

Microsoft

Exam Questions 70-767

Implementing a SQL Data Warehouse (beta)



NEW QUESTION 1

You deploy a Microsoft Azure SQL Data Warehouse instance. The instance must be available eight hours each day. You need to pause Azure resources when they are not in use to reduce costs.

What will be the impact of pausing resources? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Answer Area

What will happen to existing queries that are running?

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

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

When you pause a database:

Compute and memory resources are returned to the pool of available resources in the data center Data Warehouse Unit (DWU) costs are zero for the duration of the pause.

Data storage is not affected and your data stays intact.

SQL Data Warehouse cancels all running or queued operations. When you resume a database:

SQL Data Warehouse acquires compute and memory resources for your DWU setting. Compute charges for your DWUs resume.

Your data will be available.

You will need to restart your workload queries. References:

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


NEW QUESTION 2

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

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

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

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



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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 3

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 4

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

Pause compute with PowerShell Pause compute with REST APIs References:

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

NEW QUESTION 5

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

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

You are 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: In the Customer model, add a domain-based attribute to reference the CountryRegion entity in the Geography model.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

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 or similar answer choices. An answer choice may be correct for more than one question in the series. Each question is independent of the other questions in this series. Information and details provided in a question apply only to that question.
 You are implementing a Microsoft SQL Server data warehouse with a multi-dimensional data model. You have a fact table that includes sales data for all products. The model includes a dimension named Geography that stores all geographies. You create a dimension that has a foreign key and provides the ability to analyze sales by the following sales channels: Internet or retail store.
 You need to update the data model to allow business users to analyze Internet sales by geography without changing the overall structure of the data model. What should you do?

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

Answer: D

NEW QUESTION 8

You need to build a knowledge base in Data Quality Services (DQS).
 You need to ensure that the data is validated by using a third-party data source before DQS processes 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	Answer Area
Perform Network Discovery.	
Configure a matching policy.	
Configure reference data services.	⬅️
Perform Domain Management.	⬆️
Perform Knowledge Discovery.	⬆️

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Building a DQS knowledge base involves the following processes and components:
 Step 1: Perform Knowledge Discovery
 A computer-assisted process that builds knowledge into a knowledge base by processing a data sample
 Step 2: Perform Domain Management
 An interactive process that enables the data steward to verify and modify the knowledge that is in knowledge base domains, each of which is associated with a data field. This can include setting field-wide properties, creating rules, changing specific values, using reference data services, or setting up term-based or cross-field relationships.
 Step 3: Configure reference Data Services
 A process of domain management that enables you to validate your data against data maintained and guaranteed by a reference data provider.
 Step 4: Configure a Matching Policy
 A policy that defines how DQS processes records to identify potential duplicates and non-matches, built into the knowledge base in a computer-assisted and interactive process.
 References: <https://docs.microsoft.com/en-us/sql/data-quality-services/dqs-knowledge-bases-and-domains>

NEW QUESTION 9

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

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

Answer: A

NEW QUESTION 10

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 10

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 dw1 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 will NOT store the user database files.

Does this meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 12

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

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

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

You need to configure data loading for the tables.

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

Table

Technology

Dimension.SalesTerritory

▼

Change Data Capture (CDC)
Change Tracking
Temporal table
Microsoft SQL Server snapshot replication

Dimension.Customer

▼

Change Data Capture (CDC)
Change Tracking
Temporal table
Microsoft SQL Server snapshot replication

Dimension.Date

▼

Change Data Capture (CDC)
Change Tracking
Temporal table
Microsoft SQL Server snapshot replication

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Scenario: The Dimension.SalesTerritory and Dimension.Customer tables are frequently updated
 Optimize data loading for the Dimension.SalesTerritory, Dimension.Customer, and Dimension.Date tables. Box 1: Change Tracking
 Box 2: Change Tracking Box 3: Temporal Table
 Temporal Tables are generally useful in scenarios that require tracking history of data changes.
 We recommend you to consider Temporal Tables in the following use cases for major productivity benefits.
 * Slowly-Changing Dimensions
 Dimensions in data warehousing typically contain relatively static data about entities such as geographical locations, customers, or products.
 References:
<https://docs.microsoft.com/en-us/sql/relational-databases/tables/temporal-table-usage-scenarios>

NEW QUESTION 14

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. You run the following Transact-SQL statement:

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

The query fails to return results.
 You need to determine why the query fails.
 Solution: You run the following Transact-SQL statement:

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

To use submit_time we must use sys.dm_pdw_exec_requests table. References:
<https://docs.microsoft.com/en-us/sql/relational-databases/system-dynamic-management-views/sys-dm-pdw-exec>

NEW QUESTION 15

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

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

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

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

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

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

Solution: You create measure for the Fact.Order table. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

You should use a columnstore index.

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

References:

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

NEW QUESTION 22

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 26

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No


Answer: A

Explanation:

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

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

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



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

References:

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

NEW QUESTION 30

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the

stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this 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 control flow package part. You add an instance of the control flow package part to each data warehouse load package. Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

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 32

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 33

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



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



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

Answer Area

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

References:

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

NEW QUESTION 34

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

After you answer a question in this 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 37

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 Diamension.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 41

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

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

You 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 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 are developing a Microsoft SQL Server Integration Services (SSIS) package. You need to use XPath to extract information from documents. Which SSIS Toolbox item should you use?

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

Answer: B

NEW QUESTION 43

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

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

Answer: A

Explanation:

We use ALTER TABLE to partition the table.

NEW QUESTION 46

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

- A. Yes
- B. No

Answer: B

Explanation:

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

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

<https://docs.microsoft.com/en-us/sql/relational-databases/policy-based-management/place-data-and-log-files-on->

NEW QUESTION 50

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution. After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. Your company uses Microsoft SQL Server to deploy a data warehouse to an environment that has a SQL Server Analysis Services (SSAS) instance. The data warehouse includes the Fact.Order table as shown in the following table definition. The table has no indexes.

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

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

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

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

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

Solution You create two nonclustered indexes. The first includes the [Order Date Key] and [Tax Amount] columns. The second will include the [Order Date Key] and [Total Excluding Tax] columns.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 51

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 52

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

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

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

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

NEW QUESTION 54

You are developing a Microsoft SQL Server Integration Services (SSIS) package to incrementally load new and changed records from a data source.


The SSIS package must load new records into Table1 and updated records into Table1_Updates. After loading records, the package must call a Transact-SQL statement to process updated rows according to existing business logic.

You need to complete the design of the SSIS package.

Which tasks should you use? To answer, drag the appropriate SSIS objects to the correct targets. Each SSIS object may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.

NOTE: Each correct selection is worth one point.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Step 1: CDC Control Task Get Processing Range Step 2: Mark Processed Range
 Step 3: Data Flow

The Data Flow task encapsulates the data flow engine that moves data between sources and destinations, and lets the user transform, clean, and modify data as it is moved. Addition of a Data Flow task to a package control flow makes it possible for the package to extract, transform, and load data.

Step 4: CDC Source

The CDC source reads a range of change data from SQL Server 2017 change tables and delivers the changes downstream to other SSIS component.

Step 5: CDC Splitter

The CDC splitter splits a single flow of change rows from a CDC source data flow into different data flows for Insert, Update and Delete operations.

References:

<https://docs.microsoft.com/en-us/sql/integration-services/control-flow/cdc-control-task> <https://docs.microsoft.com/en-us/sql/integration-services/control-flow/data-flow-task> <https://docs.microsoft.com/en-us/sql/integration-services/data-flow/cdc-splitter?view=sql-server-2017>

NEW QUESTION 59

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 61

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

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

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

You need to configure the Fact.Order table.

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

Actions

Recreate the Fact.Order table on the partition scheme.

Execute an ALTER TABLE command to specify the partition function.

Create a partition scheme based on the partition function.

Execute an ALTER TABLE command to specify the partition scheme.

Recreate the Fact.Order table on the partition function.

Create a partition function.

Answer Area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

Fact.Order partition.

Step 1: Create a partition function.

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

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

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

NEW QUESTION 63

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 style="border: 1px solid gray; padding: 5px;"> <p>▼</p> <p>Turn off semi-additive behavior.</p> <p>Enable the First Child semi-additive function.</p> <p>Enable the FirstNonEmpty semi-additive function.</p> <p>Enable the LastNoneEmpty semi-additive function.</p> <p>Enable the Count semi-additive function.</p> <p>Enable the None semi-additive function.</p> </div>
Measure2	<div style="border: 1px solid gray; padding: 5px;"> <p>▼</p> <p>Turn off semi-additive behavior.</p> <p>Enable the First Child semi-additive function.</p> <p>Enable the FirstNonEmpty semi-additive function.</p> <p>Enable the LastNoneEmpty semi-additive function.</p> <p>Enable the Count semi-additive function.</p> <p>Enable the None semi-additive function.</p> </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 66

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

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

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

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

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

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

Answer: G

Explanation:

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

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

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

NEW QUESTION 70

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-databases/policy-based-management/place-data-and-log-files-on->

NEW QUESTION 73

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		Answer Area
DATABASE [OnlineSales]	TABLE [dbo].[Customers]	ALTER <input type="text"/>
CHANGE_TRACKING = ON	ENABLE CHANGE_TRACKING	SET <input type="text"/>
SYSTEM_VERSIONING = ON	QUERY_STORE = ON	(CHANGE_RETENTION = 2 DAYS, AUTO_CLEANUP = ON)
RECOVERY FULL	ENABLE_BROKER	ALTER <input type="text"/>
		<input type="text"/>

- 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 75

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

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

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

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

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

Answer: D

Explanation:

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

References:

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

NEW QUESTION 79

You are developing a Microsoft SQL Server Integration Services (SSIS) package that loads a data warehouse. You need to inspect the data that is being processed by the package. What should you do first?

- A. Set a break point on the Control Flow path.
- B. Enable SQL Trace.
- C. Enable logging on the Data Flow path.
- D. Enable a data viewer on the Data Flow path.

Answer: A

NEW QUESTION 83

You have a data warehouse named DW1.

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

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

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

How many partitions should you recommend?

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

Answer: B

NEW QUESTION 87

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

NEW QUESTION 90

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

Start of repeated scenario

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

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

The data warehouse includes the following tables:

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

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

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

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

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

End of repeated scenario

You need to configure the fact.Transaction table.

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

Results Messages

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

Transact-SQL statements

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

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

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

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

Answer area



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

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

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

Answer area

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

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

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

NEW QUESTION 92

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 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 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 97

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 100

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 102

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 105

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 to Domainbased.

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 110

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

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

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

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

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

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

NOTE: Each correct selection is worth one point.

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

Answer: AD

NEW QUESTION 115


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 118

You need to recommend a storage solution for a data warehouse that minimizes load times. The solution must provide availability if a hard disk fails.
 Which RAID configuration should you recommend for each type of database file? To answer, drag the appropriate RAID configurations to the correct database file types. Each RAID configuration may be used once, more than once, or not at all. You may need to drag the split bar between panes or scroll to view content.
 NOTE: Each correct selection is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: RAID 5

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

Box 2: RAID 10

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

Note: In general RAID 1+0 will provide better throughput for write-intensive applications. The amount of performance gained will vary based on the HW vendor's RAID implementations. Most common alternative to RAID 1+0 is RAID 5. Generally, RAID 1+0 provides better write performance than any other RAID level providing data protection, including RAID 5.

NEW QUESTION 120

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 121

.....

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](#)