

Microsoft

Exam Questions 70-767

Implementing a SQL Data Warehouse (beta)



NEW QUESTION 1

You plan to deploy several Microsoft SQL Server Integration Services (SSIS) packages to a highly available SQL Server instance. The instance is configured to use an AlwaysOn availability group that has two replicas.

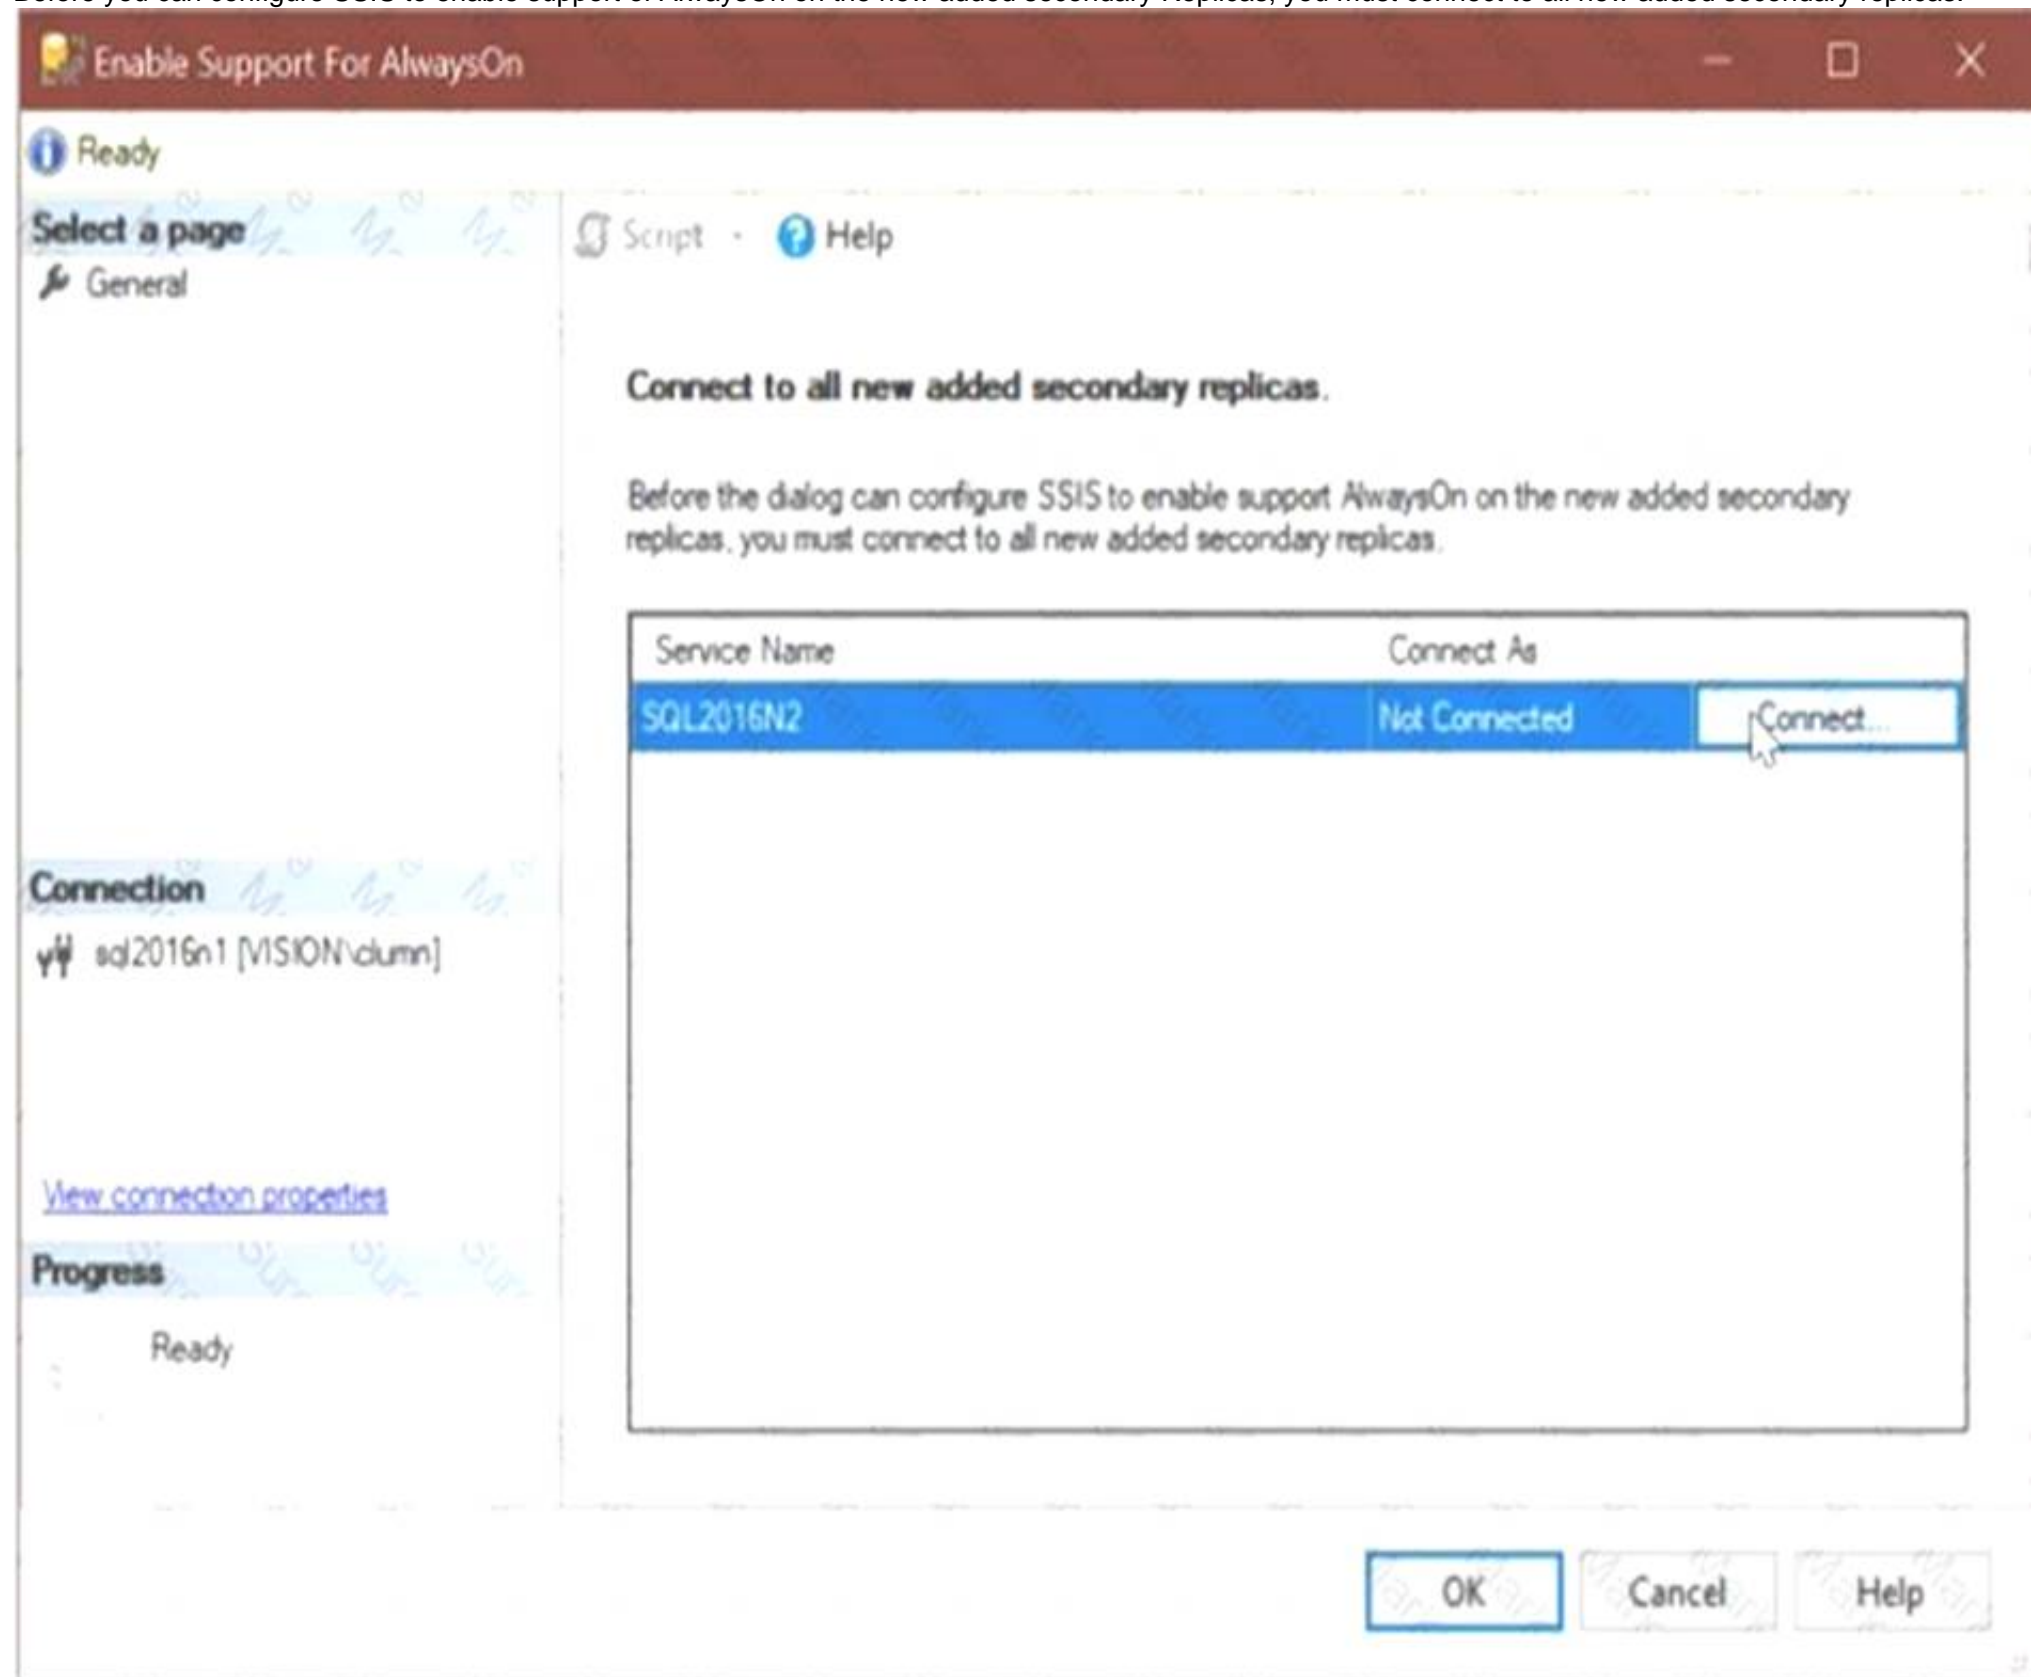
You need to identify which deployment method must be used to ensure that the packages are always accessible from all the nodes in the availability group. Which deployment method should you use for the packages?

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

Answer: A

Explanation:

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



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

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

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

NEW QUESTION 2

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 hapen 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 untill the data warehouse instance is resumed.	

What will hapen 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:

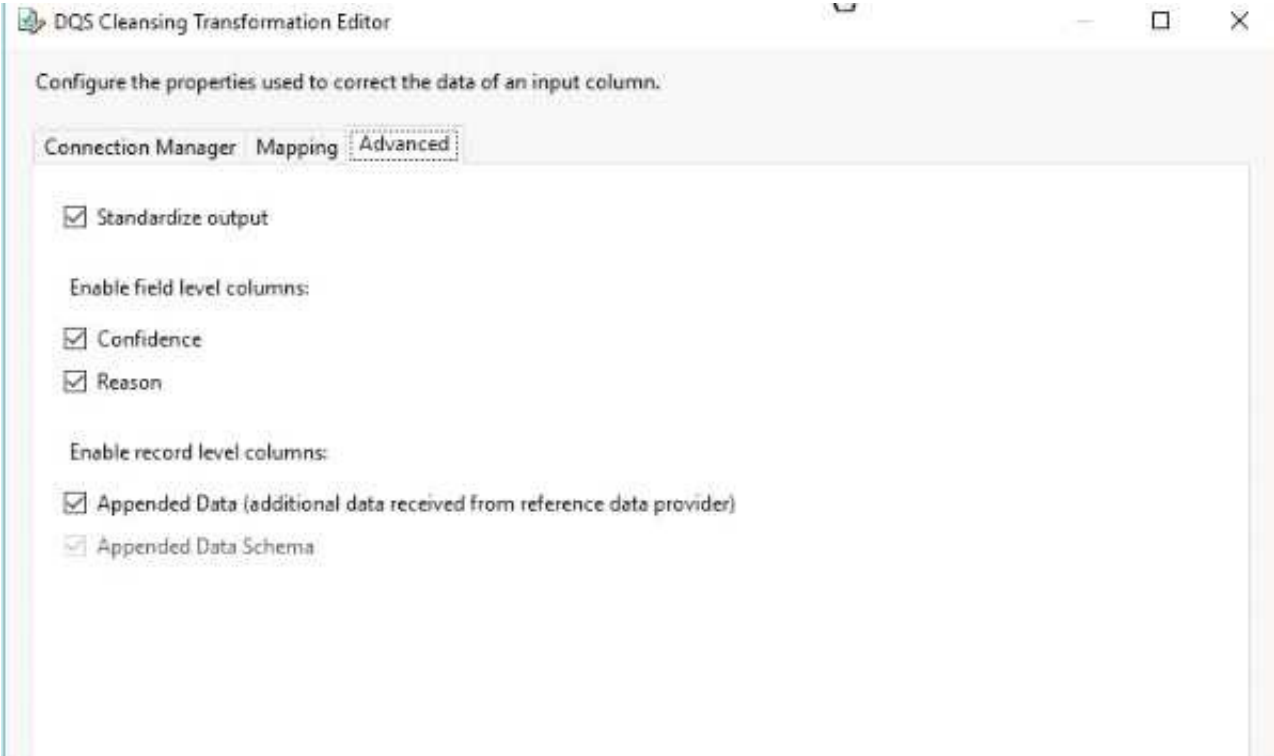
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage-compute-rest-api>

NEW QUESTION 3

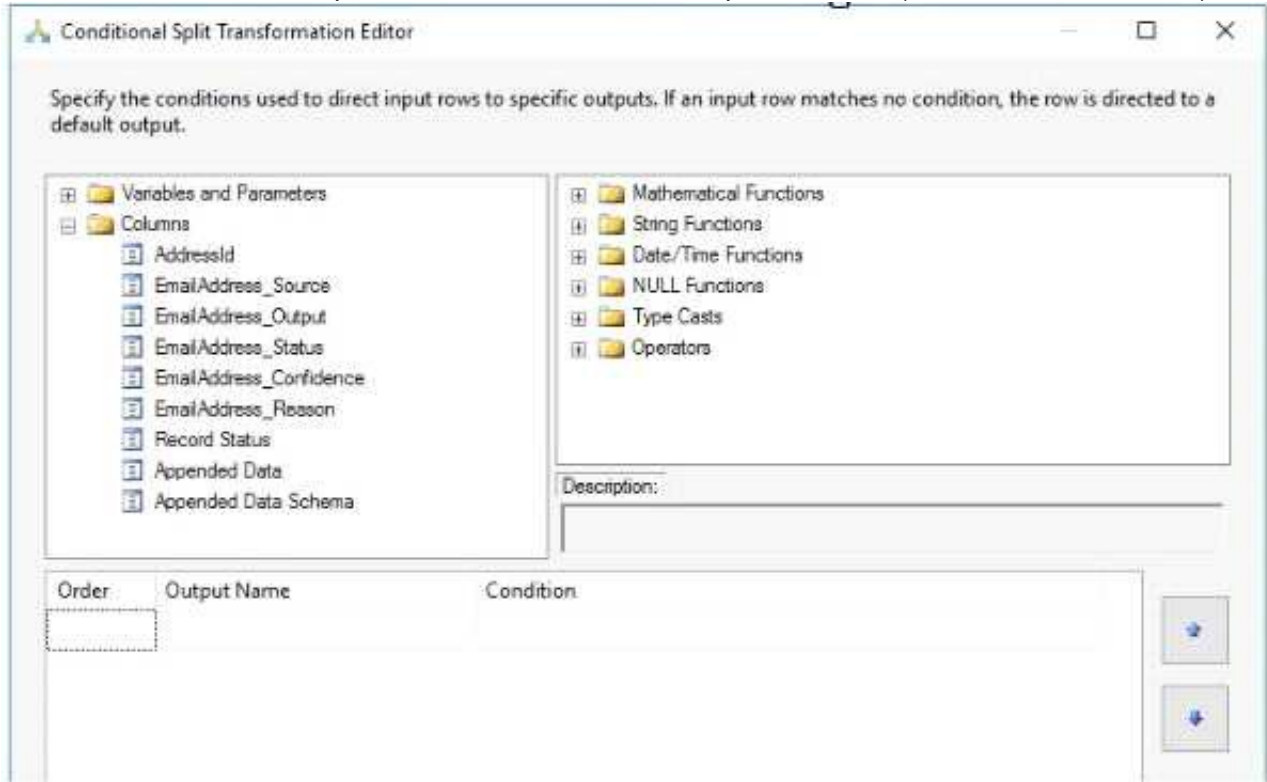
You have a Microsoft SQL Server Integration Services (SSIS) package that contains a Data Flow task as shown in the Data Flow exhibit. (Click the Exhibit button.)



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



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



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

Answer area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

	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"/>

NEW QUESTION 4

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 5

You need to ensure that a downstream system can consume data in a Master Data Services (MDS) system. What should you configure?

- A. a Data Collector

- B. a knowledgebase
- C. a matching policy
- D. a subscription view

Answer: D

Explanation:

Subscription views to consume your master data. References:
<https://docs.microsoft.com/en-us/sql/master-data-services/master-data-services-overview-mds?view=sql-server->

NEW QUESTION 6

You are building a server to host a data warehouse.
The planned disk activity for-khe data warehouse is five percent write activity and 95 percent read activity. You need to recommend a storage solution for the data files of the data warehouse. The solution must meet the following requirements:
*Ensure that the data warehouse is available if two disks fail.
*Minimize hardware costs.
Which RAID configuration should you recommend?

- A. RAID1
- B. RAID 5
- C. RAID 6
- D. RAID 10

Answer: C

Explanation:

According to the Storage Networking Industry Association (SNIA), the definition of RAID 6 is: "Any form of RAID that can continue to execute read and write requests to all of a RAID array's virtual disks in the presence of any two concurrent disk failures."

NEW QUESTION 7

Note: This question is part of a series of questions that use the same scenario. For your convenience, the scenario is repeated in each question. Each question presents a different goal and answer chokes, but the text of the scenario is exactly the same in each question in the series.
Start of repeated scenario
Contoso. Ltd. has a Microsoft SQL Server environment that includes SQL Server Integration Services (SSIS), a data warehouse, and SQL Server Analysis Services (SSAS) Tabular and multi-dimensional models.
The data warehouse stores data related to your company sales, financial transactions and financial budgets. All data for the data warehouse originates from the company s business financial system.
The data warehouse includes the following tables:

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

The company plans to use Microsoft Azure to store older records from the data warehouse. You must modify the database to enable the Stretch Database capability.
Users report that they are becoming confused about which city table to use for various queries. You plan to create a new schema named Dimension and change the name of the dbo.diamension_city table to Dimension.city. Data loss is not permissible, and you must not leave traces of the old table in the data warehouse.
The fact. Transaction table has measures named RawCost and Totaisale that calculate the wholesale cost of materials. You plan to create a measure that calculates the profit margin based on the two existing measures.
You must implement a partitioning scheme for the fact.Transaction table to move older data to less expensive storage. Each partition will store data for a single calendar year, as shown in the exhibit (Click the Exhibit button.) You must align the partitions.

	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	376	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 create the ProfitMargin measure for the fact. Transaction table.

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

Answer area

CREATE MEMBER
CREATE SET
CREATE SUBCUBE

CURRENTCUBE.Measures.ProfitMargin

AS 'Measures.

[TotalSale]
[RawCost]

/Measures.

[TotalSale]
[RawCost]

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer area

CREATE MEMBER
CREATE SET
CREATE SUBCUBE

CURRENTCUBE.Measures.ProfitMargin

AS 'Measures.

[TotalSale]
[RawCost]

/Measures.

[TotalSale]
[RawCost]

NEW QUESTION 8

You have a 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



- A. Mastered
- B. Not Mastered

Answer: A

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

NEW QUESTION 9

You have a data warehouse. You need to move a table named Fact.ErrorLog to a new filegroup named LowCost. Which three actions should you perform in sequence? To answer, move the appropriate actions from the list of actions to the answer area and arrange them in the correct order.

Actions

Add a file to the LowCost filegroup.

Rename the Fact.ErrorLog table to Fact.ErrorLogBak.

Drop the Fact.ErrorLog table.

Create a new Fact.ErrorLog table on the LowCost filegroup.

Add a filegroup named LowCost to the database.

Reorganize the clustered index on the Fact.ErrorLog table in the new filegroup.

Rebuild the clustered index on the Fact.ErrorLog table in the new filegroup.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Step 2:

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

Step 3:

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

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

This recreates the same index but on the SECONDARY filegroup.

References:

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

NEW QUESTION 10

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 on 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 are designing a warehouse named DW1.

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

```
AS RANGE LEFT FOR VALUES ('20150101', '20160101', '20170101', '20180101', '20190101', '20200101');
```

```
GO
```

```
CREATE PARTITION SCHEME schema1
```

```
AS PARTITION function1
```

```
ALL TO ([primary]);
```

```
GO
```

```
CREATE TABLE table1
```

```
(MyId BIGINT IDENTITY (1,1),
```

```
OrderDate datetime,
```

```
DueDate datetime,
```

```
AccountNumber nvarchar(15)
```

```
""
```

```
PRIMARY KEY (MyId, OrderDate))
```

```
ON schema1 (OrderDate)
```

```
GO
```

Reports are generated from the data in Table1.

You need to ensure that queries to DW1 return results as quickly as possible. Which column should appear in the WHERE statement clause of the query?

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


Answer: D


NEW QUESTION 11

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 16

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

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



Each source contains data that is not sorted.

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

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

Answer: C

NEW QUESTION 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 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

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>

NEW QUESTION 20

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

You need to recommend a solution to perform the actions.

What should you recommend?

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

Answer: A

Explanation:

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

- ▶ To enable Data Quality Services integration
- ▶ Open Master Data Services Configuration Manager.
- ▶ In the left pane, click Web Configuration.

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

References:

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

NEW QUESTION 24

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

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

NOTE: Each correct selection is worth one point.

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

References:

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

NEW QUESTION 25

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

After you answer a question in this 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 27

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

End of repeated scenario

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

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

Transact-SQL segments

EXEC sp_rename 'dbo.dim_City', 'City'

ALTER SCHEMA Dimension TRANSFER dbo.City

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

SELECT *
INTO Dimension.City
FROM dbo.dim_City

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

Answer area

CREATE SCHEMA Dimension
GO

Transact-SQL segment

Transact-SQL segment

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Transact-SQL segments

EXEC sp_rename 'dbo.dim_City', 'City'
ALTER SCHEMA Dimension TRANSFER dbo.City
DROP TABLE dbo.dim_City GO CREATE TABLE Dimension.City(...)
SELECT * INTO Dimension.City FROM dbo.dim_City
ALTER TABLE dbo.dim_City ADD Dimension.City VARCHAR(20) NULL

Answer area

CREATE SCHEMA Dimension GO
ALTER TABLE dbo.dim_City ADD Dimension.City VARCHAR(20) NULL
DROP TABLE dbo.dim_City GO CREATE TABLE Dimension.City(...)

NEW QUESTION 32

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

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

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

Does this meet the goal?

- A. Yes
- B. No

Answer: A

Explanation:

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

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

References:

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

NEW QUESTION 35

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

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

You 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 36

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 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 39

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.

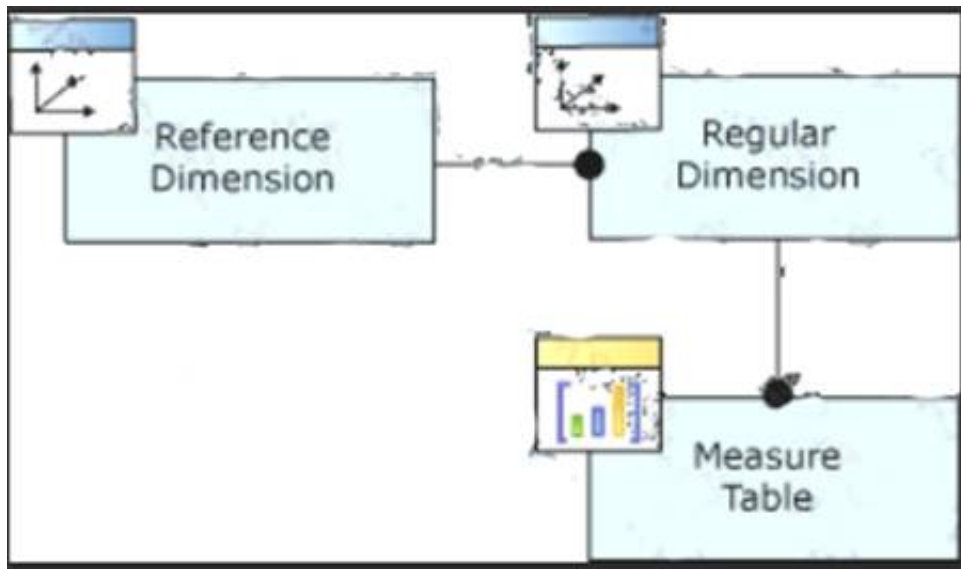
You need to recommend an appropriate dimension usage relationship. What should you recommend?

- 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 42

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 46

You have a database named DB1 that contains millions of rows. You plan to perform a weekly audit of the changes to the rows.

You need to ensure that you can view which rows were modified and the hour that the modification occurred. What should you do?

- A. Enable Policy-Based Management
- B. Configure Stretch Database.
- C. Configure an SSIS database.
- D. Enable change data capture.

Answer: D

Explanation:

SQL Server 2017 provides two features that track changes to data in a database: change data capture and change tracking.

Change data capture provides historical change information for a user table by capturing both the fact that DML changes were made and the actual data that was changed. Changes are captured by using an asynchronous process that reads the transaction log and has a low impact on the system.

References:

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

NEW QUESTION 49

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

Answer: B

Explanation:

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

Pause compute with PowerShell Pause compute with REST APIs

References:

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

NEW QUESTION 52

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

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

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

You remove the Integration Services job.

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

Which stored procedure should you execute?

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

Answer: F

Explanation:

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

References:

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

NEW QUESTION 56

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








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

Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

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

You need to configure the Fact.Order table.

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

Actions

Recreate the Fact.Order table on the partition scheme.

Execute an ALTER TABLE command to specify the partition function.

Create a partiotion scheme based on the partition function.

Execute an ALTER TABLE command to specify the partition scheme.

Recreate the Fact.Order table on the partition function.

Create a partition function.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Step 1: Create a partition function.

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

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

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

NEW QUESTION 57

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

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

You need to configure the measures.

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

NOTE: Each correct selection is worth one point.

Answer Area

Measure	Action
Measure1	<div><div></div><div>Turn off semi-additive behavior. Enable the First Child semi-additive function. Enable the FirstNonEmpty semi-additive function. Enable the LastNoneEmpty semi-additive function. Enable the Count semi-additive function. Enable the None semi-additive function.</div></div>
Measure2	<div><div></div><div>Turn off semi-additive behavior. Enable the First Child semi-additive function. Enable the FirstNonEmpty semi-additive function. Enable the LastNoneEmpty semi-additive function. Enable the Count semi-additive function. Enable the None semi-additive function.</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1:
The default setting is SUM (fully additive). Box 2:
FirstNonEmpty: The member value is evaluated as the value of its first child along the time dimension that contains data.
References:
<https://docs.microsoft.com/en-us/sql/analysis-services/multidimensional-models/define-semiadditive-behavior>

NEW QUESTION 62

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen.
You plan to deploy a Microsoft SQL server that will host a data warehouse named DB1. The server will contain four SATA drives configured as a RAID 10 array. You need to minimize write contention on the transaction log when data is being loaded to the database. Solution: You configure the server to automatically delete the transaction logs nightly.
Does this meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

NEW QUESTION 64

You have a database named OnlineSales that contains a table named Customers. You plan to copy incremental changes from the Customers table to a data warehouse every hour.
You need to enable change tracking for the Customers table.
How should you complete the Transact-SQL statements? 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

DATABASE [OnlineSales]

CHANGE_TRACKING = ON

SYSTEM_VERSIONING = ON

RECOVERY FULL

TABLE [dbo].[Customers]

ENABLE CHANGE_TRACKING

QUERY_STORE = ON

ENABLE_BROKER

Answer Area

ALTER

SET

(CHANGE_RETENTION = 2 DAYS, AUTO_CLEANUP = ON)

ALTER

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: DATABASE [OnlineSales]
Before you can use change tracking, you must enable change tracking at the database level. The following example shows how to enable change tracking by using ALTER DATABASE.
ALTER DATABASE AdventureWorks2012 SET CHANGE_TRACKING = ON
(CHANGE_RETENTION = 2 DAYS, AUTO_CLEANUP = ON) Box 2: CHANGE_TRACKING = ON
ALTER SET CHANGE_RETENTION
Box 3: ALTER TABLE [dbo].[Customers]
Change tracking must be enabled for each table that you want tracked. When change tracking is enabled, change tracking information is maintained for all rows in the table that are affected by a DML operation.
The following example shows how to enable change tracking for a table by using ALTER TABLE. ALTER TABLE Person.Contact
ENABLE CHANGE_TRACKING
WITH (TRACK_COLUMNS_UPDATED = ON) Box 4: ENABLE CHANGE_TRACKING
References:
<https://docs.microsoft.com/en-us/sql/relational-databases/track-changes/enable-and-disable-change-tracking-sql->

NEW QUESTION 69

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

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

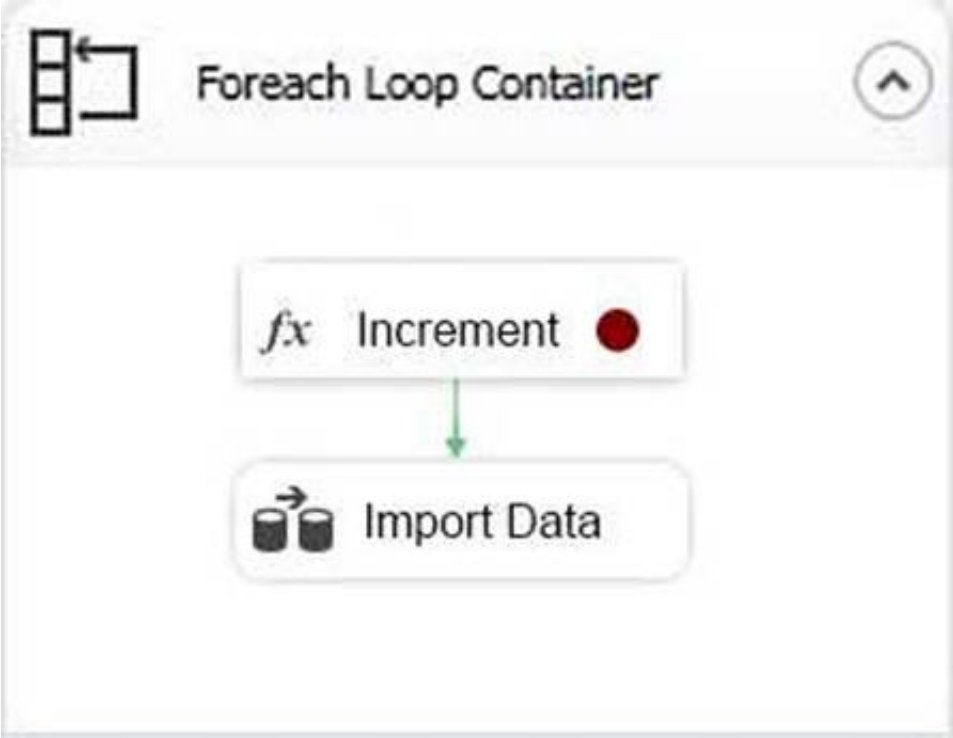
Answer: B

Explanation:

The number of files depends on the number of (logical) processors on the machine. As a general rule, if the number of logical processors is less than or equal to eight, use the same number of data files as logical processors. If the number of logical processors is greater than eight, use eight data files and then if contention continues, increase the number of data files by multiples of 4 until the contention is reduced to acceptable levels or make changes to the workload/code.
References: <https://docs.microsoft.com/en-us/sql/relational-databases/databases/tempdb-database>

NEW QUESTION 74

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



The package iterates over 100 files in a local folder. For each iteration, the package increments a variable named loop as shown in the Expression task exhibit. (Click the Exhibit button) and then imports a file. The initial value of the variable loop is 0.

Expression Builder

Specify the expression for the property: ConnectionString.

Variables

- Mathematical Functions
- String Functions
- Date/Time Functions
- NULL Functions
- Type Casts
- Operators

Description:

Expression:

`@[User::Loop] = @[User::Loop] + 10`

Evaluated value:

Evaluate Expression OK Cancel

You suspect that there may be an issue with the variable value during the loop. You define a breakpoint on the Expression task as shown in the BreakPoint exhibit. (Click the Exhibit button.)

Set Breakpoints - For Loop Container

Select the breakpoints in the task, For Loop, Foreach Loop, or Sequence to enable. Optionally, select the number of times a breakpoint is ignored before execution is suspended on the breakpoint.

Enabl...	Break Condition	Hit Count Type	Hit Count
<input checked="" type="checkbox"/>	Break when the container receives the OnPreExecute event	Hit count greater than or equal to	3
<input type="checkbox"/>	Break when the container receives the OnPostExecute event	Always	0
<input type="checkbox"/>	Break when the container receives the OnError event	Always	0
<input type="checkbox"/>	Break when the container receives the OnWarning event	Always	0
<input type="checkbox"/>	Break when the container receives the OnInformation event	Always	0
<input type="checkbox"/>	Break when the container receives the OnTaskFailed event	Always	0
<input type="checkbox"/>	Break when the container receives the OnProgress event	Always	0
<input type="checkbox"/>	Break when the container receives the OnQueryCancel event	Always	0
<input type="checkbox"/>	Break when the container receives the OnVariableValueChanged...	Always	0
<input type="checkbox"/>	Break when the container receives the OnCustomEvent event	Always	0
<input type="checkbox"/>	Break right before evaluating the expression	Always	0

OK Cancel Help

You need to check the value of the loop variable value.
 For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

	Yes	No
The value of the loop variable is 20 after the breakpoint is reached for the first time.	<input type="radio"/>	<input type="radio"/>
The loop variable resets to 0 when the breakpoint is reached.	<input type="radio"/>	<input type="radio"/>
When the code stops at a breakpoint, you can change the value of the loop variable.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

Break condition: When the task or container receives the OnPreExecute event.

Called when a task is about to execute. This event is raised by a task or a container immediately before it runs. The loop variable does not reset.

With the debugger, you can break, or suspend, execution of your program to examine your code, evaluate and edit variables in your program, etc.

NEW QUESTION 77

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

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

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

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

Answer: AD

NEW QUESTION 81

You have a data warehouse named DW1.

InDvfe 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 86

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

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

Sales contains details of records that have a type of refund and records that have a type of sales. The data warehouse design contains a table for sales data and a table for refund data.

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

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

Answer: B

Explanation:

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

References:

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

NEW QUESTION 89

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

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

You need to identify a column that can be aggregated across all dimensions. Which column should you identify?

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

Answer: C

Explanation:

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

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

NEW QUESTION 90

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"/>

- 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 94

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question. You are developing a Microsoft SQL Server Integration Services (SSIS) package. The package design consists of two differently structured sources in a single data flow. The Sales source retrieves sales transactions from a SQL Server database, and the Product source retrieves product details from an XML file. You need to combine the two data flow sources into a single output dataset. Which SSIS Toolbox item should you use?

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

Answer: G

Explanation:

The Merge Join transformation provides an output that is generated by joining two sorted datasets using a FULL, LEFT, or INNER join. For example, you can use a LEFT join to join a table that includes product information with a table that lists the country/region in which a product was manufactured. The result is a table that lists all products and their country/region of origin. References:
<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/merge-join-transformation>

NEW QUESTION 98

You are developing a Microsoft SQL Server Integration Services (SSIS) package. You create a data flow that has the following characteristics:

- The package moves data from the table [source].Tabid to DW.Tablel.
- All rows from [source].Table1 must be captured in DW.Tablel for error.Tablel.
- The table error.Tablel must accept rows that fail upon insertion into DW.Tablel due to violation of nullability or data type errors such as an invalid date, or invalid characters in a number.
- The behavior for the Error Output on the "OLE DB Destination" object is Redirect.
- The data types for all columns in [sourceJ.Tablel are VARCHAR. Null values are allowed.
- The Data access mode for both OLE DB destinations is set to Table or view - fast load.

```
The table definitions are as follows:

CREATE TABLE [source].Table1
(
    ID INT NULL,
    CreateDate VARCHAR(100) NULL,
    Date1 DATETIME2(7) NULL,
    Number1 VARCHAR(100) NULL
)
```

```
CREATE TABLE error.Table1
(
    ID INT NULL,
    CreateDate VARCHAR(100) NULL,
    Date1 DATETIME2(7) NULL,
    Number1 VARCHAR(100) NULL,
    ErrorDescription VARCHAR(255) NULL
)
```

Use the drop-down menus to select the answer choice that answers each question.

The ErrorDescription column is not yet populated in error.Table1. You must capture the error description for any rows redirected to the "Error OLE DB Destination". What should you do next?

- ☐ In "OLE DB Destination Error", map the ErrorCode field to ErrorDescription.
- ☐ Create an INSERT trigger on [Error].[Table1] to populate the ErrorDescription from ErrorCode.
- ☐ Add a Derived Column transformation before "OLE DB Destination". Use ErrorCode to populate ErrorDescription.
- ☐ Add a Script Component transformation before "OLE DB Destination Error". Capture the ErrorDescription with VB or C# code.

You execute the package. You note that all rows are redirected to OLE DB Destination Error, including both rows with bad data and rows with valid data. What is the next step?

- ☐ Uncheck the Check Constraints option in OLE DB Destination.
- ☐ Change the Data access mode for OLE DB Destination to Table or View.
- ☐ Uncheck the options Table Lock and Check Constraints for OLE DB Destination.
- ☐ Change the ValidateExternalMetadata setting for the OLE DB Destination Error object to False.
- ☐ Add a Conditional Split transformation before OLE DB Destination. Create outputs based on ErrorCode.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

The ErrorDescription column is not yet populated in error.Table1. You must capture the error description for any rows redirected to the "Error OLE DB Destination". What should you do next?

- ☒ In "OLE DB Destination Error", map the ErrorCode field to ErrorDescription.
- ☐ Create an INSERT trigger on [Error].[Table1] to populate the ErrorDescription from ErrorCode.
- ☐ Add a Derived Column transformation before "OLE DB Destination". Use ErrorCode to populate ErrorDescription.
- ☐ Add a Script Component transformation before "OLE DB Destination Error". Capture the ErrorDescription with VB or C# code.

You execute the package. You note that all rows are redirected to OLE DB Destination Error, including both rows with bad data and rows with valid data. What is the next step?

- ☒ Uncheck the Check Constraints option in OLE DB Destination.
- ☐ Change the Data access mode for OLE DB Destination to Table or View.
- ☐ Uncheck the options Table Lock and Check Constraints for OLE DB Destination.
- ☐ Change the ValidateExternalMetadata setting for the OLE DB Destination Error object to False.
- ☐ Add a Conditional Split transformation before OLE DB Destination. Create outputs based on ErrorCode.

NEW QUESTION 99

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

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

You need to ensure that the packa

ge records the current Log Sequence Number (LSN) in the source database before the package begins reading source tables.

Which SSIS Toolbox item should you use?

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

Answer: A

Explanation:

The CDC Control task is used to control the life cycle of change data capture (CDC) packages. It handles CDC package synchronization with the initial load package, the management of Log Sequence Number (LSN) ranges that are processed in a run of a CDC package.
 References: <https://docs.microsoft.com/en-us/sql/integration-services/control-flow/cdc-control-task>

NEW QUESTION 103

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

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

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

Transact-SQL segments

EmailAddress

EmailAddress_Import

NOT MATCHED BY SOURCE

NOT MATCHED BY TARGET

MATCHED

Answer area

MERGE

Transact-SQL segment

AS B

USING

Transact-SQL segment

AS A

ON

A.EmailAddressValue = B.EmailAddressValue

WHEN

Transact-SQL segment

THEN

INSERT (EmailAddressValue) VALUES (A.EmailAddressValue)

WHEN

Transact-SQL segment

THEN

DELETE

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: EmailAddress
 The EmailAddress table is the target. Box 2: EmailAddress_import
 The EmailAddress_import table is the source. Box 3: NOT MATCHED BY TARGET
 Box 4: NOT MATCHED BY SOURCE
 References: <https://docs.microsoft.com/en-us/sql/t-sql/statements/merge-transact-sql>

NEW QUESTION 104

Note: This question is part of a series of questions that use the same or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.
 You are a database administrator for an e-commerce company that runs an online store. The company has the databases described in the following table.

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

Each day, data from the table OnlineOrder in DB2 must be exported by partition. The tables must not be locked during the process.
 You need to write a Microsoft SQL Server Integration Services (SSIS) package that performs the data export. What should you use?

- A. Lookup transformation
- B. Merge transformation
- C. Merge Join transformation

- D. MERGE statement
- E. Union All transformation
- F. Balanced Data Distributor transformation
- G. Sequential container
- H. Foreach Loop container

Answer: E

Explanation:

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

References:

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

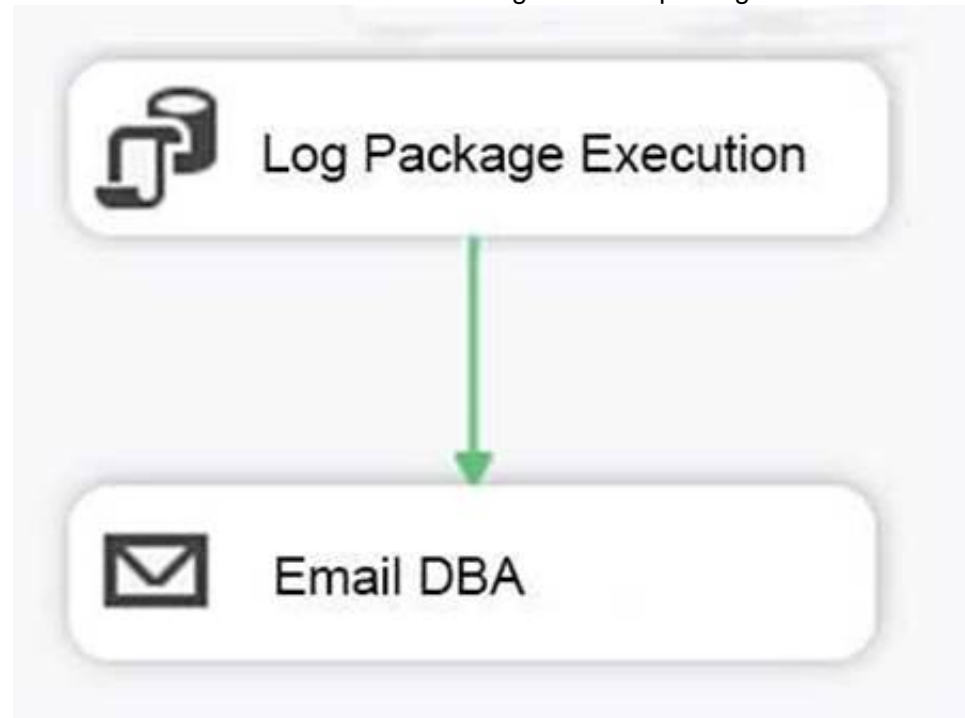
NEW QUESTION 108

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

Answer: B

Explanation:

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

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

NEW QUESTION 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 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 114

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

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

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

capture tables.

You discover that the job loads changes from the previous three days only. You need re ensure that the job loads changes from the previous week. Which stored procedure should you execute?

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

Answer: A

Explanation:

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

References:

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

NEW QUESTION 117

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

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

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

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

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

You need to load the data for the OnlineOrder table. What should you use?

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

Answer: H

Explanation:

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

References:

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

NEW QUESTION 120

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 123

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	<div></div>
Dimension.Customer	<div></div>
Dimension.Date	<div></div>

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Temporal table Box 2: Temporal table
Compared to CDC, Temporal tables are more efficient in storing historical data as it ignores insert actions. Box 3: Change Data Capture (CDC)
By using change data capture, you can track changes that have occurred over time to your table. This kind of functionality is useful for applications, like a data warehouse load process that need to identify changes, so they can correctly apply updates to track historical changes over time.
CDC is good for maintaining slowly changing dimensions.
Scenario: Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables.
The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated. References:
<https://www.mssqltips.com/sqlservertip/5212/sql-server-temporal-tables-vs-change-data-capture-vs-change-trac> <https://docs.microsoft.com/en-us/sql/relational-databases/tables/temporal-table-usage-scenarios?view=sql-server>

NEW QUESTION 125

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.
After you answer a question in this section, you will NOT be able to return to it As a result, these questions will not appear in the review screen.
You are the administrator of a Microsoft SOL Server Master Data Services (MDS) instance. The instance contains a model named Geography and a model named customer. The Geography model contains an entity named countryRegion.
You need to ensure that the countryRegion entity members are available in the customer model. Solution: In the Geography model, publish a business rule with a Change Value action.
Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 126

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 130

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

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

Which code segment should you add to the Script task?

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

Answer: D

NEW QUESTION 132

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 133

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

70-767 Practice Exam Features:

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

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