach Loop container.

Which enumerator should you use?

- A. Foreach SMO Enumerator
- B. Foreach ADO.Net Schema Rowset Enumerator

- C. Foreach NodeList Enumerator
- D. Foreach ADO Enumerator
- E. Foreach HDS File Enumerator
- F. Foreach File Enumerator

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Use the Foreach ADO enumerator to enumerate rows in tables. For example, you can get the rows in an ADO recordset.

References: <https://docs.microsoft.com/en-us/sql/integration-services/control-flow/foreach-loop-container?view=sql-server-2017>

QUESTION 16

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]?

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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.

QUESTION 17

DRAG DROP

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

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

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

You need to improve the quality of the data.

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

Select and Place:

Actions

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

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

Create a Data Quality Service (DQS) knowledge base.

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

Import the raw data from the users to perform discovery.

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

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

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

Answer Area



Correct Answer:

Actions

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

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

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

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

Answer Area

Import the raw data from the users to perform discovery.

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

⏪ Create a Data Quality Service (DQS) knowledge base.

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



Section: (none)

Explanation

Explanation/Reference:

Explanation:

References: <https://docs.microsoft.com/en-us/sql/data-quality-services/perform-knowledge-discovery> **QUESTION 18**

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.



<https://vceplus.com/>

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 19

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

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

You plan to deploy a Microsoft SQL server that will host a data warehouse named **DB1**.

The server will contain four SATA drives configured as a RAID 10 array.

You need to minimize write contention on the transaction log when data is being loaded to the database.

Solution: You replace the SATA disks with SSD disks.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

References: <https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on-separate-drives?view=sql-server-2017>

QUESTION 20

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

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

You plan to deploy a Microsoft SQL server that will host a data warehouse named **DB1**.

The server will contain four SATA drives configured as a RAID 10 array.

You need to minimize write contention on the transaction log when data is being loaded to the database.

Solution: You configure the server to automatically delete the transaction logs nightly.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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-separate-drives?view=sql-server-2017>

QUESTION 21

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

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

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

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

The query fails to return results.

You need to determine why the query fails.

Solution: You run the following Transact-SQL statement:

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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

To use submit_time we must use sys.dm_pdw_exec_requests table.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec-requests-transact-sql?view=apspd-2016-au7>

QUESTION 22

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

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

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

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

The query fails to return results.

You need to determine why the query fails.

Solution: You run the following Transact-SQL statements:

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec-requests-transact-sql?view=apspdw-2016-au7>


QUESTION 23

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

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

Your company uses Microsoft SQL Server to deploy a data warehouse to an environment that has a SQL Server Analysis Services (SSAS) instance. The data warehouse includes the **Fact.Order** table as shown in the following table definition. The table has no indexes.





The image shows a screenshot of a database management tool displaying the structure of the Fact.Order table. The table is expanded to show its columns. The columns are listed with their data types and nullability. The columns are: Order Key (bigint, not null), City Key (int, not null), Customer Key (int, not null), Stock Item Key (int, not null), Order Date Key (date, not null), Picked Date Key (date, null), Salesperson Key (int, not null), Picker Key (int, null), Quantity (int, not null), Unit Price (decimal(18,2), not null), Tax Rate (decimal(18,3), not null), Total Excluding Tax (decimal(18,2), not null), Tax Amount (decimal(18,2), not null), and Total Including Tax (decimal(18,2), not null).

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

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

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

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

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

Solution: You create a measure for the **Fact.Order** table.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

You should use a columnstore index.

Columnstore indexes are the standard for storing and querying large data warehousing fact tables. This index uses column-based data storage and query processing to achieve gains up to 10 times the query performance in your data warehouse over traditional row-oriented storage.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/indexes/columnstore-indexes-overview?view=sql-server-2017>

QUESTION 24

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

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

You have a Microsoft SQL server that has Data Quality Services (DQS) installed.

You need to review the completeness and the uniqueness of the data stored in the matching policy.

Solution: You profile the data.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Use a matching rule.

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

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 have a Microsoft SQL server that has Data Quality Services (DQS) installed.

You need to review the completeness and the uniqueness of the data stored in the matching policy.

Solution: You create a matching rule.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Use a matching rule, and use completeness and uniqueness data to determine what weight to give a field in the matching process.

If there is a high level of uniqueness in a field, using the field in a matching policy can decrease the matching results, so you may want to set the weight for that field to a relatively small value. If you have a low level of uniqueness for a column, but low completeness, you may not want to include a domain for that column.

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

QUESTION 26

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

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

You have a Microsoft SQL server that has Data Quality Services (DQS) installed.

You need to review the completeness and the uniqueness of the data stored in the matching policy.

Solution: You modify the weight of the domain in the matching rule.

Does this meet the goal?

A. Yes

B. No

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Use a matching rule, and use completeness and uniqueness data to determine what weight to give a field in the matching process.

If there is a high level of uniqueness in a field, using the field in a matching policy can decrease the matching results, so you may want to set the weight for that field to a relatively small value. If you have a low level of uniqueness for a column, but low completeness, you may not want to include a domain for that column.

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

QUESTION 27

You have a server that has Data Quality Services (DQS) installed.

You create a matching policy that contains one matching rule.

You need to configure the Similarity of Similar percentage that defines a match.

Which similarity percentage will always generate a similarity score of 0?

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

QUESTION 28

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 knowledge base
- C. a matching policy
- D. a subscription view

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

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-2017>

QUESTION 29

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_add_job**
- E. **sys.sp_cdc_change_job**
- F. **sys.sp_cdc_disable_db**
- G. **sys.sp_cdc_enable_db**
- H. **sys.sp_cdc_stop_job**



Correct Answer: F Section:
(none) Explanation

Explanation/Reference:

Explanation:

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

References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-disable-db-transact-sql?view=sql-server-2017>

QUESTION 30

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.

Users report that an application that uses **DB1** is suddenly unresponsive.

You discover that the Integration Services job causes severe blocking issues in the application.

You need to ensure that the users can run the application as quickly as possible.

Your SQL Server login is a member of only the **ssis_admin** database role.

Which stored procedure should you execute?

- 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**
- G. **sys.sp_cdc_enable_db**
- H. **sys.sp_cdc_stop_job**



Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Explanation: sys.sp_cdc_change_job modifies the configuration of a change data capture cleanup or capture job in the current database.

References: <https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-change-job-transact-sql?view=sql-server-2017>

QUESTION 31

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

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

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

You remove the Integration Services job.

You need to stop tracking changes to the database temporarily. The solution must ensure that tracking changes can be restored quickly in a few weeks.

Which stored procedure should you execute?

- A. **catalog.deploy_project**
- B. **catalog.restore_project**
- C. **catalog.stop_operation**
- D. **sys.sp_cdc_add_job**
- E. **sys.sp_cdc_change_job**
- F. **sys.sp_cdc_disable_db**
- G. **sys.sp_cdc_enable_db**
- H. **sys.sp_cdc_stop_job**



Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

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

Incorrect Answers:

H: sys.sp_cdc_stop_job stops a change data capture cleanup or capture job for the current database.

References: <https://docs.microsoft.com/en-us/sql/integration-services/system-stored-procedures/catalog-stop-operation-ssisdb-database?view=sql-server-2017>

QUESTION 32

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

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

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

You discover that the job loads changes from the previous three days only.

You need to ensure that the job loads changes from the previous week.

Which stored procedure should you execute?

- A. **catalog.deploy_project**
- B. **catalog.restore_project**
- C. **catalog.stop_operation**
- D. **sys.sp_cdc_add_job**
- E. **sys.sp_cdc_change_job**
- F. **sys.sp_cdc_disable_db**
- G. **sys.sp_cdc_enable_db**
- H. **sys.sp_cdc_stop_job**



Correct Answer: A

Section: (none)

Explanation

Explanation/Reference: Explanation: catalog.deploy_project deploys a project to a folder in the Integration Services catalog or updates an existing project that has been deployed previously.

References: <https://docs.microsoft.com/en-us/sql/integration-services/system-stored-procedures/catalog-deploy-project-ssisdb-database>

QUESTION 33

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

You are loading data from an OLTP database to a data warehouse.

The OLTP database includes a table for sales data and a table for refund data.

The data warehouse contains a single table for all the sales and refund data.

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

References: <https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/merge-transformation?view=sql-server-2017>

QUESTION 34

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

Correct Answer: G

Section: (none)

Explanation

Explanation/Reference:

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.

Incorrect Answers:

F: The Aggregate transformation applies aggregate functions, such as Average, to column values and copies the results to the transformation output. Besides aggregate functions, the transformation provides the GROUP BY clause, which you can use to specify groups to aggregate across.

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

QUESTION 35

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

You are implementing the data load process for a data warehouse.

The data warehouse uses daily partitions to store data added or modified during the last 60 days. Older data is stored in monthly partitions.

You need to ensure that the ETL process can modify the partition scheme during the data load process.

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

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

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

Incorrect Answers:

A: The Slowly Changing Dimension transformation coordinates the updating and inserting of records in data warehouse dimension tables.

References: <https://social.technet.microsoft.com/wiki/contents/articles/50781.dynamic-management-of-ssas-partitions-with-ssis.aspx>

QUESTION 36

You have a data warehouse named **DW1**.

In **DW1**, 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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

QUESTION 37

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.



Correct Answer: A

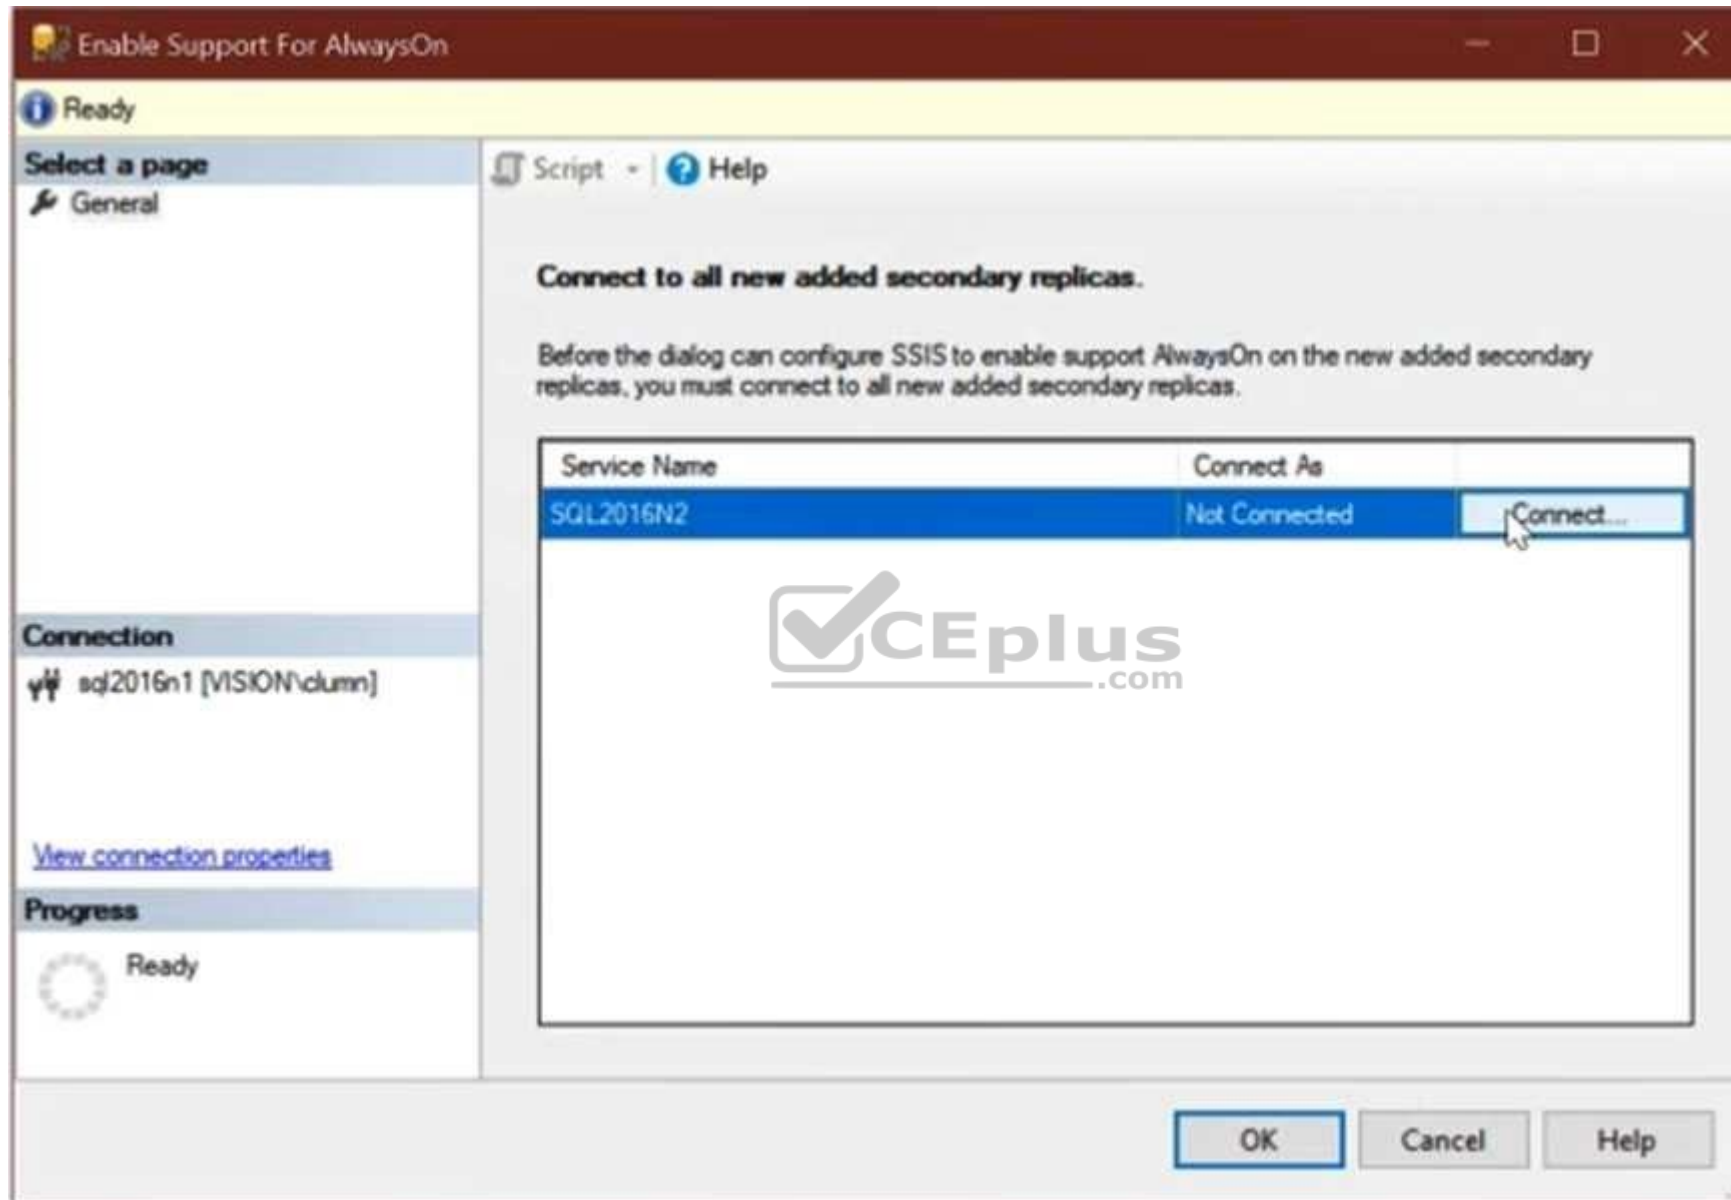
Section: (none)

Explanation

Explanation/Reference:

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://chrislumnah.com/2017/05/09/enabling-alwayson-for-ssisdb/>

QUESTION 38

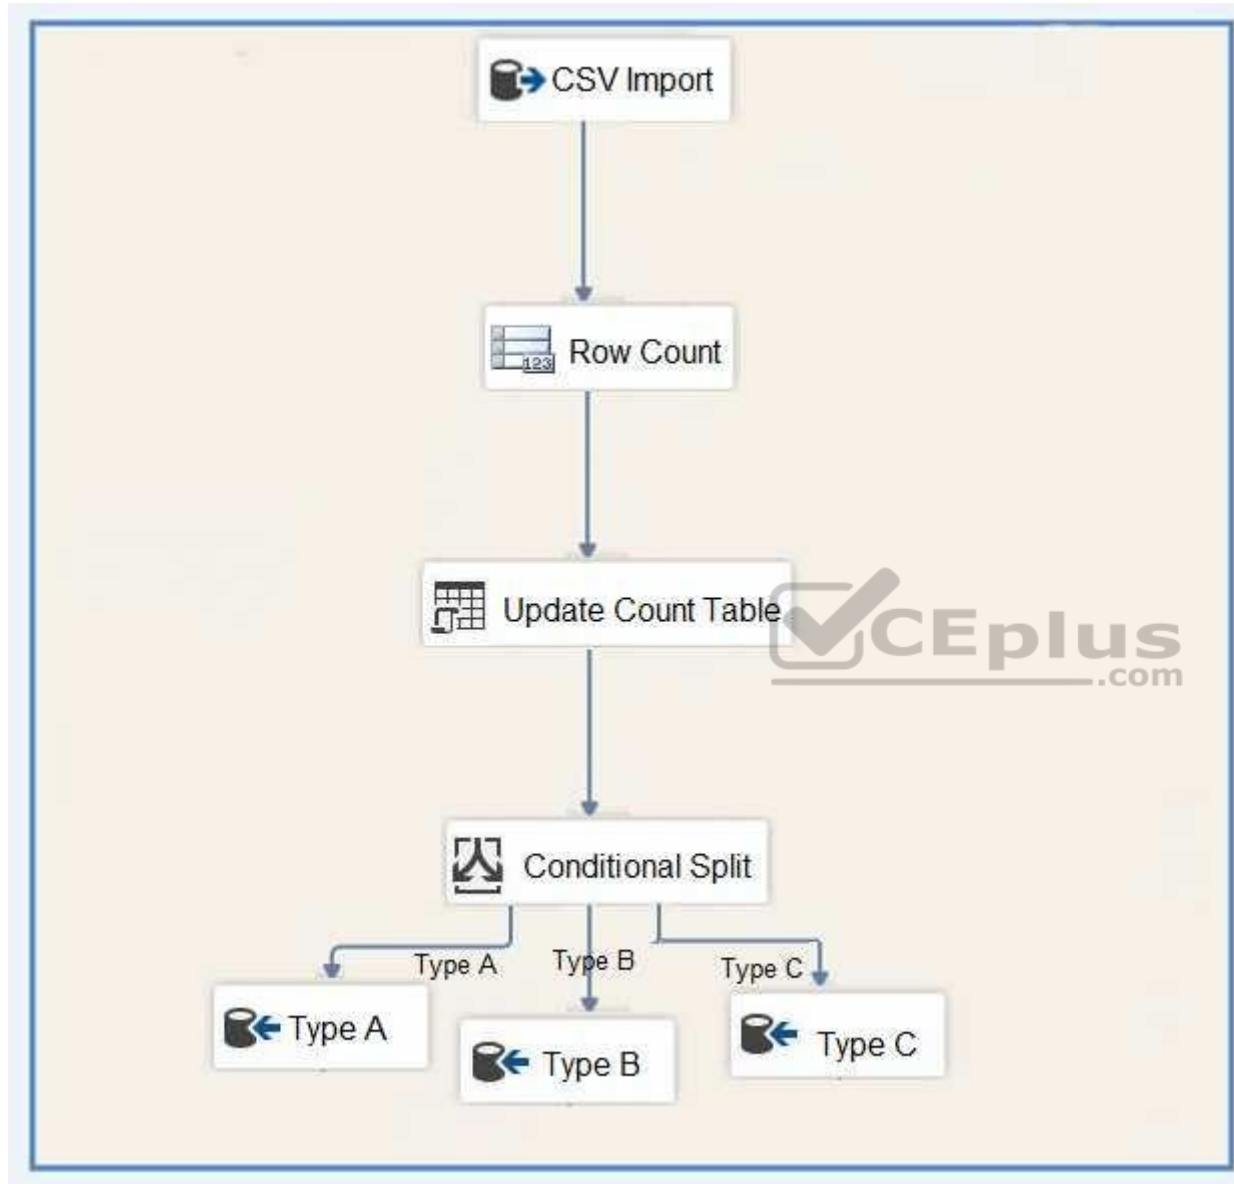
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The conditional split must be before the count.

QUESTION 39

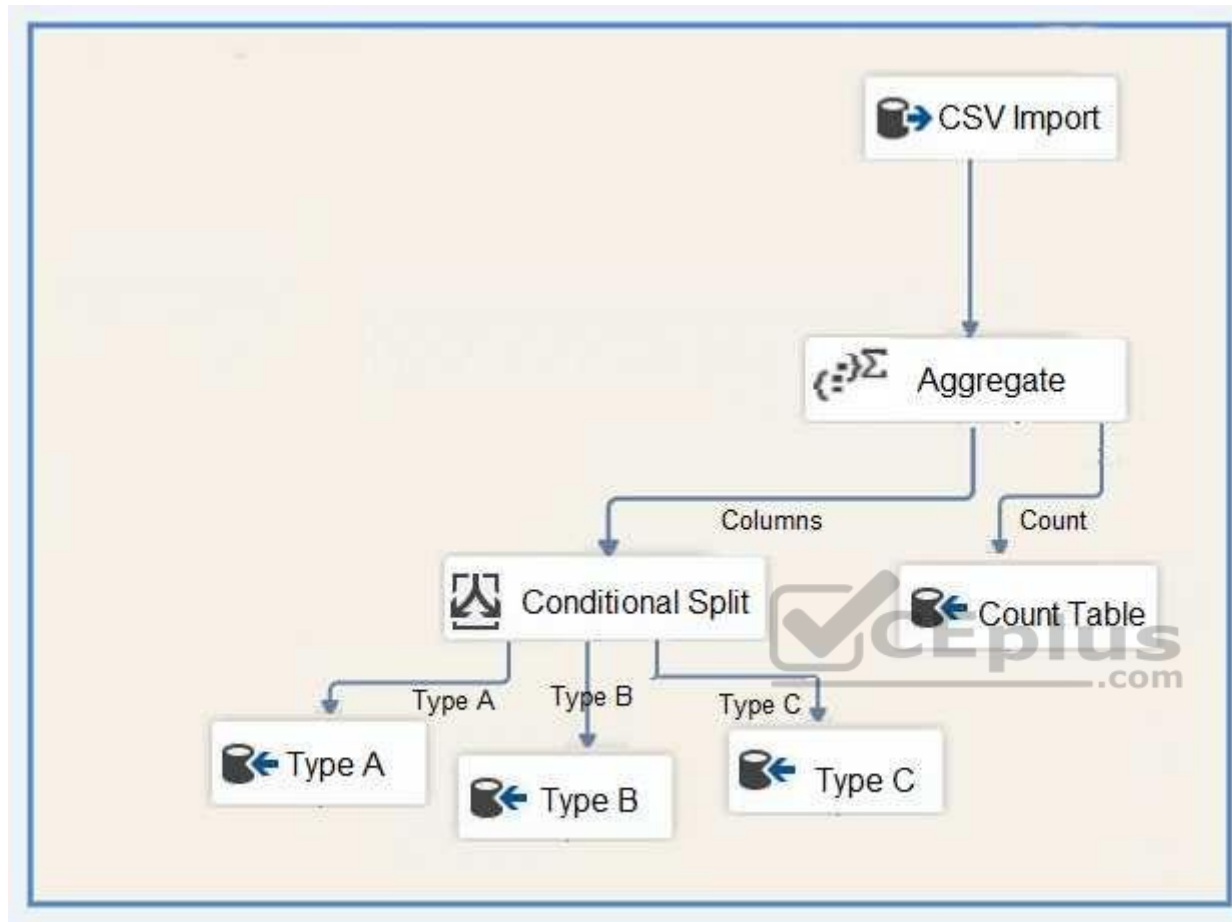
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The conditional split must be before the count.

QUESTION 40

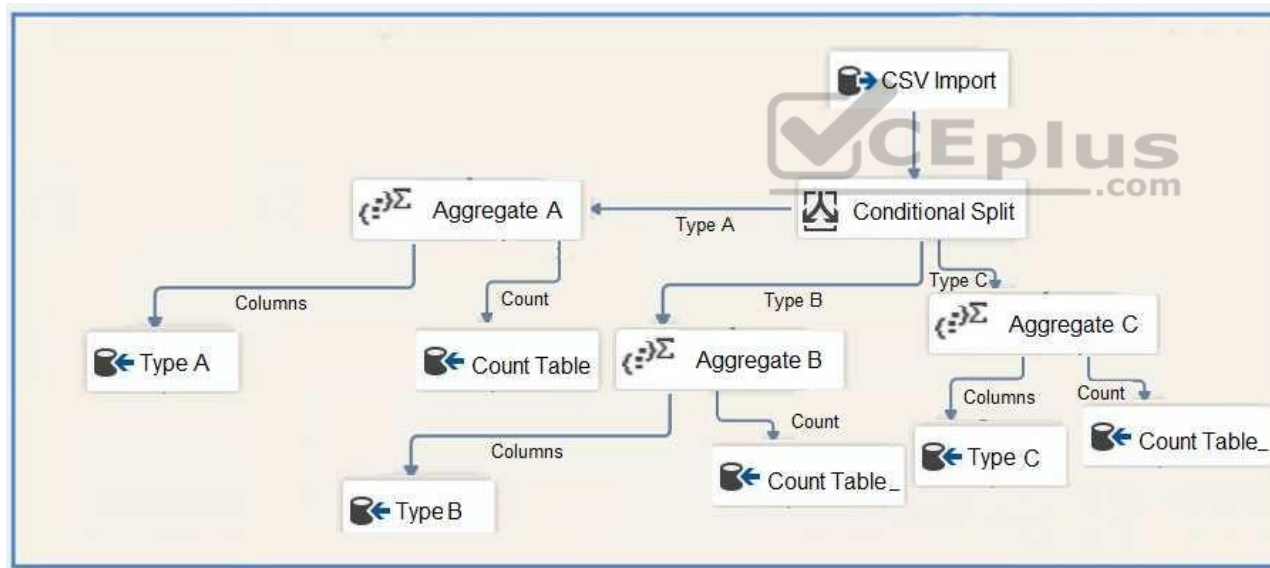
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

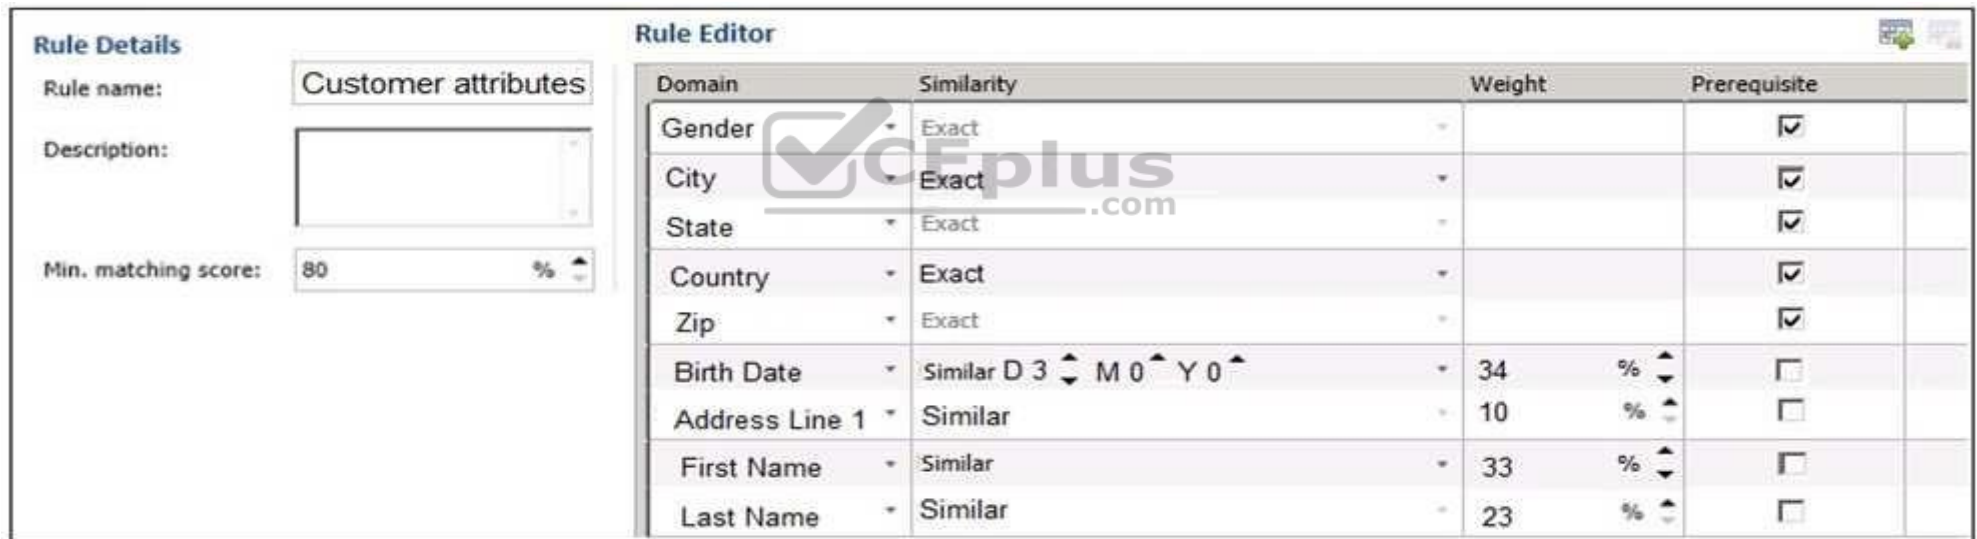
The conditional split is correctly placed before the count.

QUESTION 41

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

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

You configure a new matching policy Master Data Services (MDS) as shown in the following exhibit.



The screenshot shows the 'Rule Editor' interface for a matching policy named 'Customer attributes'. The 'Rule Details' section on the left shows the rule name, a description field, and a minimum matching score of 80%. The 'Rule Editor' section on the right displays a table of attributes and their similarity settings.

Domain	Similarity	Weight	Prerequisite
Gender	Exact		<input checked="" type="checkbox"/>
City	Exact		<input checked="" type="checkbox"/>
State	Exact		<input checked="" type="checkbox"/>
Country	Exact		<input checked="" type="checkbox"/>
Zip	Exact		<input checked="" type="checkbox"/>
Birth Date	Similar D 3 M 0 Y 0	34 %	<input type="checkbox"/>
Address Line 1	Similar	10 %	<input type="checkbox"/>
First Name	Similar	33 %	<input type="checkbox"/>
Last Name	Similar	23 %	<input type="checkbox"/>

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

You verify that the data contains multiple records that have similar address values, and you expect some of the records to match.

You need to increase the likelihood that the records will match when they have similar address values.

Solution: You increase the relative weights for Address Line 1 of the matching policy.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Decrease the Min. matching score.

A data matching project consists of a computer-assisted process and an interactive process. The matching project applies the matching rules in the matching policy to the data source to be assessed. This process assesses the likelihood that any two rows are matches in a matching score. Only those records with a probability of a match greater than a value set by the data steward in the matching policy will be considered a match.

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

QUESTION 42

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

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

You configure a new matching policy Master Data Services (MDS) as shown in the following exhibit.

Rule Details

Rule name:

Description:

Min. matching score: %

Rule Editor

Domain	Similarity	Weight	Prerequisite
Gender	Exact		<input checked="" type="checkbox"/>
City	Exact		<input checked="" type="checkbox"/>
State	Exact		<input checked="" type="checkbox"/>
Country	Exact		<input checked="" type="checkbox"/>
Zip	Exact		<input checked="" type="checkbox"/>
Birth Date	Similar D 3 M 0 Y 0	34 %	<input type="checkbox"/>
Address Line 1	Similar	10 %	<input type="checkbox"/>
First Name	Similar	33 %	<input type="checkbox"/>
Last Name	Similar	23 %	<input type="checkbox"/>

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

You verify that the data contains multiple records that have similar address values, and you expect some of the records to match.

You need to increase the likelihood that the records will match when they have similar address values.

Solution: You decrease the relative weights for Address Line 1 of the matching policy.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Decrease the Min. matching score.

A data matching project consists of a computer-assisted process and an interactive process. The matching project applies the matching rules in the matching policy to the data source to be assessed. This process assesses the likelihood that any two rows are matches in a matching score. Only those records with a probability of a match greater than a value set by the data steward in the matching policy will be considered a match.

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

QUESTION 43

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

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

You configure a new matching policy Master Data Services (MDS) as shown in the following exhibit.



The screenshot shows the 'Rule Editor' interface for a matching policy named 'Customer attributes'. The 'Rule Details' panel on the left shows the rule name, a description field, and a minimum matching score of 80%. The 'Rule Editor' panel on the right contains a table with columns: Domain, Similarity, Weight, and Prerequisite.

Domain	Similarity	Weight	Prerequisite
Gender	Exact		<input checked="" type="checkbox"/>
City	Exact		<input checked="" type="checkbox"/>
State	Exact		<input checked="" type="checkbox"/>
Country	Exact		<input checked="" type="checkbox"/>
Zip	Exact		<input checked="" type="checkbox"/>
Birth Date	Similar D 3 M 0 Y 0	34 %	<input type="checkbox"/>
Address Line 1	Similar	10 %	<input type="checkbox"/>
First Name	Similar	33 %	<input type="checkbox"/>
Last Name	Similar	23 %	<input type="checkbox"/>

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

You verify that the data contains multiple records that have similar address values, and you expect some of the records to match.

You need to increase the likelihood that the records will match when they have similar address values.

Solution: You decrease the minimum matching score of the matching policy.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: A **Section:**
(none) **Explanation**

Explanation/Reference:

Explanation:

We decrease the Min. matching score.

A data matching project consists of a computer-assisted process and an interactive process. The matching project applies the matching rules in the matching policy to the data source to be assessed. This process assesses the likelihood that any two rows are matches in a matching score. Only those records with a probability of a match greater than a value set by the data steward in the matching policy will be considered a match.

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

QUESTION 44

A database has tables named Table1, Table2, and Table3.

- Table1 has a foreign key relationship with Table2.
- Table2 has a foreign key relationship with Table3.
- Table1 does not have a direct relationship with Table3.



You need to recommend an appropriate dimension usage relationship.

What should you recommend?

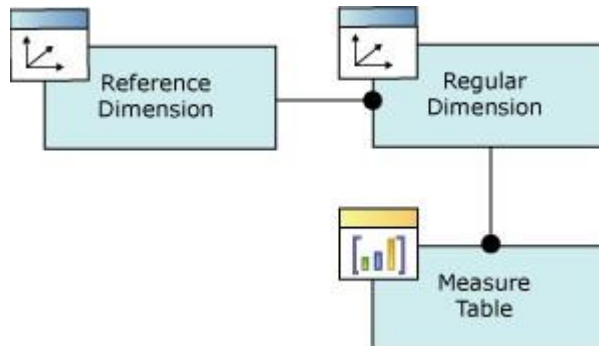
- A. many-to-one relationship
- B. referenced relationship
- C. regular dimension relationship
- D. fact relationship

Correct Answer: B
Section: (none)
Explanation

Explanation/Reference:

Explanation:

A reference dimension relationship between a cube dimension and a measure group exists when the key column for the dimension is joined indirectly to the fact table through a key in another dimension table, as shown in the following illustration.



Incorrect Answers:

C: A regular dimension relationship between a cube dimension and a measure group exists when the key column for the dimension is joined directly to the fact table.

References: <https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models-olap-logical-cube-objects/dimension-relationships>

QUESTION 45

HOTSPOT

You are a data warehouse developer.

You need to create a Microsoft SQL Server Integration Services (SSIS) catalog on a production SQL Server instance.

Which features are needed? To answer, select the appropriate options in the answer area.

Hot Area:



Answer Area

Feature	Yes	No
CLR	<input type="radio"/>	<input type="radio"/>
Automatic SSIS Package Execution	<input type="radio"/>	<input type="radio"/>
in Memory OLTP	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Feature	Yes	No
CLR	<input checked="" type="radio"/>	<input type="radio"/>
Automatic SSIS Package Execution	<input checked="" type="radio"/>	<input type="radio"/>
in Memory OLTP	<input type="radio"/>	<input checked="" type="radio"/>

Section: (none)

Explanation

Explanation/Reference:

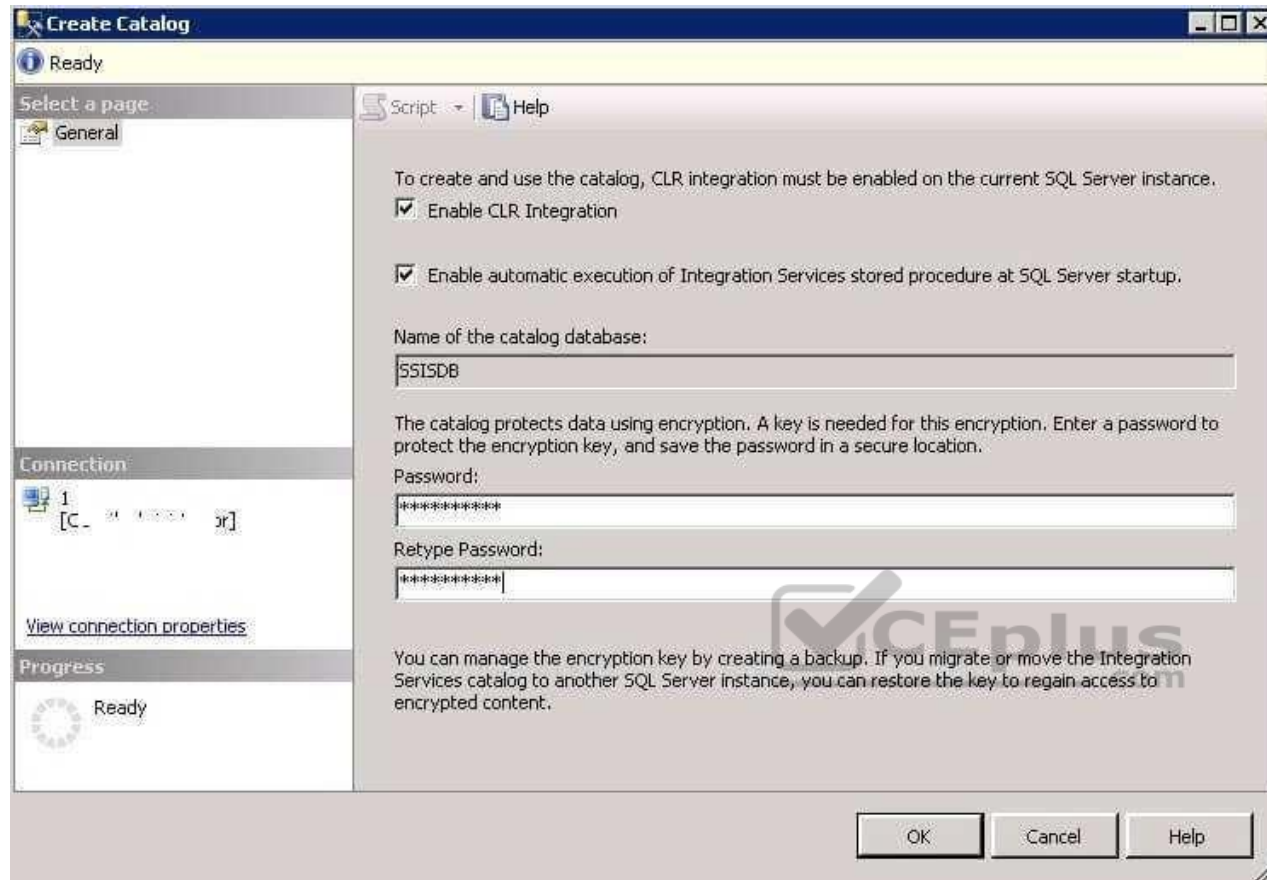
Explanation:

Box 1: Yes

"Enable CLR Integration" must be selected because the catalog uses CLR stored procedures.

Box 2: Yes

Once you have selected the "Enable CLR Integration" option, another checkbox will be enabled named "Enable automatic execution of Integration Services stored procedure at SQL Server startup". Click on this check box to enable the catalog startup stored procedure to run each time the SSIS server instance is restarted.



Box 3: No

References: <https://www.mssqltips.com/sqlservertip/4097/understanding-the-sql-server-integration-services-catalog-and-creating-the-ssisdb-catalog/>

QUESTION 46

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

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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.

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

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

QUESTION 47

DRAG DROP

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.

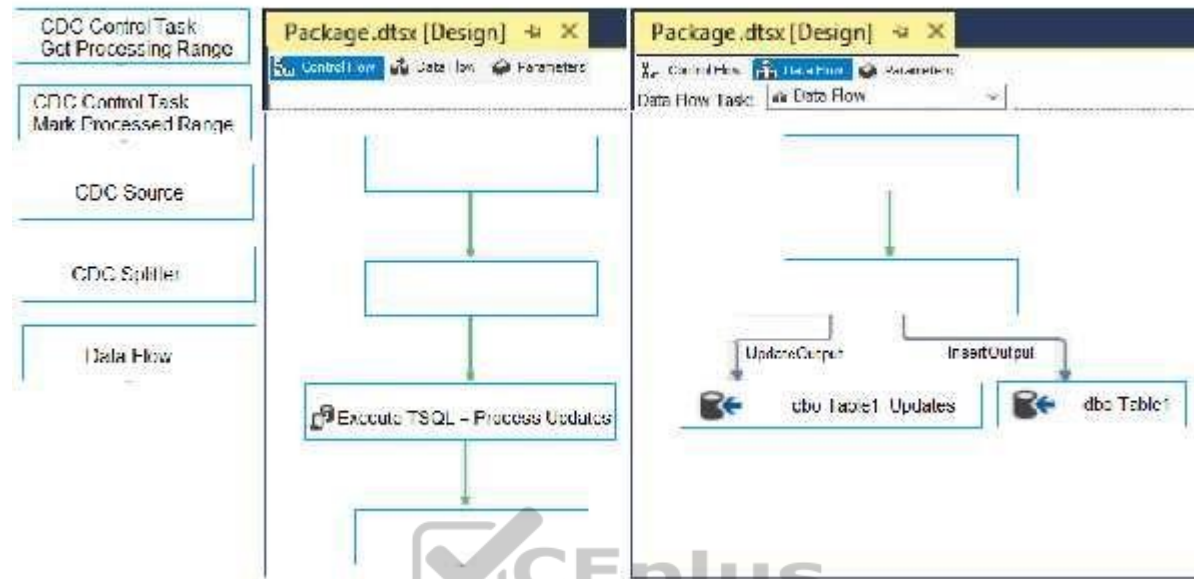
You need to complete the design of the SSIS package.

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.

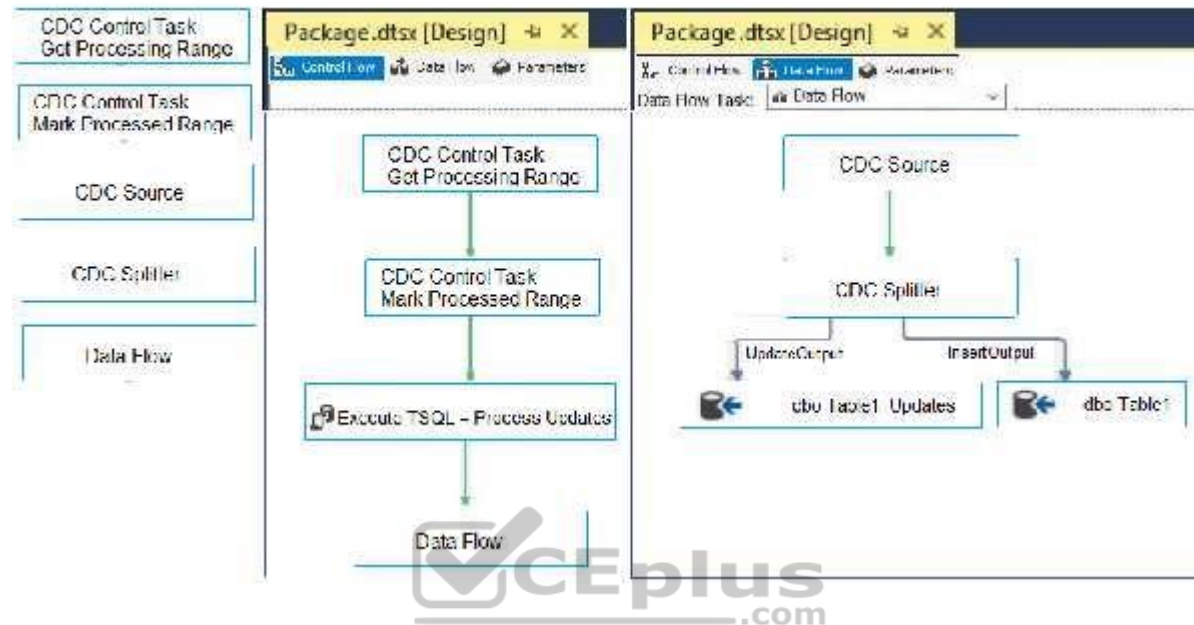
Select and Place:

Answer Area



Correct Answer:

Answer Area



Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 48

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

Many reports query `DW1` for the past year of data.

Users frequently report that the reports are slow.

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

What should you do?

- A. Move the historical data to SAS disks and move the data from the past year to SSD disks. Run the `ALTER TABLE` statement.
- B. Move all the data to SSD disks. Load and archive the data by using partition switching.
- C. Move all the data to SAS disks. Load and archive the data by using partition switching.
- D. Move the historical data to SAS disks and move the data for the past year to SSD disks. Create a distributed partitioned view.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

Explanation:

We use `ALTER TABLE` to partition the table.

Incorrect Answers:

D: A Distributed Partitioned View contains participating tables from multiple SQL Server instances, which can be used to distribute the data processing load across multiple servers. Another advantage for the SQL Server Partitioned Views is that the underlying tables can participate in more than one Partitioned View, which could be helpful in some implementations.

QUESTION 49

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`.

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 50

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

Start of repeated scenario

You have a Microsoft SQL Server data warehouse instance that supports several client applications.

The data warehouse includes the following tables: **Dimension.SalesTerritory**, **Dimension.Customer**, **Dimension.Date**, **Fact.Ticket**, and **Fact.Order**. The **Dimension.SalesTerritory** and **Dimension.Customer** tables are frequently updated. The **Fact.Order** table is optimized for weekly reporting, but the company wants to change it to daily. The **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 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.

The following requirements must be met:

- Implement table partitioning to improve the manageability of the data warehouse and to avoid the need to repopulate all transactional data each night. Use a partitioning strategy that is as granular as possible.

- Partition the **Fact.Order** table and retain a total of seven years of data.
- Partition the **Fact.Ticket** table and retain seven years of data. At the end of each month, the partition structure must apply a sliding window strategy to ensure that a new partition is available for the upcoming month, and that the oldest month of data is archived and removed.
- Optimize data loading for the **Dimension.SalesTerritory**, **Dimension.Customer**, and **Dimension.Date** tables.
- Incrementally load all tables in the database and ensure that all incremental changes are processed.
- Maximize the performance during the data loading process for the **Fact.Order** partition.
- Ensure that historical data remains online and available for querying.
- Reduce ongoing storage costs while maintaining query performance for current data.

You are not permitted to make changes to the client applications.

End of repeated scenario

You need to implement the data partitioning strategy.

How should you partition the **Fact.Order** table?

- A. Create 17,520 partitions.
- B. Create 2,557 partitions.
- C. Use a granularity of one month.
- D. Create 1,460 partitions.



Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

We create one partition for each day, which means that a granularity of one day is used.

Note: If we calculate the partitions that are needed, we get: 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.

The Fact.Order table is optimized for weekly reporting, but the company wants to change it to daily.

Maximize the performance during the data loading process for the Fact.Order partition.

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

QUESTION 51

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.

Correct Answer: A

Section: (none)

Explanation

Explanation/Reference:

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>

QUESTION 52

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?



<https://vceplus.com/>

- A. Create an attribute group that consists of the `Category` and `Subcategory` attributes. Create a custom index for the attribute group.
- B. Publish a business rule that will be used by the `Product` entity.
- C. Create a derived hierarchy based on the `Category` and `Subcategory` attributes. Use the `Category` attribute as the top level for the hierarchy.
- D. Create an attribute group that consists of the `Category` and `Subcategory` attributes. Publish a business rule for the attribute group.

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

Incorrect Answers:

A: In Master Data Services, attribute groups help organize attributes in an entity. When an entity has lots of attributes, attribute groups improve the way an entity is displayed in the Master Data Manager web application.

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

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

QUESTION 53

You create a Master Data Services (MDS) model that manages the master data for a Product dimension. The Product dimension has the following properties:

- All the members of the Product dimension have a product type, a product subtype, and a unique product name.
- Each product has a single product type and a single product subtype.
- The product type has a one-to-many relationship to the product subtype.

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

What should you add to the model?

- A. a subscription view

- B. a derived hierarchy
- C. a recursive hierarchy
- D. an explicit hierarchy

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

Incorrect Answers:

C: In Master Data Services, a recursive hierarchy is a derived hierarchy that includes a recursive relationship. A recursive relationship exists when an entity has a domain-based attribute based on the entity itself.

D: In Master Data Services, an explicit hierarchy organizes members from a single entity in any way you specify. The structure can be ragged and unlike derived hierarchies, explicit hierarchies are not based on domain-based attribute relationships.

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

QUESTION 54

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

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

`Sales` contains details of records that have a type of refund and records that have a type of sales.

The data warehouse design contains a table for sales data and a table for refund data.

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

- A. the Slowly Changing Dimension transformation
- B. the Conditional Split transformation
- C. the Merge transformation
- D. the Data Conversion transformation

- E. an Execute SQL task
- F. the Aggregate transformation
- G. the Lookup transformation

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

References: <https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/conditional-split-transformation>

QUESTION 55

DRAG DROP

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

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

NOTE: Each correct selection is worth one point.

Select and Place:

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

Correct Answer:

Transact-SQL segments

Answer Area

MERGE

BULK

INSERT Sales.Invoices

FROM

'\\share\data\file1.csv'

WITH

(FORMAT = 'CSV')

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Example:

```
BULK INSERT Sales.Orders
FROM '\\SystemX\DiskZ\Sales\data\orders.csv'
WITH ( FORMAT='CSV');
```



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 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

Correct Answer: E

Section: (none)

Explanation

Explanation/Reference:

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.

Incorrect Answers:

C: UNION ALL. Incorporates all rows into the results. This includes duplicates. If not specified, duplicate rows are removed.

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

QUESTION 57

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

You are developing a Microsoft SQL Server Integration Services (SSIS) package. The package design consists of two differently structured sources in a single data flow. The Sales source retrieves sales transactions from a SQL Server database, and the Product source retrieves product details from an XML file.

You need to combine the two data flow sources into a single output dataset.

Which SSIS Toolbox item should you use?

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

Correct Answer: G

Section: (none)

Explanation

Explanation/Reference:

Explanation:

The Merge Join transformation provides an output that is generated by joining two sorted datasets using a FULL, LEFT, or INNER join. For example, you can use a LEFT join to join a table that includes product information with a table that lists the country/region in which a product was manufactured. The result is a table that lists all products and their country/region of origin.

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

QUESTION 58

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

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

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

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

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

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

NOTE: Each correct selection is worth one point.

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

- E. Drop the `Color` column from `Dim1`.
- F. Modify the data type of the `ProductKey` column in `Dim1` to `char (18)`.

Correct Answer: AD

Section: (none)

Explanation

Explanation/Reference:

QUESTION 59

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

DateID	Stock-ID	Open-ingPrice	Closing-Price	Quanti-tyTraded	Bro-kerID	Num-berOfTra-des	Market-ID
20170301	22	30.20	34.23	100	10	1	1
20170301	31	10.05	12.23	110	10	2	2
20170302	22	30.89	34.76	899	5	1	1

You need to identify a column that can be aggregated across all dimensions.

Which column should you identify?

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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

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

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

QUESTION 60

You are designing a warehouse named DW1.

A table named Table1 is partitioned by using the following partitioning scheme and function.

```
AS RANGE LEFT FOR VALUES ('20150101', '20160101', '20170101', '20180101', '20190101',
'20200101');
GO
CREATE PARTITION SCHEME schema1
AS PARTITION function1
ALL TO ([primary]);
GO

CREATE TABLE table1
(MyId BIGINT IDENTITY (1,1),
OrderDate datetime,
DueDate datetime,
AccountNumber nvarchar(15)
...
PRIMARY KEY (MyId, OrderDate))
ON schema1 (OrderDate)
GO
```



Reports are generated from the data in Table1.

You need to ensure that queries to DW1 return results as quickly as possible.

Which column should appear in the **WHERE** statement clause of the query?

- A. AccountNumber
- B. MyId
- C. DueDate
- D. OrderDate

Correct Answer: D

Section: (none)

Explanation

Explanation/Reference:

QUESTION 61

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

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

You create a set of Microsoft SQL Server Integration Services (SSIS) packages to support an ETL process.

You need to deploy the SSIS packages and meet the following requirements:

- Maximize ease of use for backup processes.
 - Manage security of packages from a single place.
 - Maintain versioning of packages.
 - Implement parameters to assign values to packages.
 - Use stored procedures to run SSIS packages.
- Use SQL Server Agent to run packages.

Solution: You use the Project Deployment model. You save the packages to team Foundation Server.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Built-in stored procedures only work for with an SSIS catalog to store.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/packages/run-integration-services-ssis-packages>

QUESTION 62

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

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

You are the administrator of a Microsoft SQL Server Master Data Services (MDS) instance. The instance contains a model named *Geography* and a model named *Customer*. The *Geography* model contains an entity named *CountryRegion*.

You need to ensure that the *CountryRegion* entity members are available in the customer model.

Solution: Configure an entity sync relationship to replicate the *CountryRegion* entity.

Does this meet the goal?

- A. Yes
- B. No

**Correct Answer: A Section:
(none) Explanation**

Explanation/Reference:

Explanation:

Entity sync is a one-way and repeatable synchronization between entity versions. It enables you to share entity data between different models. You can keep a single source of truth in one model and reuse this master data in other models. For example, you can store US state data in one model entity and reuse that data in other models.

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

QUESTION 63

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: Create a `CountryRegion` entity in the `Customer` model. In the `Geography` model, create a subscription view and load data to the entity-based staging table that contains the `CountryRegion` entity.

Does this meet the goal?

- A. Yes
- B. No

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

Explanation:

Instead configure an entity sync relationship to replicate the `CountryRegion` entity.

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

QUESTION 64

You manage the user accounts in master data Services (MDS).



You need to assign a user access to the MDS data and functions.

Which two components must you assign? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. file share permissions
- B. model object permissions
- C. functional area permissions
- D. SQL Database permissions

Correct Answer: BC

Section: (none)

Explanation

Explanation/Reference:

B: In Master Data Services, assign permissions to model objects when you need to give a user or group access to data in the Explorer functional area of Master Data Manager, or when you need to make a user or group an administrator.

C: Assign functional area permission to grant users or groups access to the functional areas of Master Data Manager.

To assign functional area permissions

1. In Master Data Manager, click User and Group Permissions.
2. On the Users or Groups page, select the row for the user or group that you want to edit.
3. Click Edit selected user.
4. Click the Functions tab.
5. Click Edit.
6. Click a functional area and click the Add arrow.
7. When you are done, click Save.

References: <https://docs.microsoft.com/en-us/sql/master-data-services/assign-model-object-permissions-master-data-services>

<https://docs.microsoft.com/en-us/sql/master-data-services/assign-functional-area-permissions-master-data-services>

QUESTION 65

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

You are implementing a Microsoft SQL Server data warehouse with a multi-dimensional data model. You are optimizing the data model across multiple cubes.

You examine the existing data model and observe that each cube contains multiple date dimension tables.

You need to optimize the data model in and across cubes without changing the overall table structure.

What 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

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

In data warehousing, a conformed dimension is a dimension that has the same meaning to every fact with which it relates. Conformed dimensions allow facts and measures to be categorized and described in the same way across multiple facts and/or data marts, ensuring consistent reporting across the enterprise.

A conformed dimension can exist as a single dimension table that relates to multiple fact tables within the same data warehouse, or as identical dimension tables in separate data marts. Date is a common conformed dimension because its attributes (day, week, month, quarter, year, etc.) have the same meaning when joined to any fact table.

Incorrect Answers:

H: Dimension usage defines the relationships between a cube dimension and the measure groups in a cube.

References: <https://searchdatamanagement.techtarget.com/definition/conformed-dimension>

QUESTION 66

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.

Business users observe that the value displayed for the profit margin is the average profit margin across all products. The profit margin must be calculated per product.

You need to modify the existing model without changing the overall structure of the data model.

What 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

Correct Answer: G

Section: (none)

Explanation

Explanation/Reference:

Explanation:

There are three types of facts:

Non-Additive: Non-additive facts are facts that cannot be summed up for any of the dimensions present in the fact table.

Additive: Additive facts are facts that can be summed up through all of the dimensions in the fact table.

Semi-Additive: Semi-additive facts are facts that can be summed up for some of the dimensions in the fact table, but not the others.

References: <https://www.1keydata.com/datawarehousing/fact-table-types.html>

QUESTION 67

You have a Microsoft SQL Server Integration Services (SSIS) package that includes a custom Script task. You set a breakpoint at line 97 as shown in the exhibit. (Click the **Exhibit** tab.)

C# Methods



```
91     public void Main()
92     {
93         //TODO: Add your code here
94         string convertedvalue = "";
95         int myvar = 0;
96
97         convertedvalue = ConvertToString(myvar);
98
99         if (convertedvalue.Length > 0)
100         {
101             Dts.TaskResult = (int)ScriptResults.Success;
102         }
103         else
104         {
105             Dts.TaskResult = (int)ScriptResults.Failure;
106         }
107     }
108 }
109
110 private string ConvertToString(int avalue)
111 {
112     string returnValue = "";
113
114     returnValue = avalue.ToString();
115
116     return returnValue;
117 }
118
119 ScriptResult is declaration
```

You run the code and the debugger pauses at line 97.

You need to debug the `ConvertToString` method.

What should you do?

- A. Use the Step Over command.
- B. Use the Toggle Breakpoint command.
- C. Use the Step Into command.
- D. Modify the Exception Settings options.

Correct Answer: C

Section: (none)

Explanation

Explanation/Reference:

Explanation:

On a nested function call, Step Into steps into the most deeply nested function. For example, if you use Step Into on a call like `Func1(Func2())`, the debugger steps into the function `Func2`.

References: <https://docs.microsoft.com/en-us/visualstudio/debugger/navigating-through-code-with-the-debugger>

QUESTION 68

You are developing a Microsoft SQL Server Integration Service (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. `System.Diagnostics.EventLog.WriteEntry("SSIS", "Script executed with return result" + ScriptResults.Success, System.Diagnostics.EventLogEntryType.Information)`
- B. `System.Diagnostics.EventLog.WriteEntry("SSIS", "Script executed with return result" + Dts.TaskResult, System.Diagnostics.EventLogEntryType.Information)`
- C. `Dts.TaskResult = (int)ScriptResults.Failure`

D. `Dts.Events.FireInformation(0, "SSIS", "Script executed with return result" + Dts.TaskResult, String.Empty, 0)`

Correct Answer: B

Section: (none)

Explanation

Explanation/Reference:

References:

<https://support.microsoft.com/en-gb/help/906560/how-to-write-information-to-the-application-event-log-by-using-a-scrip>

