



Vendor: Microsoft

Exam Code: 70-767

Exam Name: Implementing a Data Warehouse using SQL

Version: Demo

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 have the following line-of-business solutions:

- ERP system
- Online WebStore
- Partner extranet

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

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

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

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

Solution: Perform the following actions:

- Enable the Change Tracking feature for the Products table in the three source databases.
- Query the CHANGETABLE function from the sources for the deleted rows.
- Set the IsDisabled column to True on the data warehouse Products table for the listed rows.

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: B

Explanation:

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

References:

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

QUESTION 2

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

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

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

You need to pause the compute resources when the instance is not being used.

Solution: You use SQL Server Management Studio (SSMS).

Does the solution meet the goal?

- A. Yes
- B. No

Correct 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>

QUESTION 3

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

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

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

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

You need to update the SSIS package to add current prices to the Products table.

What should you use?

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

Correct Answer: D

Explanation:

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

References:

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

QUESTION 4

DRAG DROP

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 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



Correct Answer:

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

Create a partition function.

Create a partition scheme based on the partition function.

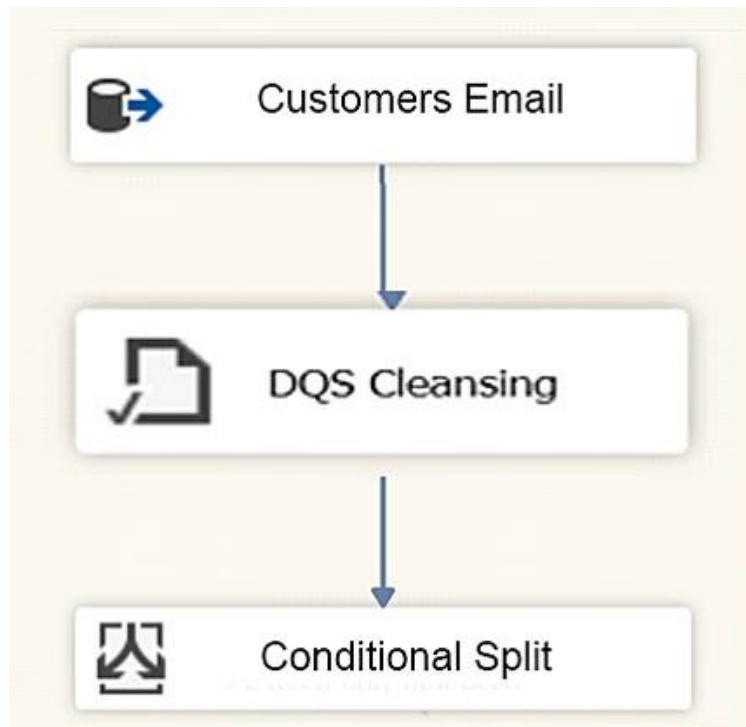
⏪ Execute an ALTER TABLE command to specify the partition function.



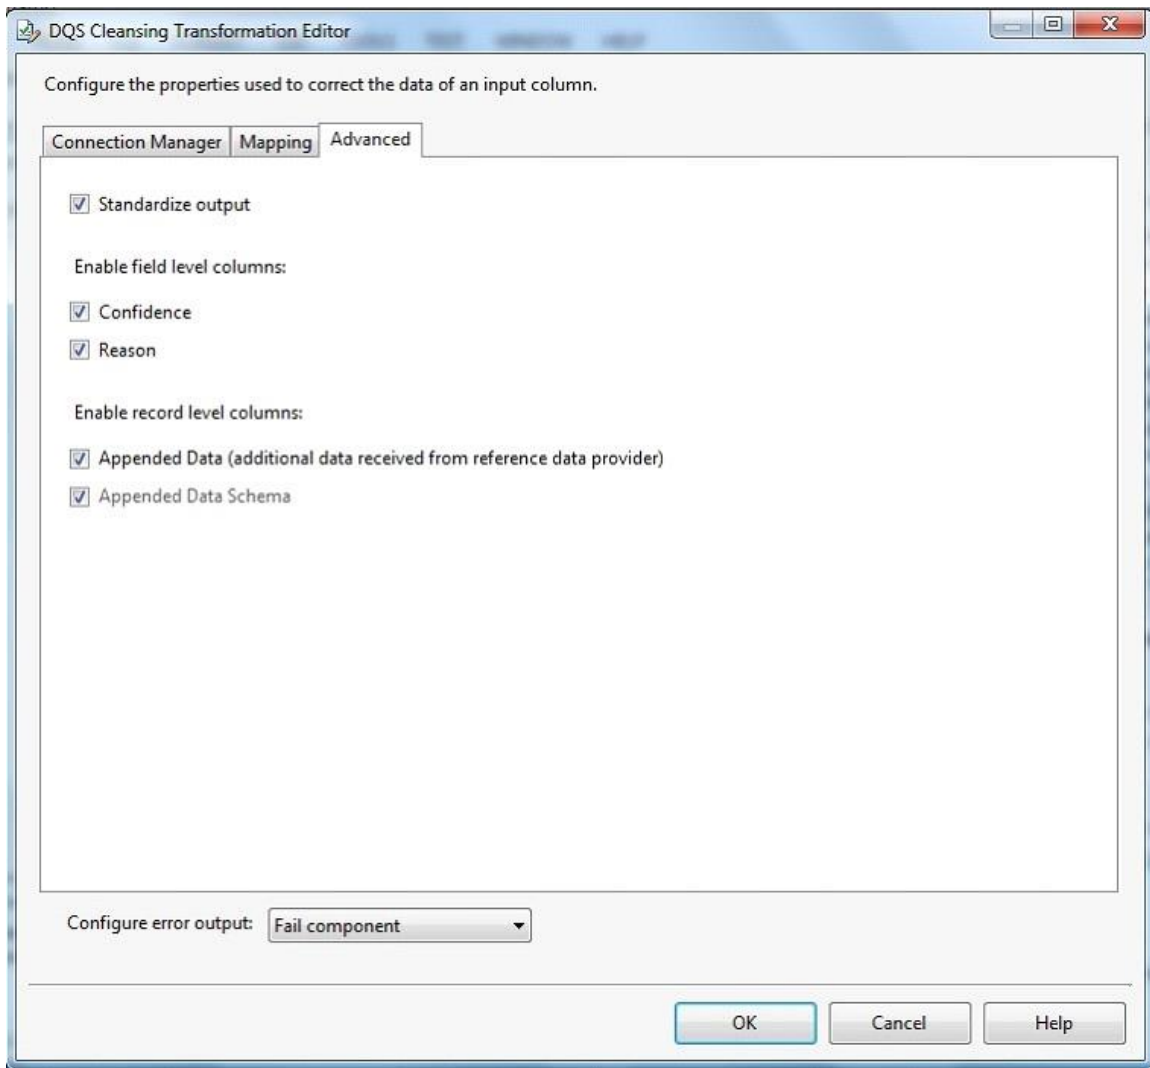
QUESTION 5

HOTSPOT

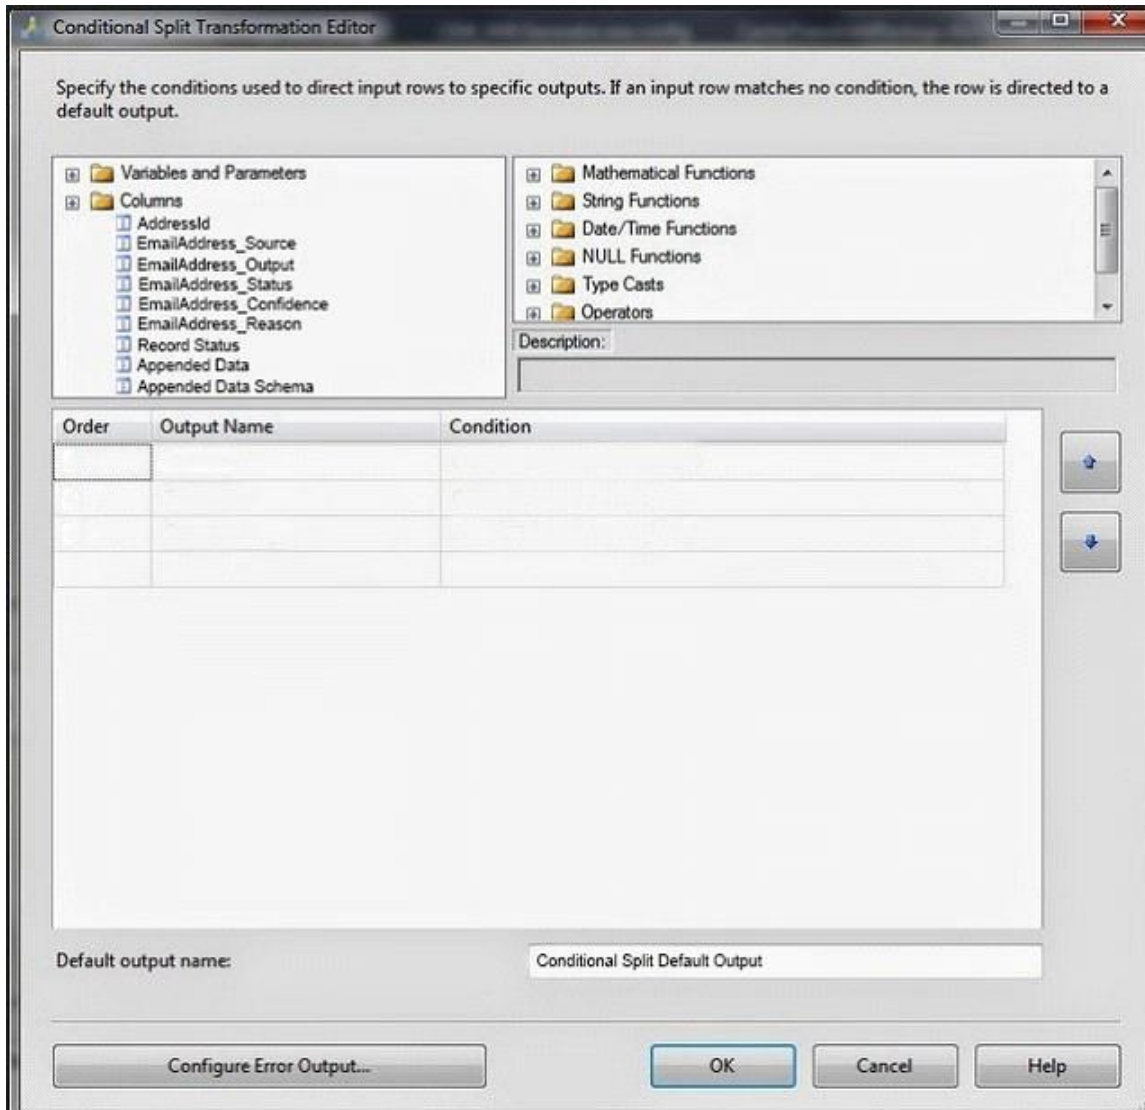
You have a Microsoft SQL Server Integration Services (SSIS) package that contains a Data Flow task as shown in the Data Flow exhibit. (Click the Exhibit button.)



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



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



You need to split the output of the DQS Cleansing task to obtain only Correct values from the EmailAddress column.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

Answer Area

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

Correct Answer:

Answer Area

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

QUESTION 6

DRAG DROP

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

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

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

You need to improve the quality of the data.

Which four actions should you perform in sequence?

To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

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 identified as being incorrectly entered.

Answer Area



Correct Answer:

Actions

- Add an external link from the Data Quality Service (DQS) knowledge base to the SQL Server table of replacement values.
- Publish the knowledge base and modify the ETL package to call it by using a .NET Script transformation.
- 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

- Import the raw data from the users to perform discovery.
- Map alternative values for entries that have been indentified as being incorrectly entered.
- Create a Data Quality Service (DQS) knowledge base.
- Publish the knowledge base and modify the ETL package to call it by using a .NET Script transformation.

QUESTION 7

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

Correct 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/>

QUESTION 8

DRAG DROP

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	Answer Area
<pre>EXEC sys.sp_cdc_enable_table @source_schema = N 'schema', @source_name = N 'Dimension.Customer', @role_name = NULL, @supports_net_changes = 1</pre>	
<pre>EXEC sys.sp_cdc_enable_db</pre>	
<pre>USE DB2</pre>	
<pre>EXEC sys.sp_cdc_enable_table</pre>	
<pre>USE DB1</pre>	
<pre>EXEC sys.sp_cdc_enable_db @source_schema = N 'schema', @source_name = N 'Dimension.Customer', @role_name = NULL, @supports_net_changes = 1</pre>	

Correct Answer:

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

```
USE DB1
```

```
EXEC sys.sp_cdc_enable_db
```

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



QUESTION 9

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

Correct 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>

QUESTION 10

DRAG DROP

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

Correct Answer:

Transact-SQL segments

- EmailAddress
- EmailAddress_Import
- NOT MATCHED BY SOURCE
- NOT MATCHED BY TARGET
- MATCHED

Answer area

```

MERC EmailAddress AS B
USIN EmailAddress_Import AS A
ON A.EmailAddressValue = B.EmailAddressValue
WHEN NOT MATCHED BY TARGET
THEN INSERT (EmailAddressValue) VALUES (A.EmailAddressValue)
WHEN NOT MATCHED BY SOURCE
THEN DELETE
    
```

QUESTION 11

HOTSPOT

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

Correct Answer:

Answer Area

UPDATE dbo.LastKeyByTable

SET

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

WHERE

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

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

Correct 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/>

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

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

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

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

You need to optimize performance.

Solution:

You run the following Transact-SQL statement:

```
CREATE STATISTICS Stat1
On Sales.SalesOrderHeader (OrderDate)
WITH SAMPLE 5 PERCENT
```

Does the solution meet the goal?

- A. Yes
- B. No

Correct Answer: A

Explanation:

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

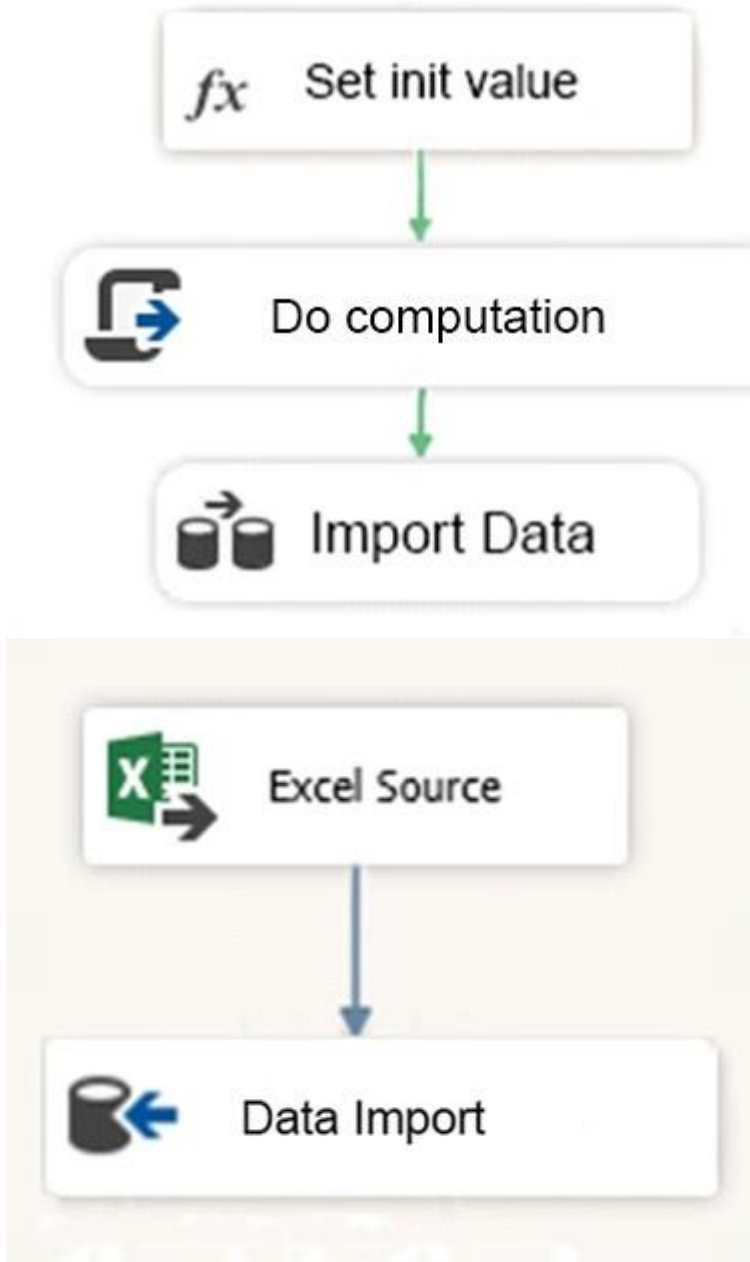
References:

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

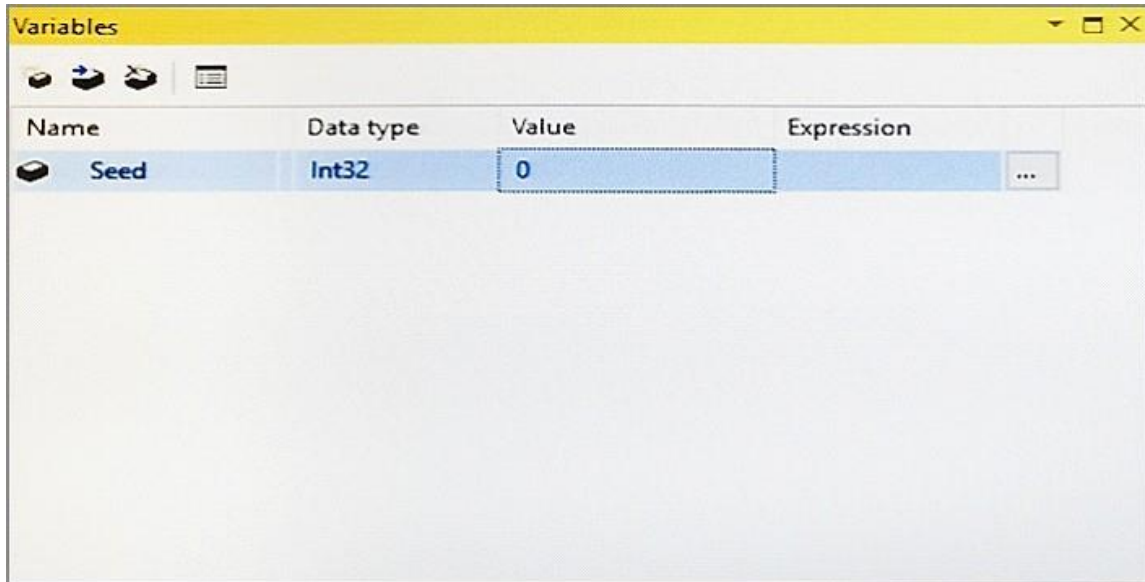
QUESTION 14

HOTSPOT

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"/>

Correct Answer:

Answer Area

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

QUESTION 15

HOTSPOT

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

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

The personal identification number is stored in a column named EmployeeSSN. All values in the EmployeeSSN column must be unique.

When importing employee data, you receive the error message shown in the SQL Error exhibit. (Click the Exhibit button.)

 Messages

Msg 2627, Level 14, State 1, Line 13

Violation of UNIQUE Key constraint 'UQ_DimEmplo_8549FE539cf2eca'. Cannot insert duplicate key object 'dbo.DimEmployee'. The duplicate key value is (<NULL>).

The statement has been terminated.

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

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

You remove the constraint on the EmployeeSSN column. You need to ensure that values in the EmployeeSSN column are unique.

For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Answer Area

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

Correct Answer:

Answer Area

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

QUESTION 16

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

Correct Answer: A

Explanation:

To pause a SQL Data Warehouse database, use any of these individual methods.

Pause compute with Azure portal

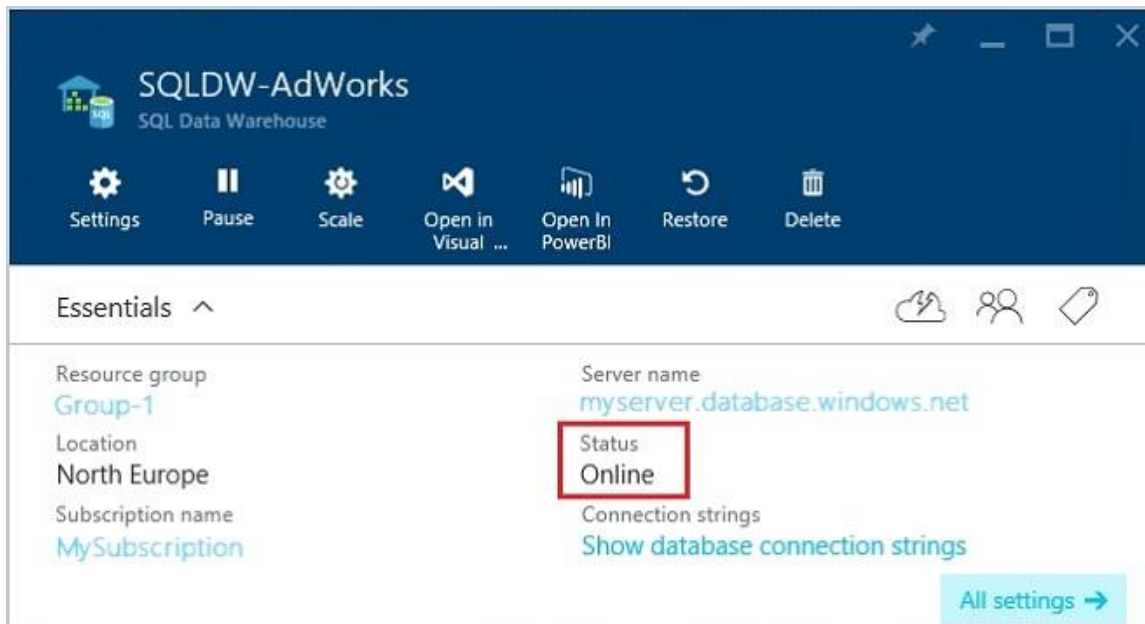
Pause compute with PowerShell

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-compute-bk>

QUESTION 17

DRAG DROP

You administer a Microsoft SQL Server Master Data Services (MDS) model. All model entity members have passed validation.

The current model version should be committed to form a record of master data that can be audited and create a new version to allow the ongoing management of the master data.

You lock the current version. You need to manage the model versions.

Which three actions should you perform in sequence?

To answer, move the appropriate actions from the list of actions to the answer area, and arrange them in the correct order.

Actions

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

Answer Area



Correct Answer:

Actions

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

Answer Area

- Validate the current version.
 - Commit the current version.
 - Create a copy of the current version.
- Navigation icons: Left arrow, Right arrow, Up arrow, Down arrow.

QUESTION 18

DRAG DROP

You have a Microsoft SQL Server Integration Services (SSIS) package that loads data into a data warehouse each night from a transactional system. The package also loads data from a set of Comma-Separated Values (CSV) files that are provided by your company's finance department.

The SSIS package processes each CSV file in a folder. The package reads the file name for the current file into a variable and uses that value to write a log entry to a database table.

You need to debug the package and determine the value of the variable before each file is processed.

Which four actions should you perform in sequence?

To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Click the **Start** toolbar button to commerce debugging the package.

When a breakpoint is reached, view the value of the variable by using the Variables window.

Open the Control Flow editor for the package.

When a breakpoint is reached, view the value of the variable by using the Locals window.

Set a breakpoint on the For Loop container.

Set a breakpoint on the Sequence container.

Open the Data Flow editor for the package.

Set a breakpoint on the Foreach Loop container.

Answer Area



Correct Answer:

Actions

- Click the **Start** toolbar button to commerce debugging the package.
- When a breakpoint is reached, view the value of the variable by using the Variables window.
- Open the Control Flow editor for the package.
- When a breakpoint is reached, view the value of the variable by using the Locals window.
- Set a breakpoint on the For Loop container.
- Set a breakpoint on the Sequence container.
- Open the Data Flow editor for the package.
- Set a breakpoint on the Foreach Loop container.

Answer Area

Open the Control Flow editor for the package.

Set a breakpoint on the Foreach Loop container.

Click the **Start** toolbar button to commerce debugging the package.

When a breakpoint is reached, view the value of the variable by using the Locals window.



QUESTION 19

HOTSPOT

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.

Knowledge Base Management

New Knowledge Base

Name: Sport Names

Description:

Create Knowledge Base from:

- None
- Existing Knowledge Base
- Import from DQS file

Select Activity:

- Domain Management
- Knowledge Discovery
- Matching Policy

Cancel Next

Correct Answer:

Knowledge Base Management

New Knowledge Base

Name: Sport Names

Description:

Create Knowledge Base from:

- None
- Existing Knowledge Base
- Import from DQS file

Select Activity:

- Domain Management
- Knowledge Discovery
- Matching Policy

Cancel Next

QUESTION 20

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

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

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

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



Solution:

You add the control flow to an ASP.NET assembly. You add a script task that references this assembly to each data warehouse load package.

Does the solution meet the goal?

- A. Yes
- B. No

Correct 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>