

Exam Questions 70-767

Implementing a SQL Data Warehouse (beta)

<https://www.2passeasy.com/dumps/70-767/>



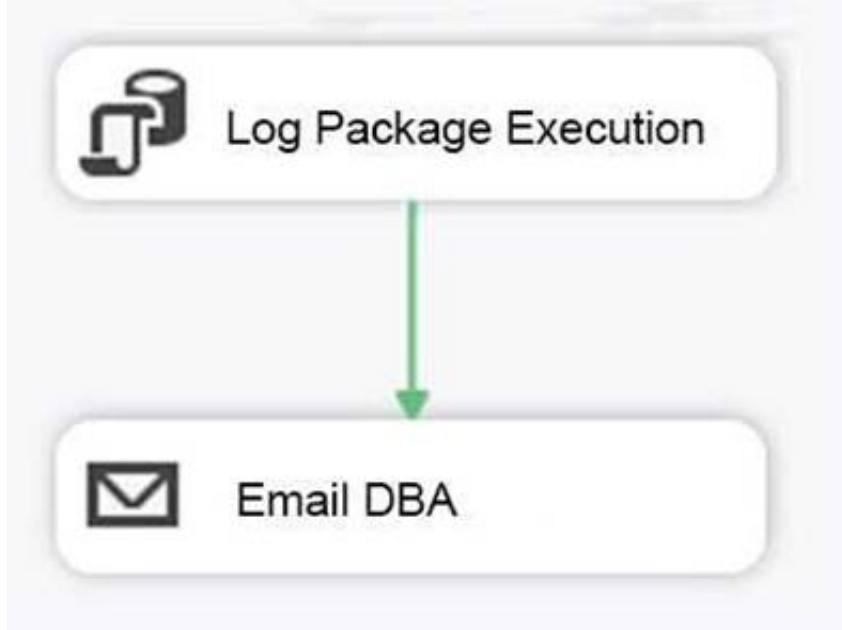
NEW QUESTION 1

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

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

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

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



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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 2

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

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

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

What change should you make?

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

Answer: A

Explanation:

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

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

References:

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

NEW QUESTION 3

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

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

You have a Microsoft Azure SQL Data Warehouse instance that must be available six months a day for reporting.

You need to pause the compute resources when the instance is not being used. Solution: You use SQL Server Configuration Manager.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

NEW QUESTION 4

You have a database named DB1. You create a Microsoft SQL Server Integration Services (SSIS) package that incrementally imports data from a table named Customers. The package uses an OLE DB data source for connections to DB1. The package defines the following variables.

Variable name	Data type	Description
LastKey	Int64	LastKey stores the last identifier used in the imported table.
TableName	String	TableName stores the name of the imported table.

To support incremental data loading, you create a table by running the following Transact-SQL segment:

```
CREATE TABLE LastKeyByTable (
    Id int IDENTITY(1,1) PRIMARY KEY,
    TableName sysname UNIQUE,
    LastKey bigint
)
```

You need to create a DML statements that updates the LastKeyByTable table.

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

Answer Area

UPDATE dbo.LastKeyByTable

SET

LastKey = ?

LastKey = @A

LastKey = @B

LastKey = @LastKey

WHERE

TableName = ?

TableName = @A

TableName = @B

TableName = @TableName

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

```
UPDATE dbo.LastKeyByTable
```

```
SET
```

LastKey =?

LastKey = @A

LastKey = @B

LastKey = @LastKey

```
WHERE
```

TableName =?

TableName = @A

TableName = @B

TableName = @TableName

NEW QUESTION 5

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

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

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

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

Answer: D

NEW QUESTION 6

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

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

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

You need to review the completeness and the uniqueness of the data stored in the matching policy. Solution: You create a matching rule.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

References:

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

NEW QUESTION 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 the series.

Start of repeated scenario

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

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

The data warehouse includes the following tables:

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

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

Users report that they are becoming confused about which city table to use for various queries. You plan to create a new schema named Dimension and change the name of the dbo.dimension_city table to Dimension.city. Data loss is not permissible, and you must not leave traces of the old table in the data warehouse. The fact.Transaction table has measures named RawCost and Totaisale that calculate the wholesale cost of materials. You plan to create a measure that calculates the profit margin based on the two existing measures.

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

Results Messages

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

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

End of repeated scenario

You need to create the ProfitMargin measure for the fact.Transaction table.

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

Answer area

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer area

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

NEW QUESTION 8

You have a data warehouse.

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

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

Actions

- Add a file to the LowCost filegroup.
- Rename the Fact.ErrorLog table to Fact.ErrorLogBak.
- Drop the Fact.ErrorLog table.
- Create a new Fact.ErrorLog table on the LowCost filegroup.
- Add a filegroup named LowCost to the database.
- Reorganize the clustered index on the Fact.ErrorLog table in the new filegroup.
- Rebuild the clustered index on the Fact.ErrorLog table in the new filegroup.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Step 2:

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

Step 3:

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

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

This recreates the same index but on the SECONDARY filegroup.

References:

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

NEW QUESTION 9

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

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

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

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

Answer: H

NEW QUESTION 10

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

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

Answer: A

NEW QUESTION 10

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

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

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

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

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

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

Actions

INSERT SELECT

MERGE

SWITCH

DELETE

SPLIT

Answer area

First action

Action

Action

Second action

Action

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

From scenario: - Partition the Fact.Ticket table and retain seven years of data. At the end of each month, the partition structure must apply a sliding window strategy to ensure that a new partition is available for the upcoming month, and that the oldest month of data is archived and removed.

The detailed steps for the recurring partition maintenance tasks are: References:

<https://docs.microsoft.com/en-us/sql/relational-databases/tables/manage-retention-of-historical-data-in-system-v>

NEW QUESTION 12

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

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

You have the following line-of-business solutions:

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

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

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

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

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

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

- A. Yes
 B. No

Answer: B

Explanation:

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

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

NEW QUESTION 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 developing a Microsoft SQL Server Integration Services (SSIS) package. The package design consists of the sources shown in the following diagram:



Each source contains data that is not sorted.

You need to combine data from all of the sources into a single dataset. Which SSIS Toolbox item should you use?

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

Answer: C

NEW QUESTION 16

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

You need to recommend a solution to perform the actions.

What should you recommend?

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

Answer: A

Explanation:

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

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

References:

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

NEW QUESTION 20

You manage Master Data Services (MDS).

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

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

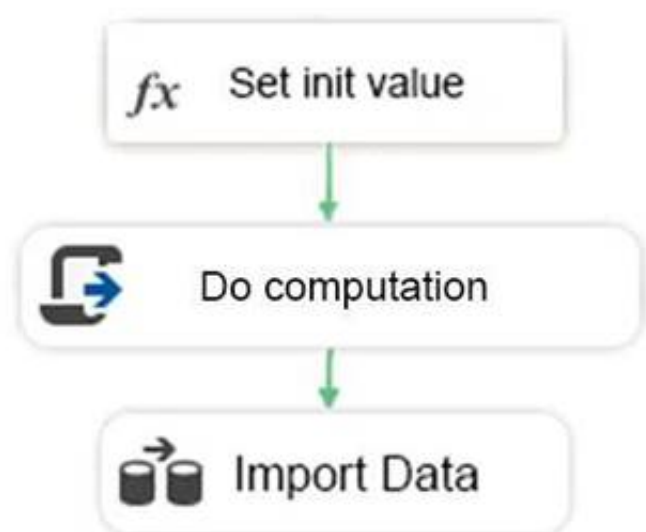
You need to configure the Transaction Log Type setting. Which type should you use?


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

Answer: D

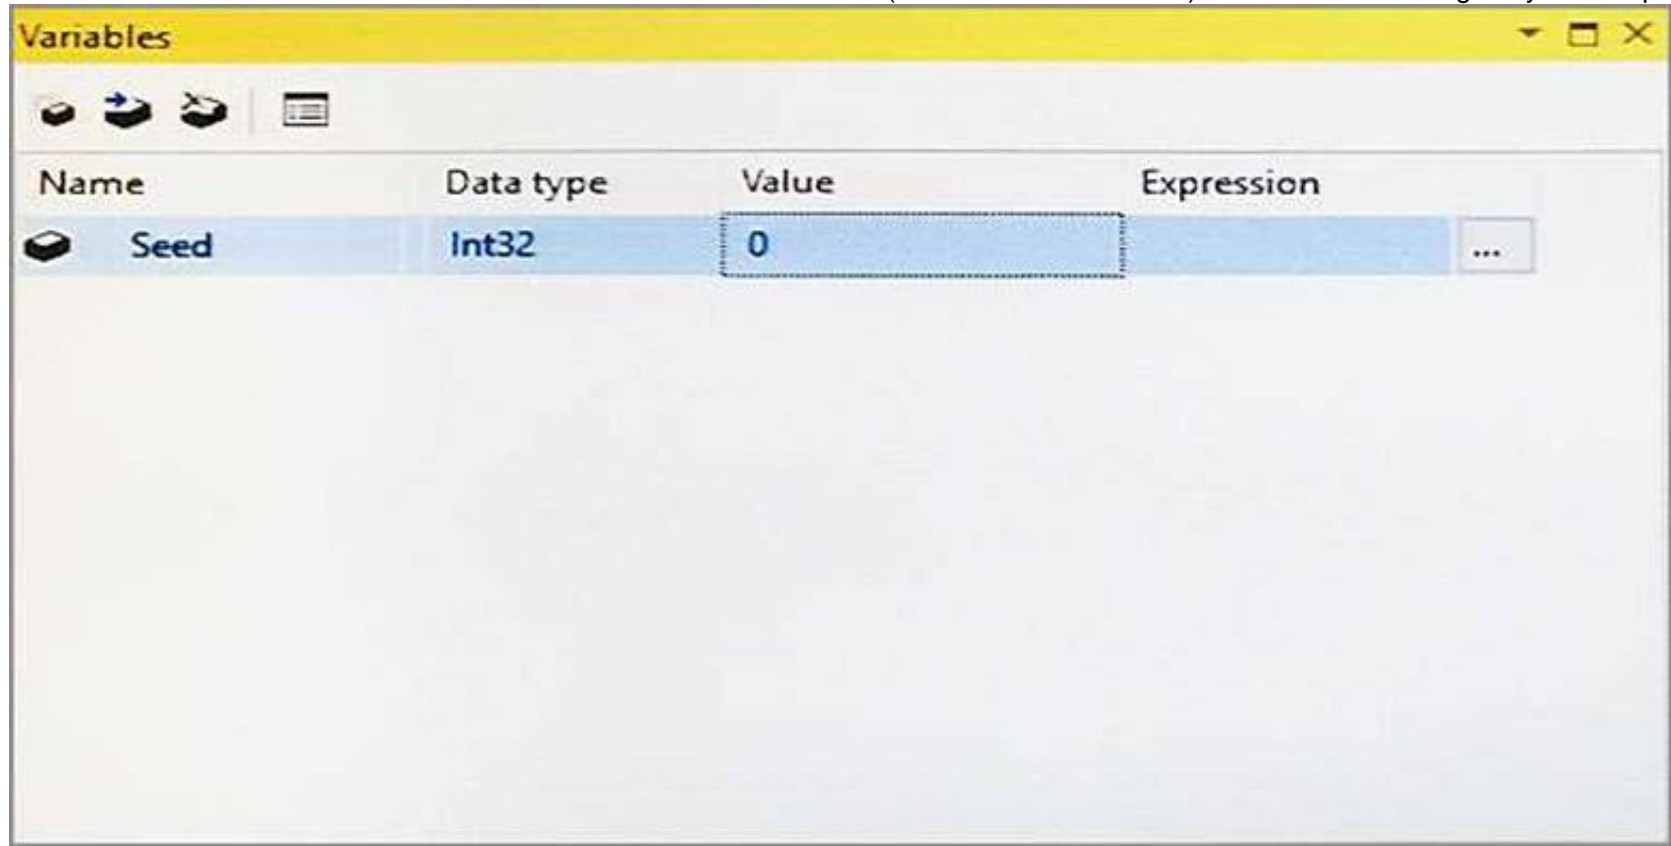
NEW QUESTION 21

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





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



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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 23

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

References:

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

NEW QUESTION 27

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

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

You have the following line-of-business solutions:

- ☐ ERP system
- ☐ Online WebStore
- ☐ Partner extranet

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 28

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

Start of repeated scenario

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

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

The data warehouse includes the following tables:

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

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

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

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

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

End of repeated scenario

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

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

Transact-SQL segments

EXEC sp_rename 'dbo.dim_City', 'City'

ALTER SCHEMA Dimension TRANSFER dbo.City

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

SELECT *
INTO Dimension.City
FROM dbo.dim_City

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

Answer area

CREATE SCHEMA Dimension
GO

Transact-SQL segment

Transact-SQL segment

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Transact-SQL segments

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

```
ALTER SCHEMA Dimension TRANSFER dbo.City
```

```
DROP TABLE dbo.dim_City
```

```
GO
CREATE TABLE Dimension.City( ... )
```

```
SELECT *
INTO Dimension.City
FROM dbo.dim_City
```

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

Answer area

```
CREATE SCHEMA Dimension
GO
```

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

```
DROP TABLE dbo.dim_City
```

```
GO
CREATE TABLE Dimension.City( ... )
```

NEW QUESTION 32

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

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

You have a Microsoft SQL server that has Data Quality Services (DQS) installed. You need to review the completeness and the uniqueness of the data stored in the matching policy. Solution: You modify the weight of the domain in the matching rule.

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

References:

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

NEW QUESTION 34

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

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

You have an on-premises Microsoft SQL Server instance and a Microsoft Azure SQL Data Warehouse instance. You move data from the on-premises database to the data warehouse once each day by using a SQL Server Integration Services (SSIS) package.

You observe that the package no longer completes within the allotted time. You need to determine which tasks are taking a long time to complete.

Solution: You alter the package to log the start and completion times for a task to a table in the on-premises SQL Server instance.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

NEW QUESTION 35

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

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

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

You need to ensure that the countryRegion entity members are available in the customer model. Solution: Configure an entity sync relationship to replicate the CountryRegion entity.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 38

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

You have a database named DB1.

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

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

Answer: D

Explanation:

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

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

Note: sys.sp_cdc_change_job

References:

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

NEW QUESTION 40

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

You are developing a Microsoft SQL Server Integration Services (SSIS) package. You need to use XPath to extract information from documents.

Which SSIS Toolbox item should you use?

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

Answer: B

NEW QUESTION 45

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

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

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

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

Answer Area

CREATE CLUSTERED INDEX lx_product

NONCLUSTERED

ON dbo.Products (ProductReference)

(QuantityInStock)

(ProductName)

INCLUDE (Products)

(ProductReference)

(ProductName, ProductReference)

A. Mastered

B. Not Mastered

Answer: A

Explanation:

Answer Area

CREATE INDEX lx_product

CLUSTERED
 NONCLUSTERED

ON dbo.Products

(ProductReference)
 (QuantityInStock)
 (ProductName)

INCLUDE

(Products)
 (ProductReference)
 (ProductName, ProductReference)

NEW QUESTION 46

You have a data warehouse named DW1 that contains 20 years of data. DW1 contains a very large fact table. New data is loaded to the fact table monthly. Many reports query DW1 for the past year of data. Users frequently report that the reports are slow. You need to modify the fact table to minimize the amount of time it takes to run the reports. The solution must ensure that other reports can continue to be generated from DW1. What should you do?

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

Answer: A

Explanation:

We use ALTER TABLE to partition the table.

NEW QUESTION 47

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

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

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

You remove the Integration Services job.

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

Which stored procedure should you execute?

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

Answer: C

Explanation:

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

References:

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

NEW QUESTION 49

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

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

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

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

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



The image shows a 'Knowledge Base Management' dialog box. It has a title bar with a database icon and the text 'Knowledge Base Management'. Below the title bar is a section titled 'New Knowledge Base'. This section contains several fields and a list box:

- Name:** A text box containing 'Sport Names'.
- Description:** A text box that is currently empty.
- Create Knowledge Base from:** A dropdown menu with three options: 'None', 'Existing Knowledge Base', and 'Import from DQS file'. The 'None' option is currently selected.
- Select Activity:** A dropdown menu with three options: 'Domain Management', 'Knowledge Discovery', and 'Matching Policy'. The 'Knowledge Discovery' option is currently selected.

At the bottom right of the dialog box are two buttons: 'Cancel' and 'Next'.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Spot 1: Create Knowledge base from: None

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

NEW QUESTION 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 this series.

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


Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

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

 - Reduce ongoing storage costs while maintaining query performance for current data. You are not permitted to make changes to the client applications. You need to configure the Fact.Order table.
 Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Recreate the Fact.Order table on the partition scheme.

Execute an ALTER TABLE command to specify the partition function.

Create a partition scheme based on the partition function.

Execute an ALTER TABLE command to specify the partition scheme.

Recreate the Fact.Order table on the partition function.

Create a partition function.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Step 1: Create a partition function.

Using CREATE PARTITION FUNCTION is the first step in creating a partitioned table or index. Step 2: Create a partition scheme based on the partition function.

To migrate SQL Server partition definitions to SQL Data Warehouse simply: Step 3: Execute an ALTER TABLE command to specify the partition function.

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

NEW QUESTION 52

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

You are designing a data warehouse and the load process for the data warehouse.

You have a source system that contains two tables named Table1 and Table2. All the rows in each table have a corresponding row in the other table.

The primary key for Table1 is named Key1. The primary key for Table2 is named Key2.

You need to combine both tables into a single table named Table3 in the data warehouse. The solution must ensure that all the nonkey columns in Table1 and Table2 exist in Table3. Which component should you use to load the data to the data warehouse?

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

Answer: G

Explanation:

The Lookup transformation performs lookups by joining data in input columns with columns in a reference dataset. You use the lookup to access additional information in a related table that is based on values in common columns.

You can configure the Lookup transformation in the following ways: Specify joins between the input and the reference dataset.

Add columns from the reference dataset to the Lookup transformation output. Etc.

NEW QUESTION 53

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

You are implementing a Microsoft SQL Server data warehouse with a multi-dimensional data model. When testing a pilot version of the data warehouse, business users observe that the number of products in stock is inaccurate. The number of products in stock always increases and represents the total number of products that have ever been in stock.

You need to correct the existing model and ensure that it reflects the number of in-stock products. You must not change the overall structure of the data model. What should you do?

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

Answer: H

NEW QUESTION 57

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

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

You plan to deploy a Microsoft SQL server that will host a data warehouse named DB1. The server will contain four SATA drives configured as a RAID 10 array. You need to minimize write contention on the transaction log when data is being loaded to the database. Solution: You configure the server to automatically delete the transaction logs nightly.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

<https://www.red-gate.com/simple-talk/sql/database-administration/optimizing-transaction-log-throughput/> [https://docs.microsoft.com/en-us/sql/relational-](https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on-)
[databases/policy-based-management/place-data-and-log-files-on-](https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on-)

NEW QUESTION 58

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

What should you do?

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

Answer: B

Explanation:

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

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

NEW QUESTION 60

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

Answer: B

NEW QUESTION 64

You have a data warehouse named DW1.

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

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

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

How many partitions should you recommend?

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

Answer: B

NEW QUESTION 67

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

The model contains an entity named Product. The Product entity has three user-defined attributes named category, Subcategory, and Price, respectively. You need to ensure that combinations of values stored in the category and subcategory attributes are unique. What should you do?

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

Answer: D

NEW QUESTION 69

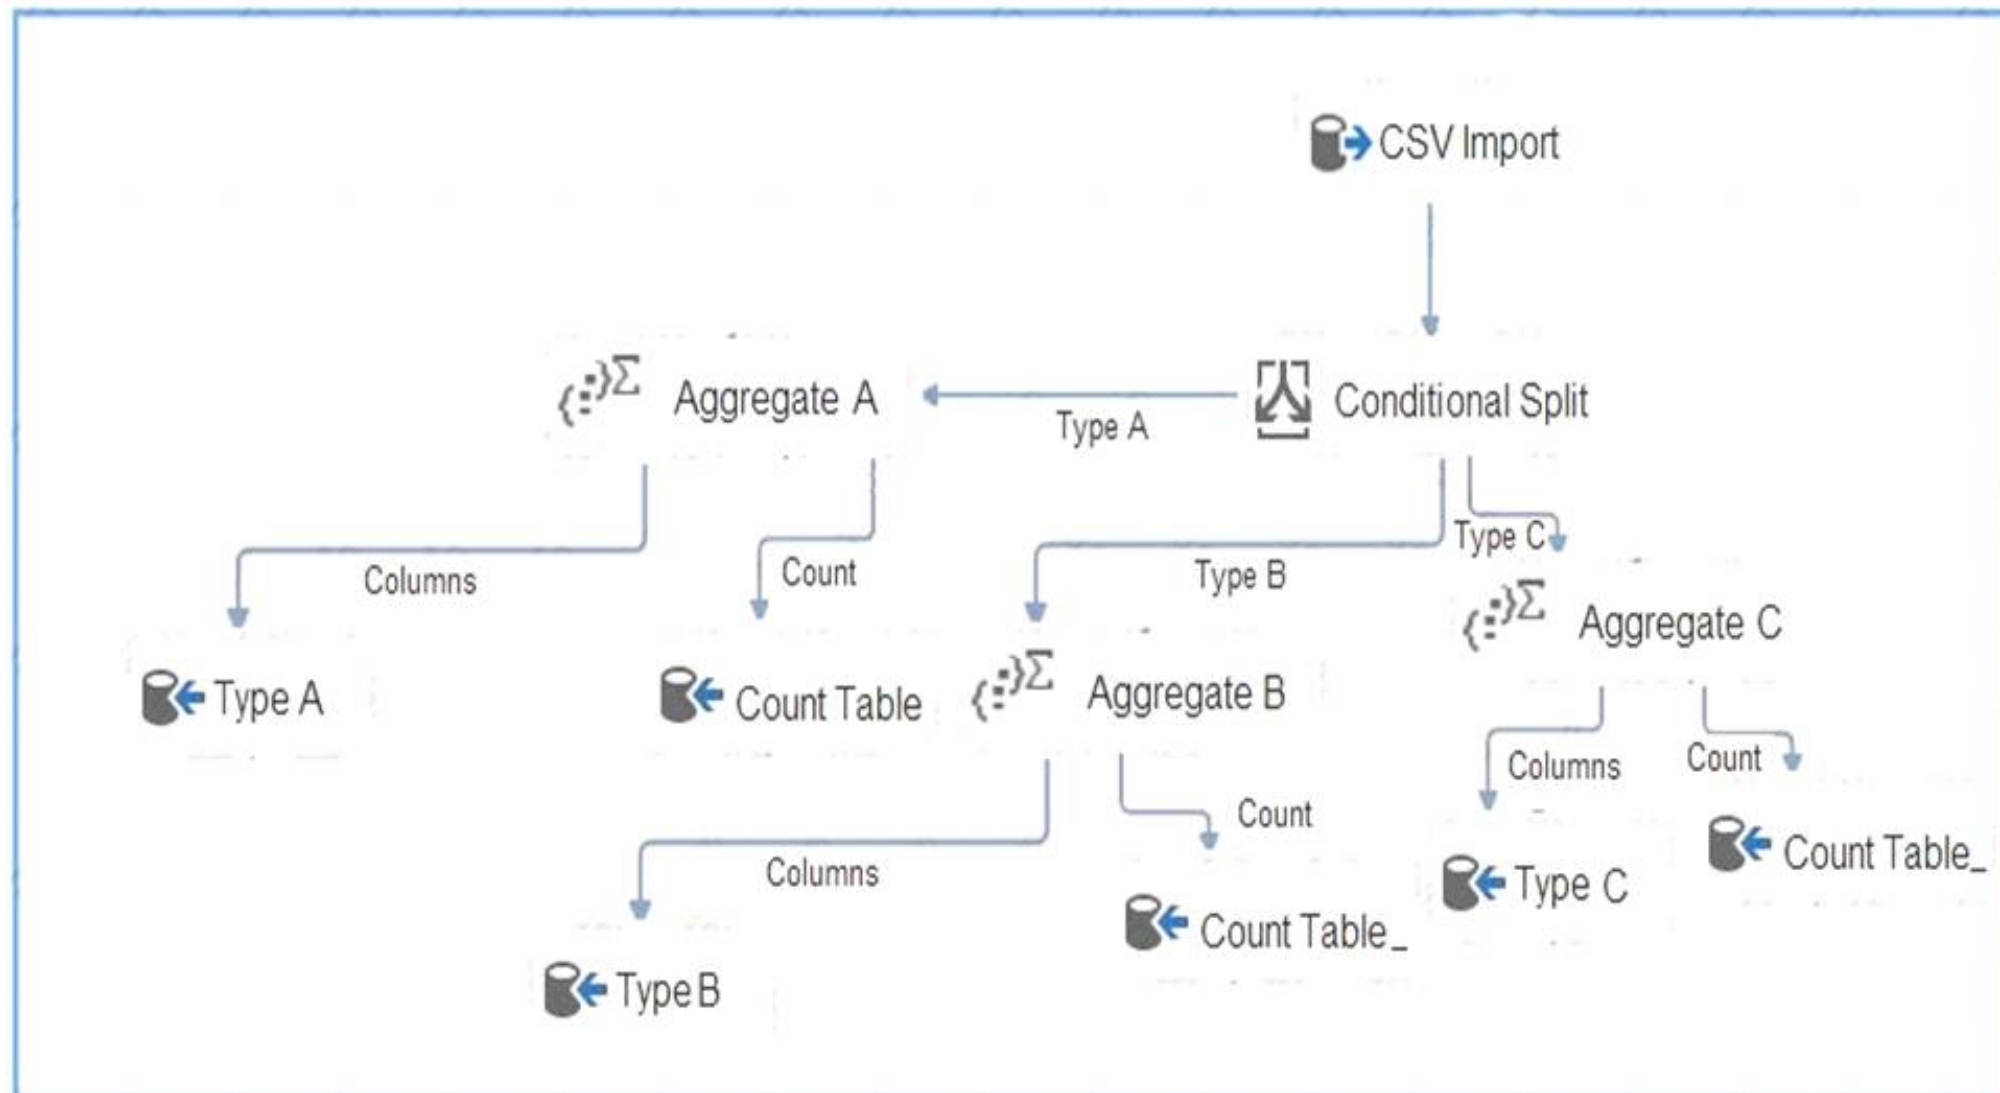
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

Each night you receive a comma separated values (CSV) file that contains different types of rows. Each row type has a different structure. Each row in the CSV file is unique. The first column in every row is named Type. This column identifies the data type.

For each data type, you need to load data from the CSV file to a target table. A separate table must contain the number of rows loaded for each data type.

Solution: You create a SQL Server Integration Services (SSIS) package as shown in the exhibit. (Click the Exhibit tab.)



Does the solution meet the goal?

- A. Yes
- B. NO

Answer: A

Explanation:

The conditional split is correctly placed before the count.

NEW QUESTION 71

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

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

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

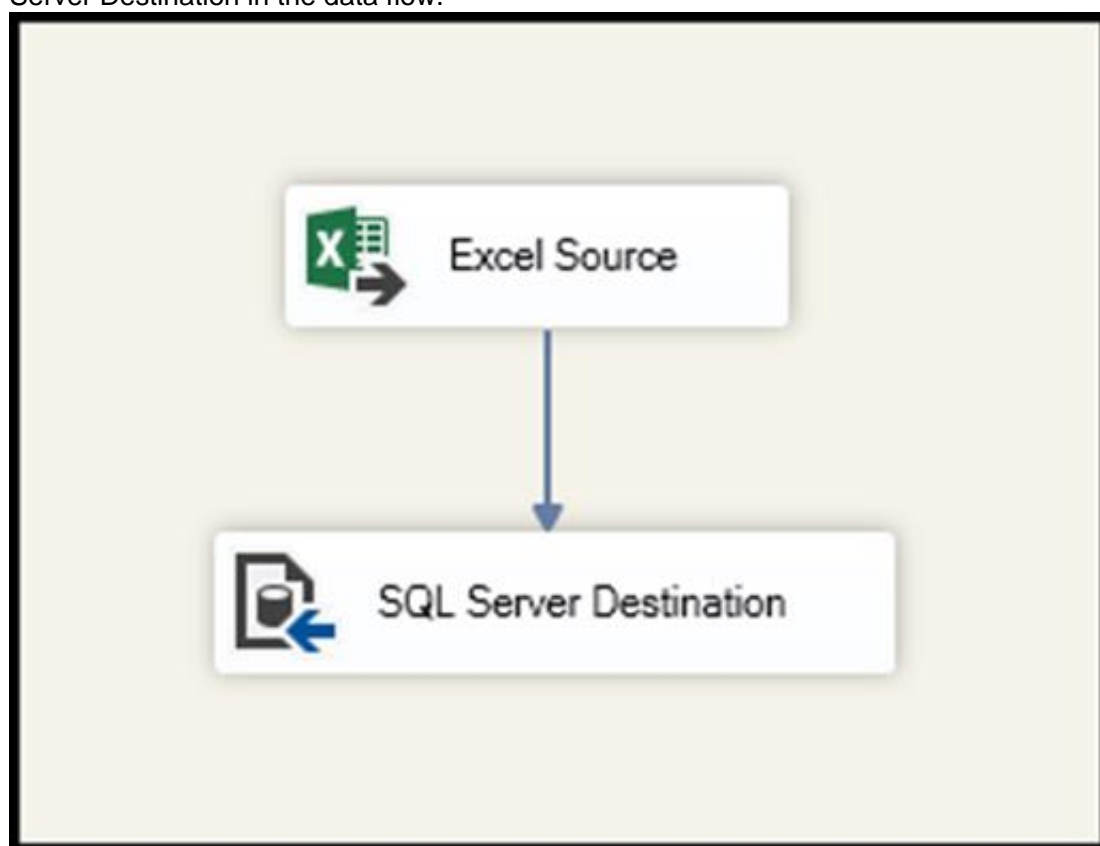
Each day, you publish a Microsoft Excel workbook that contains a list of product names and current prices to an external website. Suppliers update pricing information in the workbook. Each supplier saves the workbook with a unique name. Each night, the Products table is deleted and refreshed from MDS 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

Answer: A

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>

NEW QUESTION 72

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

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

You have the following line-of-business solutions:

- ▶ ERP system
- ▶ Online WebStore
- ▶ Partner extranet

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

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

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

- ▶ If a change is made to the ReferenceNr column in any of the sources, set the value of IsDisabled to True and create a new row in the Products table.
- ▶ If a row is deleted in any of the sources, set the value of IsDisabled to True in the data warehouse. Solution: Perform the following actions:
- ▶ Enable the Change Tracking feature for the Products table in the three source databases.

- ▶ Query the CHANGETABLE function from the sources for the deleted rows.
- ▶ Set the IsDisabled column to True on the data warehouse Products table for the listed rows. Does the solution meet the goal?

A. Yes
 B. No

Answer: B

Explanation:

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

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

NEW QUESTION 77

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

Start of repeated scenario

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

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

The data warehouse includes the following tables:

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

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

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

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

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

End of repeated scenario

You need to configure the fact. Transaction table.

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

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

Transact-SQL segments

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

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

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

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

Answer area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Transact-SQL segments

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

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

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

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

Answer area

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

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

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

NEW QUESTION 82

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

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

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

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

Transact-SQL segments

EmailAddress

EmailAddress_Import

NOT MATCHED BY SOURCE

NOT MATCHED BY TARGET

MATCHED

Answer area

```

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Box 1: EmailAddress

The EmailAddress table is the target. Box 2: EmailAddress_import

The EmailAddress_import table is the source. Box 3: NOT MATCHED BY TARGET

Box 4: NOT MATCHED BY SOURCE

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

NEW QUESTION 86

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

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

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

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

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

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

Answer: E

Explanation:

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

References:

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

NEW QUESTION 90

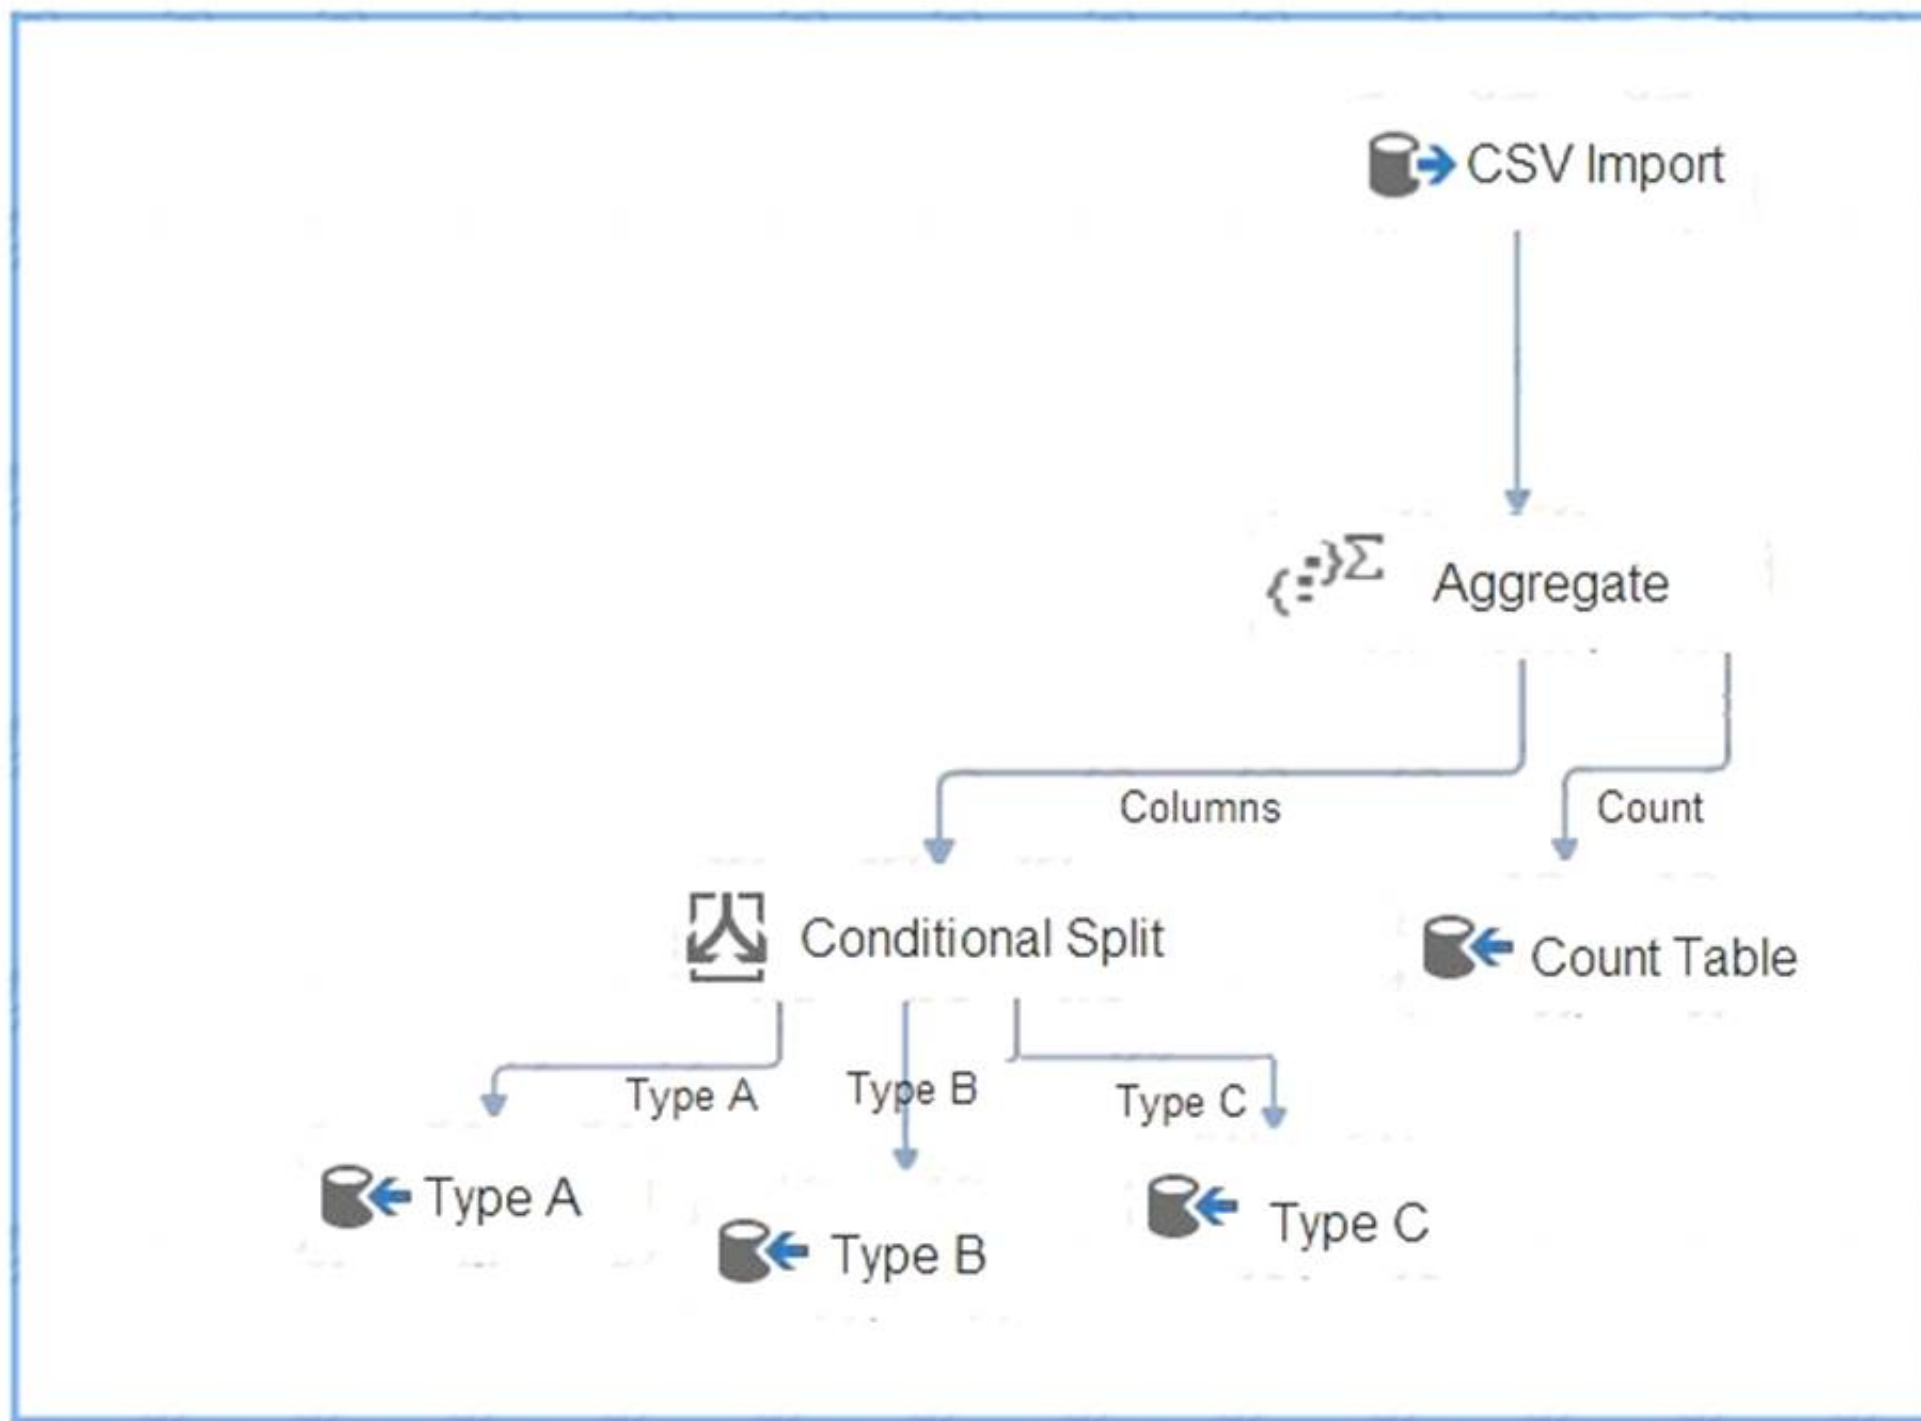
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

Each night you receive a comma separated values (CSV) file that contains different types of rows. Each row type has a different structure. Each row in the CSV file is unique. The first column in every row is named Type. This column identifies the data type.

For each data type, you need to load data from the CSV file to a target table. A separate table must contain the number of rows loaded for each data type.

Solution: You create a SQL Server Integration Services (SSIS) package as shown in the exhibit. (Click the Exhibit tab.)



Does the solution meet the goal?

- A. Yes
- B. NO

Answer: B

Explanation:

The conditional split must be before the count.

NEW QUESTION 94

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

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

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 95

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

You 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 re 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.addJob
- E. sys.sp.cdc.changejob
- F. sys.sp_cdc_disable_db
- G. sys.sp_cdc_enable_db
- H. sys.sp_cdc.stopJob

Answer: A

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-ssisd>

NEW QUESTION 97

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 101

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

Start of repeated scenario

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

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

and is considered historical.

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

End of repeated scenario

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

Which technology should you use for each table?

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

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

Table	Technology
Dimension.SalesTerritory	<div> <div>Change Data Capture (CDC)</div> <div>Change Tracking</div> <div>Temporal table</div> <div>Microsoft SQL Server snapshot replication</div> </div>
Dimension.Customer	<div> <div>Change Data Capture (CDC)</div> <div>Change Tracking</div> <div>Temporal table</div> <div>Microsoft SQL Server snapshot replication</div> </div>
Dimension.Date	<div> <div>Change Data Capture (CDC)</div> <div>Change Tracking</div> <div>Temporal table</div> <div>Microsoft SQL Server snapshot replication</div> </div>

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Box 1: Temporal table Box 2: Temporal table

Compared to CDC, Temporal tables are more efficient in storing historical data as it ignores insert actions. Box 3: Change Data Capture (CDC)

By using change data capture, you can track changes that have occurred over time to your table. This kind of functionality is useful for applications, like a data warehouse load process that need to identify changes, so they can correctly apply updates to track historical changes over time.

CDC is good for maintaining slowly changing dimensions.

Scenario: Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.

The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. References:

<https://www.mssqltips.com/sqlservertip/5212/sql-server-temporal-tables-vs-change-data-capture-vs-change-trac> <https://docs.microsoft.com/en-us/sql/relational-databases/tables/temporal-table-usage-scenarios?view=sql-server>

NEW QUESTION 104

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

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

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

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

Does the solution meet the goal?

- A. Yes
 B. No

Answer: B

NEW QUESTION 109

You need to recommend a storage solution for a data warehouse that minimizes load times. The solution must provide availability if a hard disk fails.

Which RAID configuration should you recommend for each type of database file? To answer, drag the appropriate RAID configurations to the correct database file types. Each RAID configuration 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.

RAID Configurations	Answer Area
<div>RAID 0</div> <div>RAID 5</div> <div>RAID 6</div> <div>RAID 10</div>	<div>Data file: RAID configuration</div> <div>Transaction log: RAID configuration</div>

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Box 1: RAID 5

RAID 5 is the similar to that of RAID 0 provided that the number of disks is the same. However, due to the fact that it is useless to read the parity data, the read speed is just (N-1) times faster but not N times as in RAID 0.

Box 2: RAID 10

Always place log files on RAID 1+0 (or RAID 1) disks. This provides better protection from hardware failure, and better write performance.

Note: In general RAID 1+0 will provide better throughput for write-intensive applications. The amount of performance gained will vary based on the HW vendor's

RAID implementations. Most common alternative to RAID 1+0 is RAID 5. Generally, RAID 1+0 provides better write performance than any other RAID level providing data protection, including RAID 5.

NEW QUESTION 111

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

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

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

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

Answer: E

Explanation:

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

NEW QUESTION 115

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 70-767 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 70-767 Product From:

<https://www.2passeasy.com/dumps/70-767/>

Money Back Guarantee

70-767 Practice Exam Features:

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