

70-767 Dumps

Implementing a SQL Data Warehouse (beta)

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



NEW QUESTION 1

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.

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.	

- A. Mastered
- B. Not Mastered

Answer: A

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>

NEW QUESTION 2

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

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

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

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

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

Answer Area

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

▼

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

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

```
ADD CONSTRAINT PK_Dimension_Employee
```

▼

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Answer Area

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

▼

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

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

```
ADD CONSTRAINT PK_Dimension_Employee
```

▼

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

NEW QUESTION 3

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

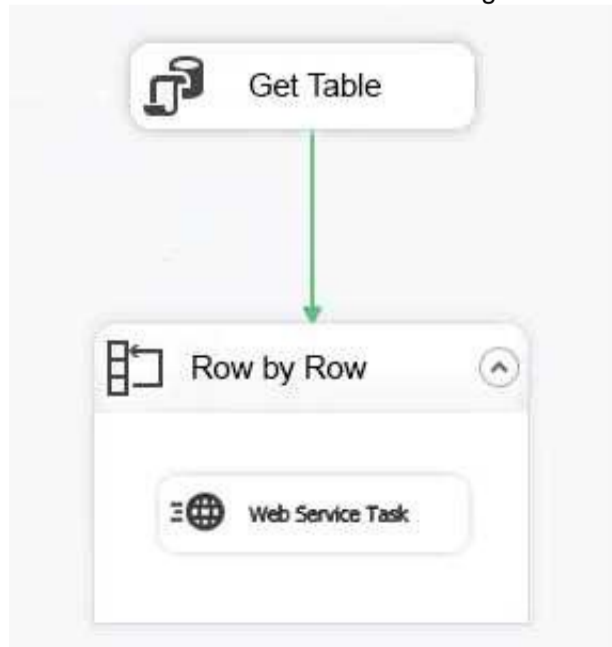
LastKey = ?
LastKey = @A
LastKey = @B
LastKey = @LastKey

WHERE

TableName = ?
TableName = @A
TableName = @B
TableName = @TableName

NEW QUESTION 4

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 Azure Blob Enumerator
- C. Foreach NodeList Enumerator
- D. Foreach ADO Enumerator

Answer: D

Explanation:

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

NEW QUESTION 5

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

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

You 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 6

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

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

- ▶ 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 data loading for the Dimension.Customer 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.

NOTE: You will not need all of the Transact-SQL segments.

Transact-SQL segments

```
EXEC sys.sp_cdc_enable_table
@source_schema = N 'schema',
@source_name   = N 'Dimension.Customer',
@role_name     = NULL,
@supports_net_changes = 1
```

```
EXEC sys.sp_cdc_enable_db
```

```
USE DB2
```

```
EXEC sys.sp_cdc_enable_table
```

```
USE DB1
```

```
EXEC sys.sp_cdc_enable_db
@source_schema = N 'schema',
@source_name   = N 'Dimension.Customer',
@role_name     = NULL,
@supports_net_changes = 1
```

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: USE DB1

From Scenario: 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.

Step 2: EXEC sys.sp_cdc_enable_db

Before you can enable a table for change data capture, the database must be enabled. To enable the database, use the sys.sp_cdc_enable_db stored procedure.

sys.sp_cdc_enable_db has no parameters. Step 3: EXEC sys.sp_cdc_enable_table

@source_schema = N 'schema' etc.

sys.sp_cdc_enable_table enables change data capture for the specified source table in the current database. Partial syntax:

sys.sp_cdc_enable_table

[@source_schema =] 'source_schema',

[@source_name =] 'source_name' , [, [@capture_instance =] 'capture_instance'] [, [@supports_net_changes =] supports_net_changes]

Etc.

References:

<https://docs.microsoft.com/en-us/sql/relational-databases/system-stored-procedures/sys-sp-cdc-enable-table-trans>

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

NEW QUESTION 7

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.

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 alternative values for entries that have been indentified as being incorrectly entered.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

NEW QUESTION 8

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 FoctAddress 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 9

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

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

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

You need to implement the data partitioning strategy. How should you partition the Fact.Order table?

- A. Create 17,520 partitions.
- B. Use a granularity of two days.
- C. Create 2,557 partitions.
- D. Create 730 partitions.

Answer: C

Explanation:

We create one partition for each day. 7 years times 365 days is 2,555. Make that 2,557 to provide for leap years.

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

NEW QUESTION 10

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 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.
- ▶ 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 data loading for the tables.

Which data loading technology should you use for each table? To answer, select the appropriate options in the answer area.

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

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

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

Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables. Box 1: Change Tracking

Box 2: Change Tracking Box 3: Temporal Table

Temporal Tables are generally useful in scenarios that require tracking history of data changes.

We recommend you to consider Temporal Tables in the following use cases for major productivity benefits.

* Slowly-Changing Dimensions

Dimensions in data warehousing typically contain relatively static data about entities such as geographical locations, customers, or products.

References:

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

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 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.

Columns
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)

```
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 measure for the Fact.Order table. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

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-serv>

NEW QUESTION 17

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

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

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

You need to pause the compute resources when the instance is not being used. Solution: You use the Azure portal.

Does the solution meet the goal?

- A. Yes
- B. No

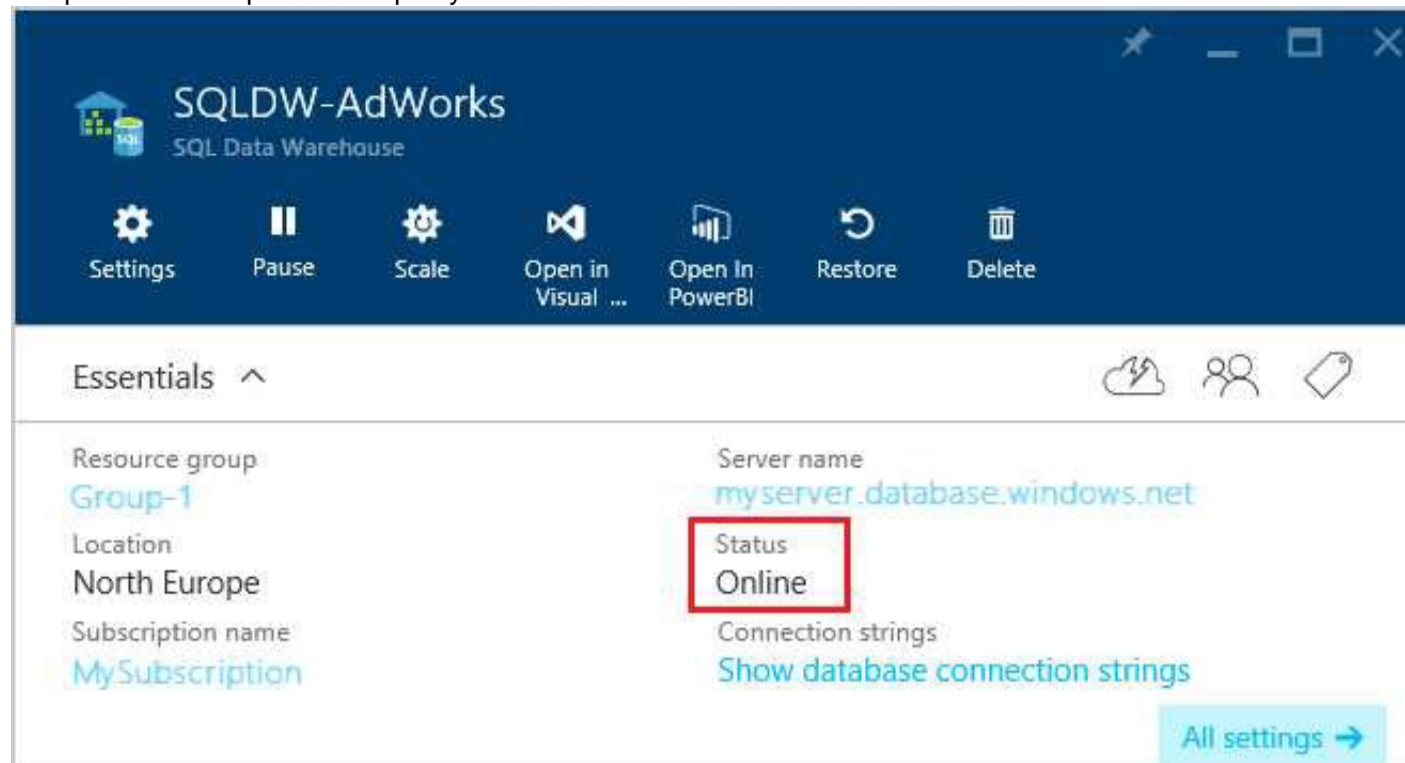
Answer: A

Explanation:

To pause a SQL Data Warehouse database, use any of these individual methods. Pause compute with Azure portal

Pause compute with PowerShell Pause compute with REST APIs Note: To pause a database:

1. Open the Azure portal and open your database. Notice that the Status is Online.



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

References:

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

NEW QUESTION 22

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 24

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 29

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 30

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

After you answer a question in this 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 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 37

You are developing a data warehouse. You run the following Transact-SQL statement:

```
USE AdventureWorks
GO
CREATE TABLE Production.TransactionHistoryArchive(
    TransactionID INT IDENTITY (1, 1) NOT NULL,
    CONSTRAINT PK_TransactionHistoryArchive_TransactionID PRIMARY KEY CLUSTERED (TransactionID)
)
```

Use the drop-down menus to select the answer choice that answers each question based on the information presented in the graphic.

NOTE: Each correct selection is worth one point.

What is the name of the table created?

▼
AdventureWorks
Production
TransactionHistoryArchive

What is the name of the primary key?

▼
Identity
Production
TransactionID

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

What is the name of the table created?

▼
AdventureWorks
Production
TransactionHistoryArchive

What is the name of the primary key?

▼
Identity
Production
TransactionID

NEW QUESTION 40

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.

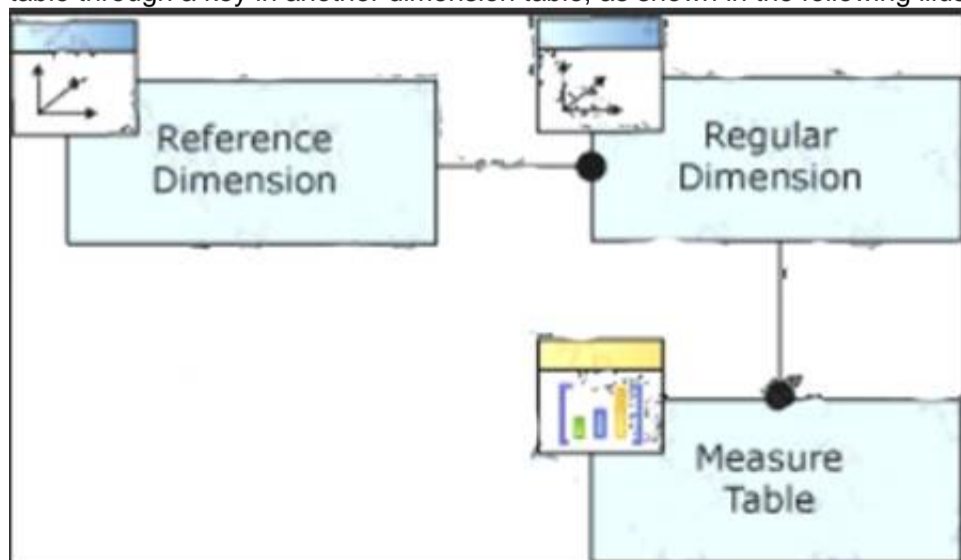
You need to recommend an appropriate dimension usage relationship. What should you recommend?

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

Answer: B

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.



NEW QUESTION 43

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it As a result, these questions will not appear in the review screen.
You 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 enable package logging within SSIS.
Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 45

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 48

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

	Transaction Key	Date Key	Customer Key	Bill To Customer Key	Supplier Key	Transaction Type Key	Payment Method Key	WWI Invoice ID
1	7	2013-01-01	375	202	0	1	0	7
2	11	2013-01-01	387	202	0	1	0	11
3	12	2013-01-01	330	202	0	1	0	12
4	13	2013-01-01	274	202	0	1	0	13
5	16	2013-01-01	215	202	0	1	0	16
6	25	2013-01-01	298	202	0	1	0	25
7	26	2013-01-01	285	202	0	1	0	26
8	30	2013-01-01	368	202	0	1	0	30
9	35	2013-01-01	232	202	0	1	0	35
10	39	2013-01-01	346	202	0	1	0	39
11	41	2013-01-01	216	202	0	1	0	41
12	63	2013-01-02	224	202	0	1	0	42
13	64	2013-01-02	264	202	0	1	0	43
14	65	2013-01-02	268	202	0	1	0	44
15	70	2013-01-02	376	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 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.

Answer area

```
CREATE CLUSTERED COLUMNSTORE INDEX idx_fact_sale ON fact.Sale
CREATE NONCLUSTERED COLUMNSTORE INDEX idx_fact_sale ON fact.Sale
ALTER INDEX idx_fact_sale ON fact.Sale DISABLE
```

```
WITH (DROP_EXISTING = ON)
DROP INDEX idx_fact_sale ON fact.Sale
ALTER INDEX idx_fact_sale ON fact.Sale REBUILD
CREATE CLUSTERED COLUMNSTORE INDEX idx_fact_sale_cs ON fact.Sale
```

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer area

```
CREATE CLUSTERED COLUMNSTORE INDEX idx_fact_sale ON fact.Sale
CREATE NONCLUSTERED COLUMNSTORE INDEX idx_fact_sale ON fact.Sale
ALTER INDEX idx_fact_sale ON fact.Sale DISABLE
```

```
WITH (DROP_EXISTING = ON)
DROP INDEX idx_fact_sale ON fact.Sale
ALTER INDEX idx_fact_sale ON fact.Sale REBUILD
CREATE CLUSTERED COLUMNSTORE INDEX idx_fact_sale_cs ON fact.Sale
```

NEW QUESTION 52

You create a Master Data Services (MDS) model that manages the master data for a Product dimension. The Product dimension has the following properties: All the members of the Product dimension have a product type, a product subtype, and a unique product name. Each product has a single product type and a single product subtype. The product type has a one-to-many relationship to the product subtype. You need to ensure that the relationship between the product name, the product type, and the product subtype is maintained when products are added to or updates in the database. What should you add to the model?

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

Answer: B

Explanation:

A Master Data Services derived hierarchy is derived from the domain-based attribute relationships that already exist between entities in a model. You can create a derived hierarchy to highlight any of the existing domain-based attribute relationships in the model.

NEW QUESTION 55

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.

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"/>

- A. Mastered
- B. Not Mastered

Answer: A

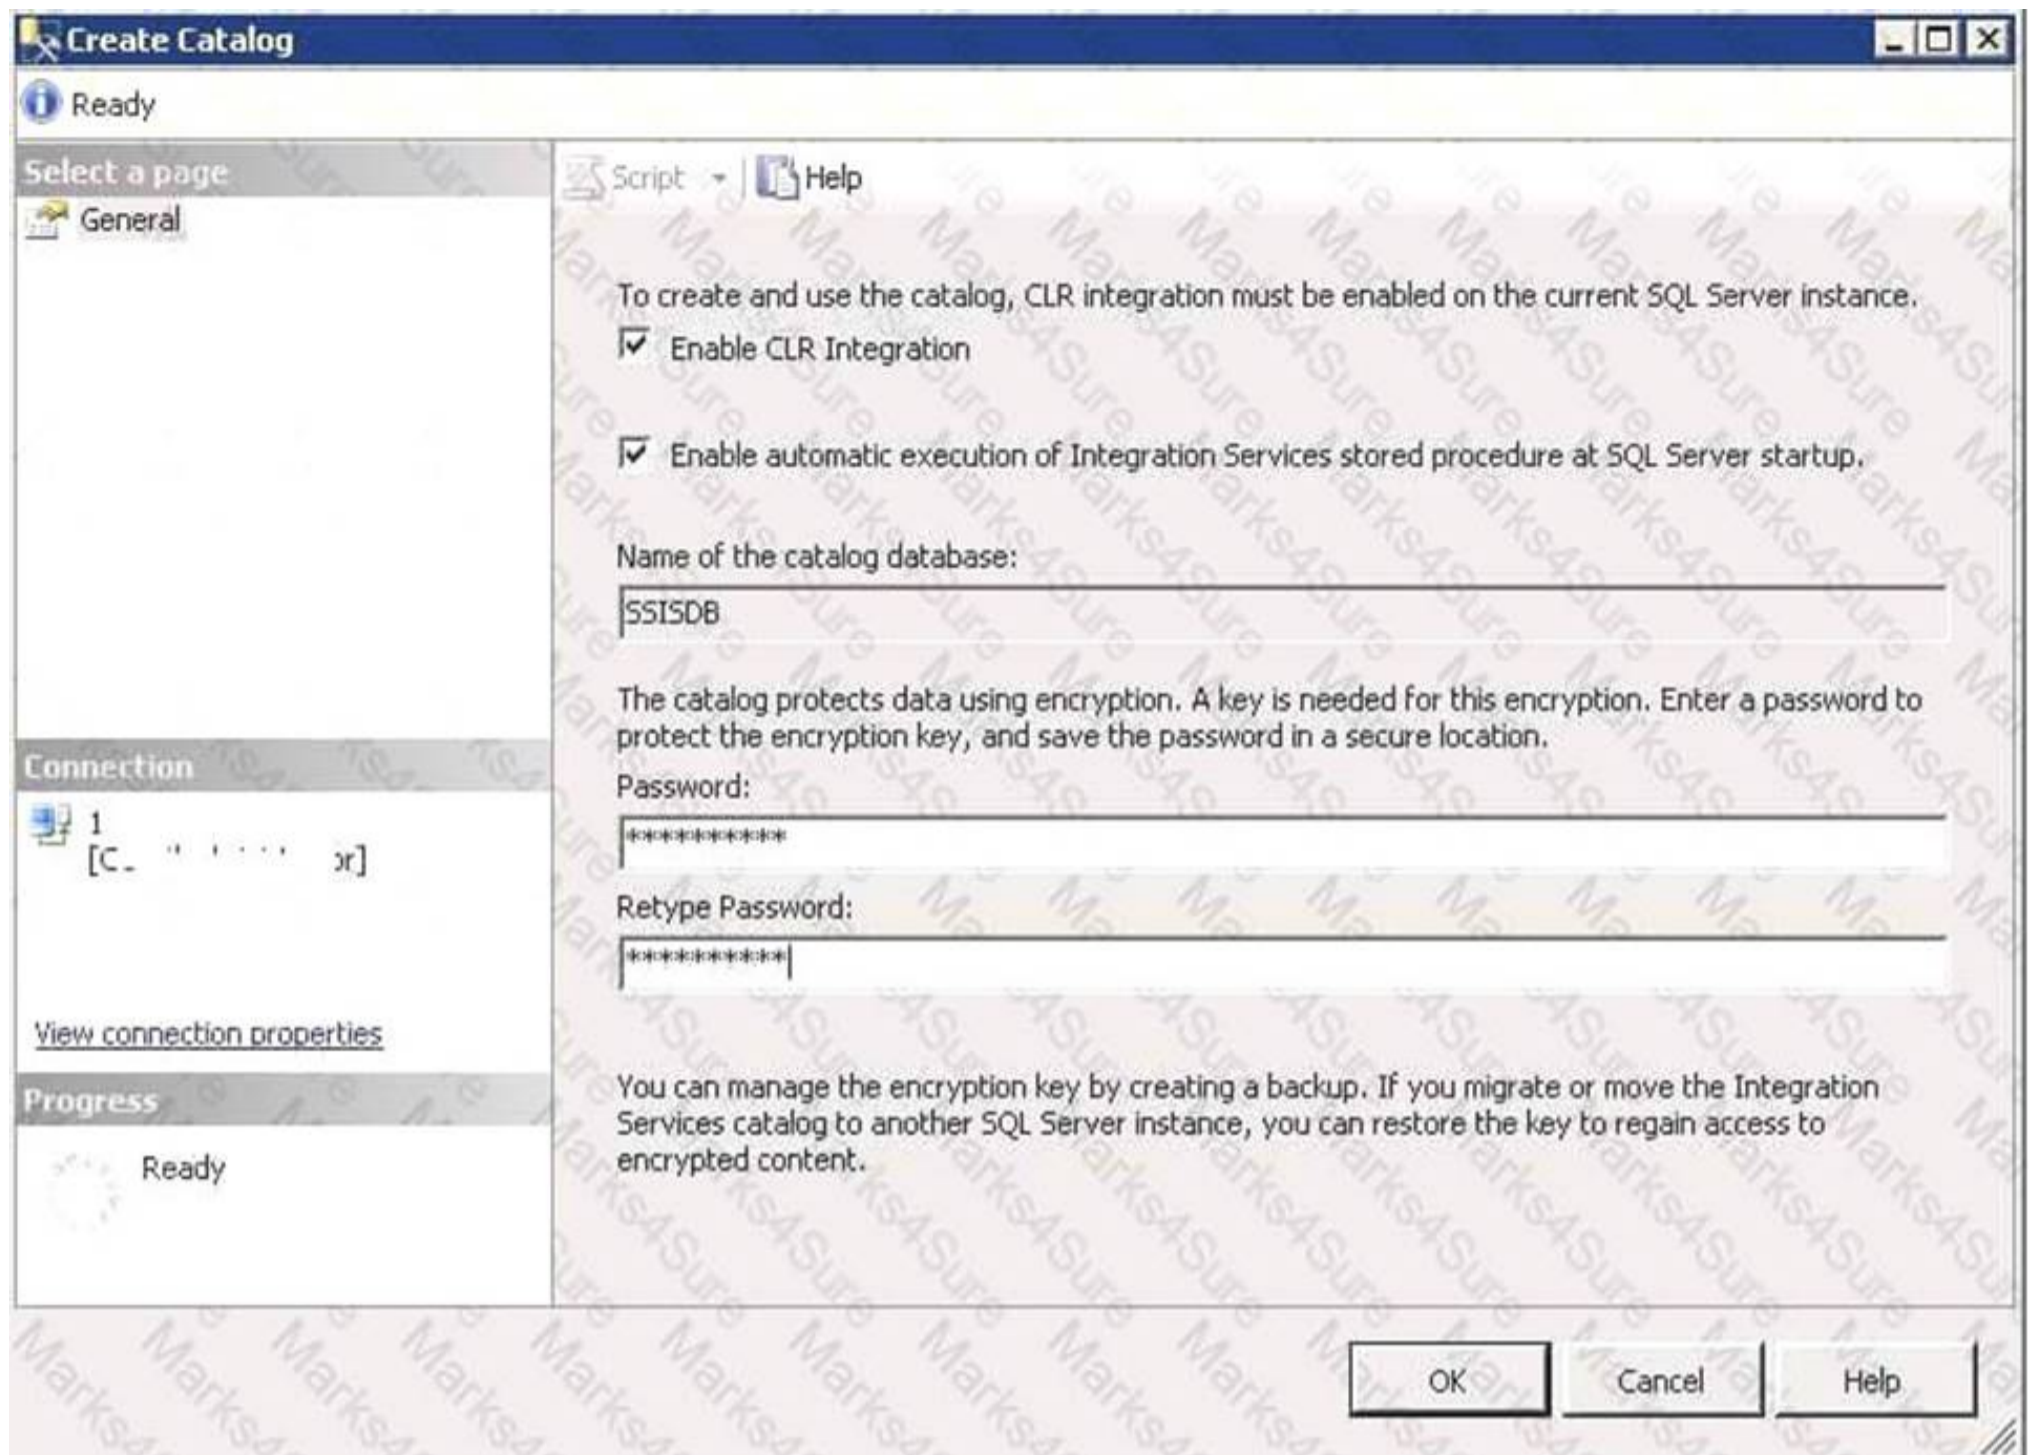
Explanation:

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-crea>

NEW QUESTION 60

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 64

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?

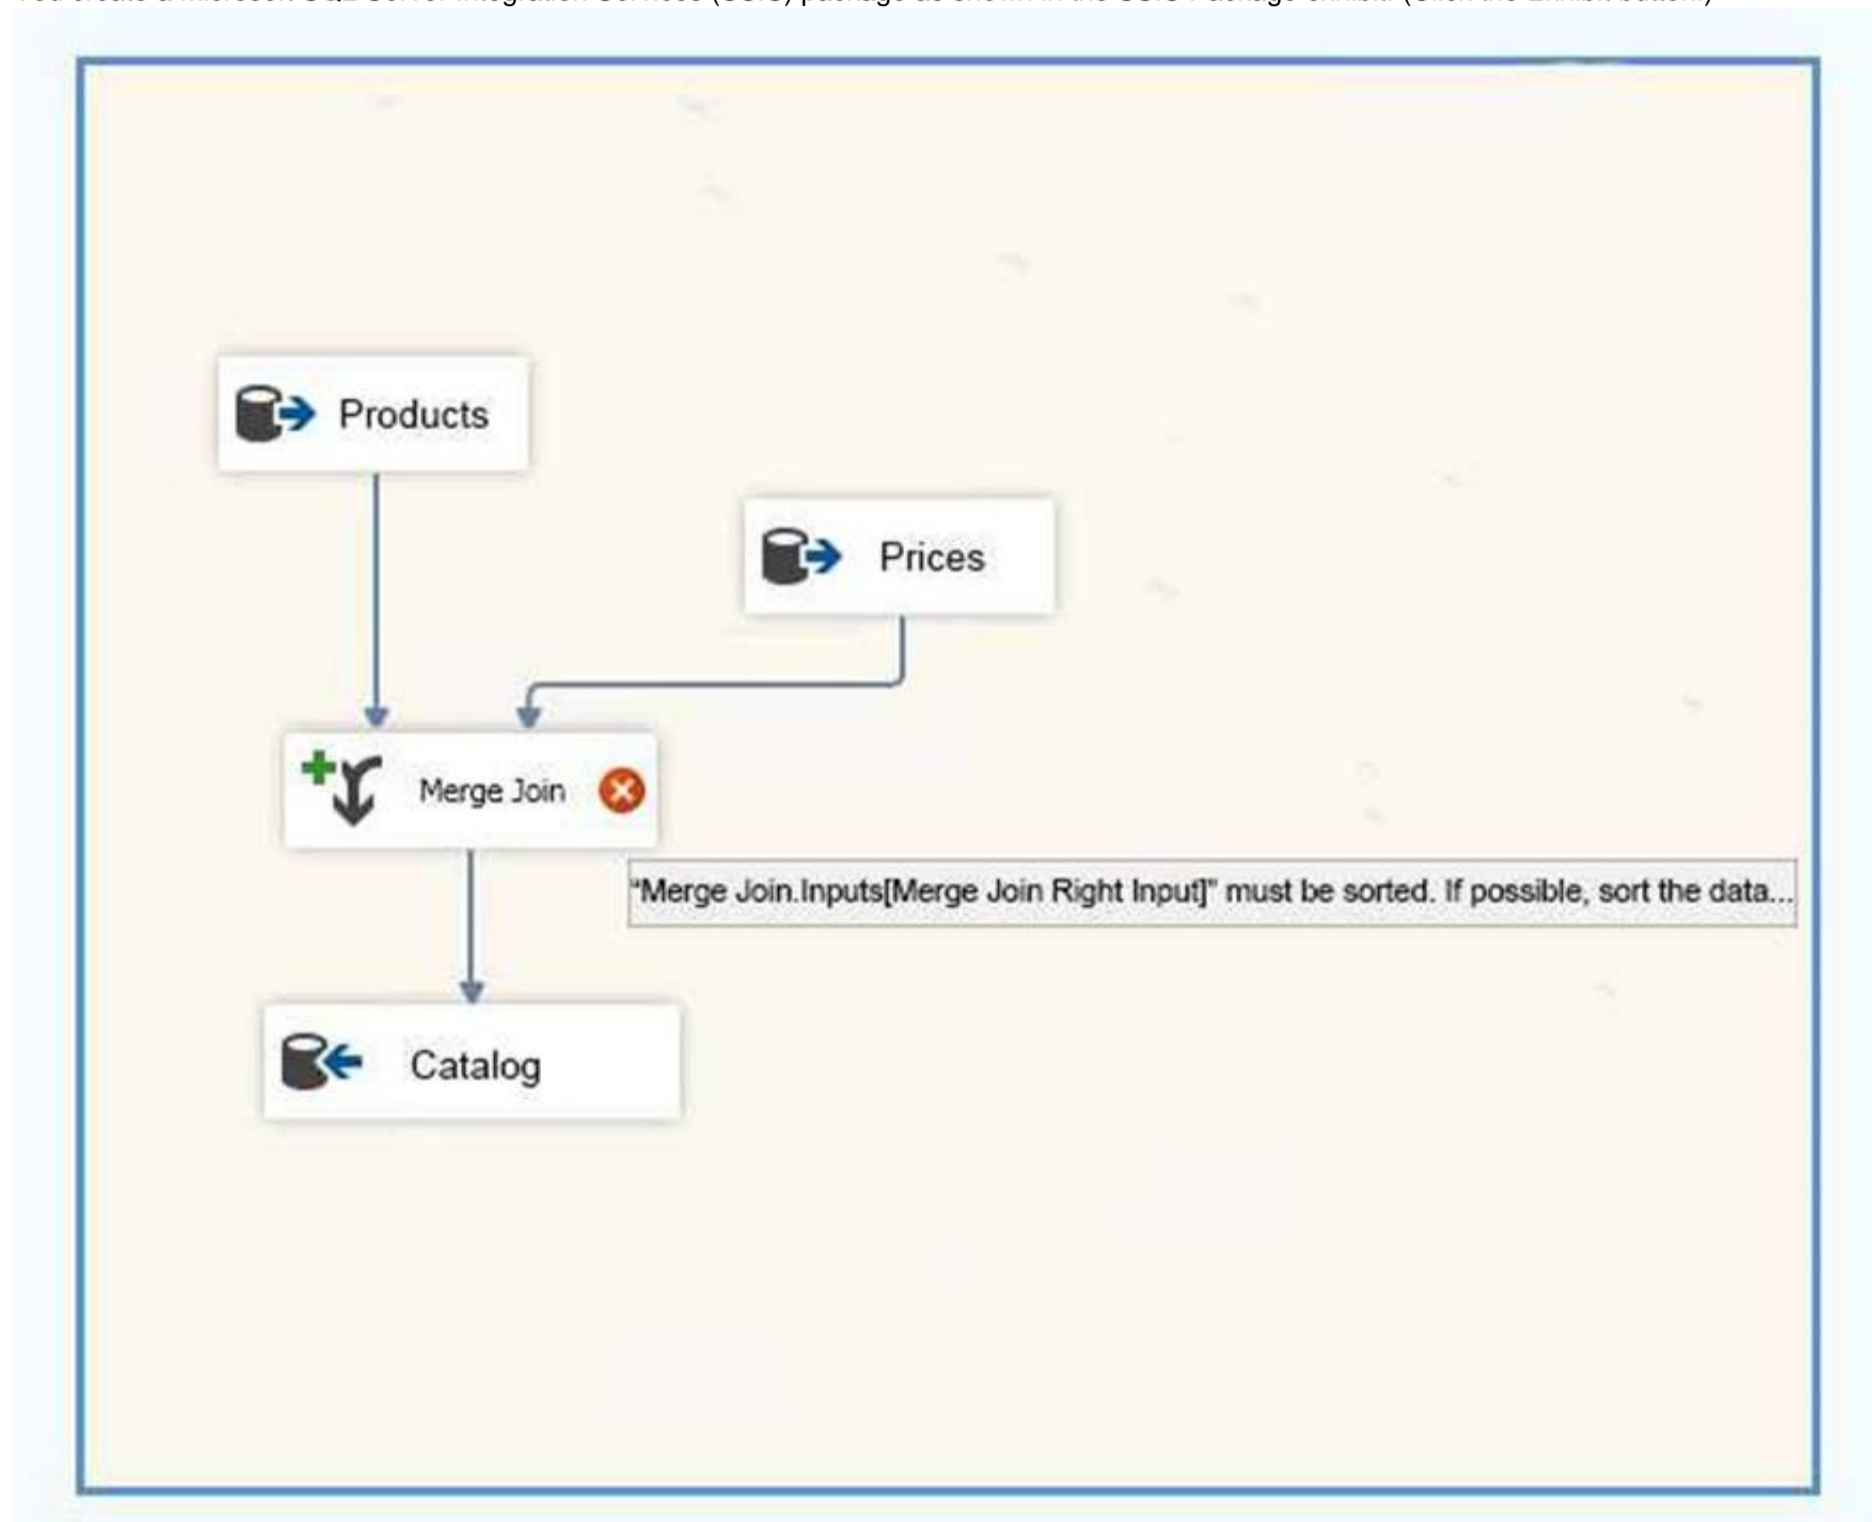
- A. star schema

- 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 67

You create a Microsoft SQL Server Integration Services (SSIS) package as shown in the SSIS Package exhibit. (Click the Exhibit button.)



The package uses data from the Products table and the Prices table. Properties of the Prices source are shown in the OLE DB Source Editor exhibit (Click the Exhibit Button.) and the Advanced Editor for Prices exhibit (Click the Exhibit button.)

OLE DB Source Editor

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

Connection Manager
Columns
Error Output

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

OLE DB connection manager:
Catalog New...

Data access mode:
SQL command

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

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

Preview...

OK Cancel Help

Advanced Editor for OLE DB Source

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

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

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

Inputs and outputs:

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

Add Output Add Column
Remove Output Remove Column

Common Properties

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

Data Type Properties

CodePage	1252
DataType	string [DT_STR]
Length	100
Precision	0

ID

Refresh OK Cancel Help

You join the Products and Prices tables by using the ReferenceNr column. You need to resolve the error with the package.

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

	Yes	No
You can resolve the error by adding a Sort transform between the OLE DB source and the Merge Join transform.	<input type="radio"/>	<input type="radio"/>
You can resolve the error by changing the SortKeyPosition setting for the ReferenceNr column and the value of the IsSorted setting for the OLE DB Source Output.	<input type="radio"/>	<input type="radio"/>
You can resolve the error by adding an Aggregate transform between the OLE DB source and the Merge Join transform.	<input type="radio"/>	<input type="radio"/>
You can resolve the error by replacing the Merge Join transform with a Lookup transform.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

There are two important sort properties that must be set for the source or upstream transformation that supplies data to the Merge and Merge Join transformations: The Merge Join Transformation requires sorted data for its inputs.

If you do not use a Sort transformation to sort the data, you must set these sort properties manually on the source or the upstream transformation.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/sort-data-for-the-merge-and->

NEW QUESTION 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 sections, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.

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

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

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

You need to optimize performance.

Solution: You run the following Transact-SQL statement:

```
CREATE STATISTICS Stat1
On Sales.SalesOrderHeader (OrderDate)
WITH SAMPLE 100 ROWS
```

Does the solution meet the goal?

- A. Yes
B. No

Answer: B

Explanation:

100 out of 500,000 rows is a too small sample size.

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

NEW QUESTION 71

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

You create a matching policy that contains one matching rule.

You need to configure the Similarity of Similar percentage that defines a match. Which similarity percentage will always generate a similarity score of 0?

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

Answer: A

Explanation:

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

References:

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

NEW QUESTION 74

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

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

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

Product prices are updated and are stored in a table named Products on DB1. The Products table is deleted and refreshed each night from MDS by using a Microsoft SQL Server Integration Services (SSIS) package. None of the data sources are sorted.

You need to update the SSIS package to add current prices to the Products table. What should you use?

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

Answer: D

Explanation:

In the current release of SQL Server Integration Services, the SQL statement in an Execute SQL task can contain a MERGE statement. This MERGE statement enables you to accomplish multiple INSERT, UPDATE, and DELETE operations in a single statement.

References:

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

NEW QUESTION 79

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

Start of repeated scenario

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

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

The data warehouse includes the following tables:

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

You must implement a partitioning scheme for the fact. Transaction table to move older data to less expensive storage. Each partition will store data for a single

calendar year, as shown in the exhibit (Click the Exhibit button.) You must align the partitions.

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

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

End of repeated scenario

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

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

Answer area

CREATE PARTITION FUNCTION
CREATE PARTITION SCHEME
CREATE PROCEDURE
CREATE RESOURCE POOL

[DateRange] (DATETIME) AS RANGE RIGHT
[DateRange] (DATETIME) AS RANGE LEFT
[DateRange] (INT) AS RANGE RIGHT
[DateRange] (INT) AS RANGE LEFT

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

CREATE PROCEDURE
[DateRange] (INT) AS RANGE LEFT

NEW QUESTION 80

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 84

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

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

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

Messages

Msg 2627, Level 14, State 1, Line 13
Violation of UNIQUE Key constraint 'UQ_DimEmplo_8549FE539cf2eca'. Cannot insert duplicate ke
object 'dbo.DimEmployee'. The duplicate key value is (<NULL>).

The statement has been terminated.

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

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

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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

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

NEW QUESTION 85

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).

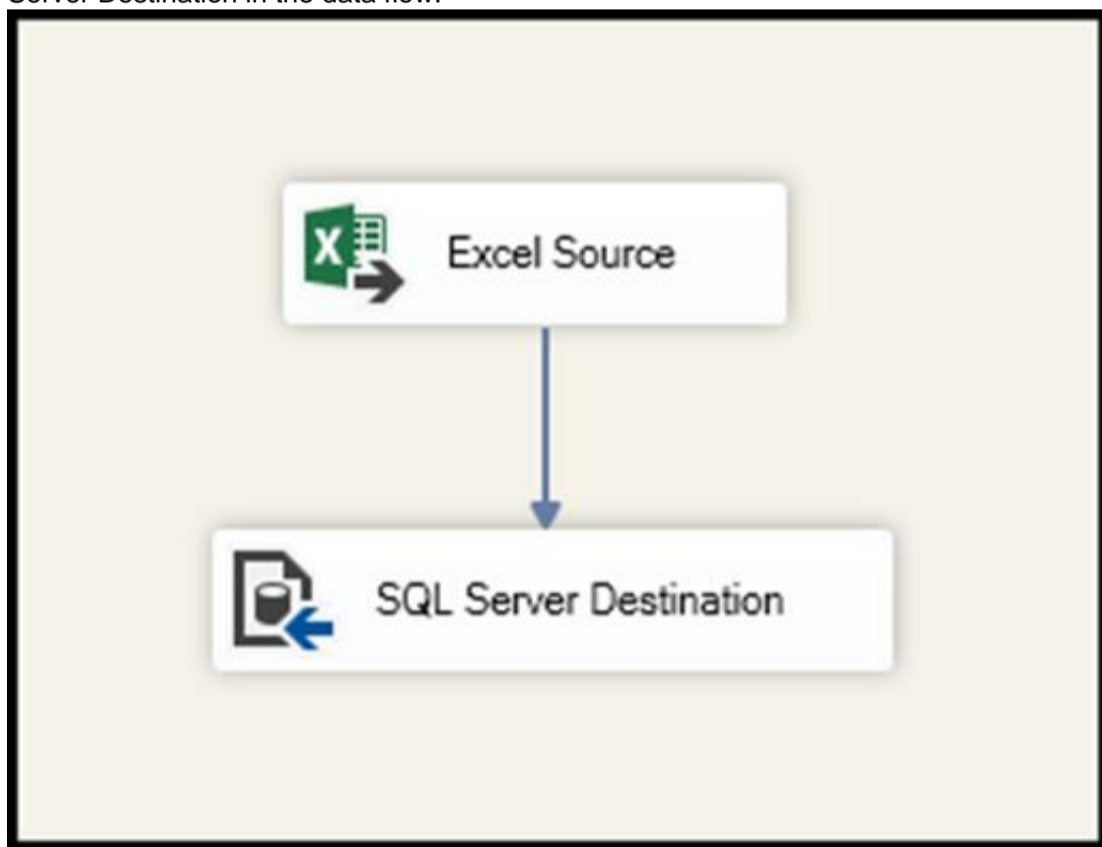
Each day, you publish a Microsoft Excel workbook that contains a list of product names and current prices to an external website. Suppliers update pricing information in the workbook. Each supplier saves the workbook with a unique name. Each night, the Products table is deleted and refreshed from MDS 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 89

You have a database that includes a table named dbo.sales. The table contains two billion rows. You created the table by running the following Transact-SQL statement:

```

CREATE TABLE dbo.Sales (
    SaleId BIGINT PRIMARY KEY,
    StoreId INT,
    EmployeeId INT,
    SaleAmount MONEY,
    TaxAmount MONEY,
    SubTotalAmount MONEY,
    LineItems XML,
    Refund BIT,
    SaleDate DATE,
    SaleTime TIME
)
  
```

You run the following queries against the dbo.sales table. All of the queries perform poorly.

Query name	Query text
Query1	<pre>SELECT StoreId, SUM(SaleAmount) SaleTotal, SUM(TaxAmount) TaxTotal FROM dbo.Sales WHERE SaleDate BETWEEN '1/1/2015' AND '1/1/2016' GROUP BY StoreId</pre>
Query2	<pre>SELECT StoreId, datepart(hh, SaleTime) SaleHour, count(*) FROM dbo.Sales WHERE SaleDate = convert(varchar(10), getdate()-1, 111) GROUP BY StoreId, datepart(hh, SaleTime)</pre>
Query3	<pre>SELECT SaleId, StoreId, EmployeeId, SaleAmount FROM dbo.Sales WHERE Refund = 1 AND SaleDate = convert(varchar(10), getdate()-1, 111)</pre>

The ETL process that populates the table uses bulk insert to load 10 million rows each day. The process currently takes six hours to load the records. The value of the Refund column is equal to 1 for only 0.01 percent of the rows in the table. For all other rows, the value of the Refund column is equal to 0. You need to maximize the performance of queries and the ETL process. Which index type should you use for each query? To answer, select the appropriate index types in the answer area. NOTE: Each correct selection is worth one point.

Query name	Index type
Query1	<input type="checkbox"/> Clustered ColumnStore Index <input type="checkbox"/> Clustered Index <input type="checkbox"/> Nonclustered Index <input type="checkbox"/> Filtered nonclustered Index
Query2	<input type="checkbox"/> Clustered ColumnStore Index <input type="checkbox"/> Clustered Index <input type="checkbox"/> Nonclustered Index <input type="checkbox"/> Filtered nonclustered Index
Query3	<input type="checkbox"/> Clustered ColumnStore Index <input type="checkbox"/> Clustered Index <input type="checkbox"/> Nonclustered Index <input type="checkbox"/> Filtered nonclustered Index

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Query name	Index type
Query1	<input checked="" type="checkbox"/> Clustered ColumnStore Index <input type="checkbox"/> Clustered Index <input type="checkbox"/> Nonclustered Index <input type="checkbox"/> Filtered nonclustered Index
Query2	<input checked="" type="checkbox"/> Clustered ColumnStore Index <input type="checkbox"/> Clustered Index <input type="checkbox"/> Nonclustered Index <input type="checkbox"/> Filtered nonclustered Index
Query3	<input checked="" type="checkbox"/> Clustered ColumnStore Index <input type="checkbox"/> Clustered Index <input type="checkbox"/> Nonclustered Index <input type="checkbox"/> Filtered nonclustered Index

NEW QUESTION 92

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have a Microsoft SQL server that has Data Quality Services (DQS) installed. You need to review the completeness and the uniqueness of the data stored in the matching policy. Solution: You profile the data. Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

Use a matching rule. References:

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

NEW QUESTION 95

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	Answer area
EmailAddress	MERGE Transact-SQL segment AS B
EmailAddress_Import	USING Transact-SQL segment AS A
NOT MATCHED BY SOURCE	ON A.EmailAddressValue = B.EmailAddressValue
NOT MATCHED BY TARGET	WHEN Transact-SQL segment
MATCHED	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 98

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 99

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

Start of repeated scenario

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

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

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

End of repeated scenario

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

Which technology should you use for each table?

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

Answer area

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Temporal table Box 2: Temporal table

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

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

CDC is good for maintaining slowly changing dimensions.

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

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

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

NEW QUESTION 104

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

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

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

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

You need to modify Dim1 to ensure that the table can be used for both fact tables.
Which two actions should you perform? Each correct answer presents part of the solution.
NOTE: Each correct selection is worth one point.

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

Answer: AD

NEW QUESTION 106

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 configuring a Microsoft SQL server named ow1 for a new data warehouse. The server contains eight drives and eight processor cores. Each drive uses a separate physical disk.
You need to configure storage for the tempdb database. The solution must minimize the amount of time it takes to process daily ETL jobs.
Solution: You configure eight files for the tempdb database. You place the files on a drive that contains the operating system files.
Does this meet the goal?

- A. Yes
- B. No

Answer: B

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 116

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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