

70-767 Dumps

Implementing a SQL Data Warehouse (beta)

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



NEW QUESTION 1

You are designing an indexing strategy for a data warehouse. The data warehouse contains a table named Table1. Data is bulk inserted into Table1. You plan to create the indexes configured as shown in the following table.

Index name	Indexing specifications
Index1	<ul style="list-style-type: none">Index1 contains all the data in Table1.Queries against Index1 perform aggregation operations against hundreds of millions of rows.
Index2	<ul style="list-style-type: none">Index2 returns all the columns in this index.Index2 contains 80 percent of the columns in Table1.Index2 is used to assist with queries against other tables by performing point lookups against Table1.

Which type of index should you use to minimize the query times of each index? To answer, drag the appropriate index types to the correct indexes. Each index type 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.

Index Types

Clustered	Clustered columnstore
Hash	Heap
Nonclustered	Nonclustered columnstore

Answer Area

Index1:	Index type
Index2:	Index type

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Index Types

Clustered	Clustered columnstore
Hash	Heap
Nonclustered	Nonclustered columnstore

Answer Area

Index1:	Clustered columnstore
Index2:	Nonclustered columnstore

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 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.

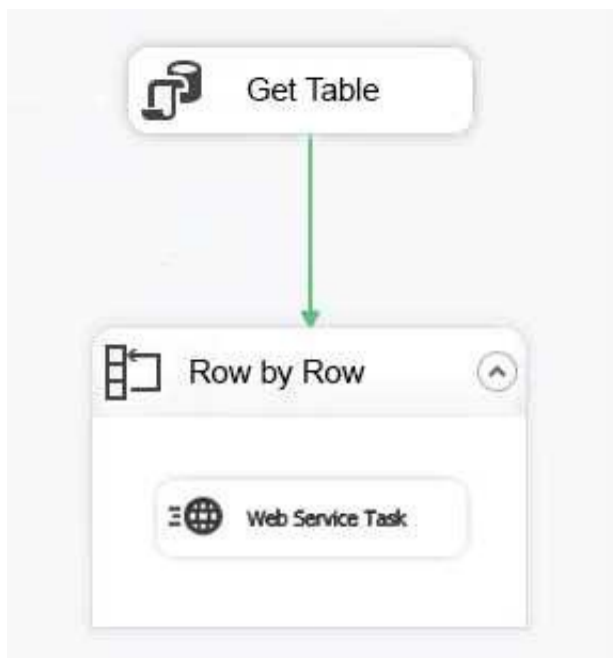
Does the solution meet the goal?

- A. Yes
B. No

Answer: A

NEW QUESTION 3

You have a Microsoft SQL Server Integration Services (SSIS) package that includes the control flow shown in the following diagram.



You need to choose the enumerator for the Foreach Loop container. Which enumerator should you use?

- A. Foreach SMO Enumerator
- B. Foreach Azure Blob Enumerator
- C. Foreach NodeList Enumerator
- D. Foreach ADO Enumerator

Answer: D

Explanation:

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

NEW QUESTION 4

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 5

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 6

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 7

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

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

You have the following requirements:

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

You need to configure data loading for the tables.

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

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

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

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

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

Box 2: Change Tracking Box 3: Temporal Table

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

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

* Slowly-Changing Dimensions

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

References:

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

NEW QUESTION 8

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 9

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 10

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 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 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 12

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

You have a database named DB1.

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

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

Answer: D

Explanation:

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

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

Note: sys.sp_cdc_change_job

References:

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

NEW QUESTION 13

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 16

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

After you answer a question in this 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 17

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.

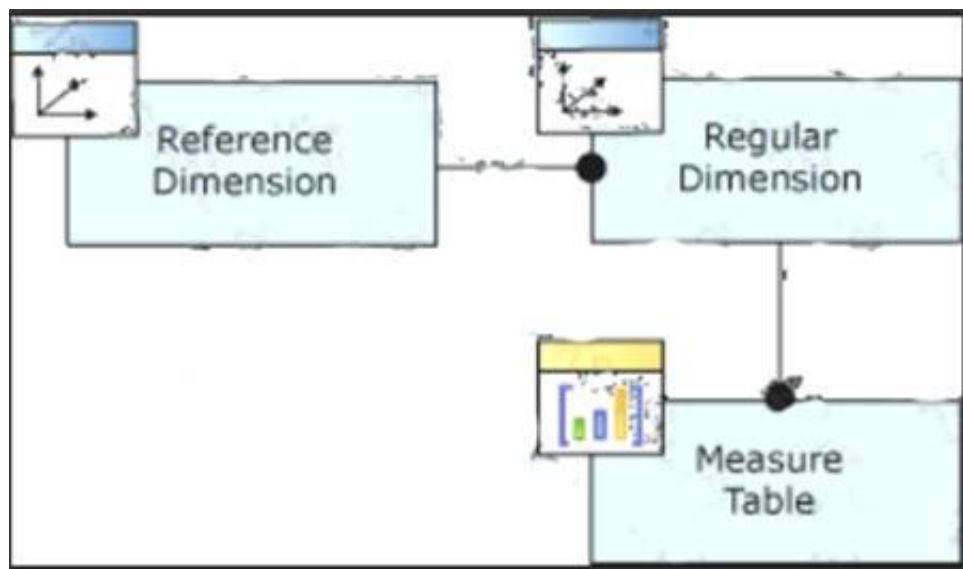
You need to recommend an appropriate dimension usage relationship. What should you recommend?

- 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 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 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 24

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 28

You create a Master Data Services (MDS) model that manages the master data for a Product dimension. The Product dimension has the following properties:

All the members of the Product dimension have a product type, a product subtype, and a unique product name.

Each product has a single product type and a single product subtype. The product type has a one-to-many relationship to the product subtype.

You need to ensure that the relationship between the product name, the product type, and the product subtype is maintained when products are added to or updates in the database.

What should you add to the model?

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

Answer: B

Explanation:

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

NEW QUESTION 31

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> 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>
Measure2	<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>

- 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 33

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 36

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

You are implementing a Microsoft SQL Server data warehouse with a multi-dimensional data model. When testing a pilot version of the data warehouse, business users observe that the number of products in

stock is inaccurate. The number of products in stock always increases and represents the total number of

products that have ever been in stock.

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

What should you do?

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

Answer: H

NEW QUESTION 39

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

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

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

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

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

You need to optimize performance.

Solution: You run the following Transact-SQL statement:

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

Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 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 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 46

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

Start of repeated scenario

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

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

The data warehouse includes the following tables:

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

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

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

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

End of repeated scenario

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

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

Answer area

CREATE PARTITION FUNCTION
CREATE PARTITION SCHEME
CREATE PROCEDURE
CREATE RESOURCE POOL

[DateRange] (DATETIME) AS RANGE RIGHT
[DateRange] (DATETIME) AS RANGE LEFT
[DateRange] (INT) AS RANGE RIGHT
[DateRange] (INT) AS RANGE LEFT
'20160101', '20170101')

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

CREATE PROCEDURE
[DateRange] (INT) AS RANGE LEFT

NEW QUESTION 51

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 53

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 one columnstore index that includes the [Order Date Key], [Tax Amount], and [Total Excluding Tax] columns.

Does the solution meet the goal?

- A. Yes
- B. No

Answer: A

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 55

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 56

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

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

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

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

Answer: D

NEW QUESTION 57

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 58

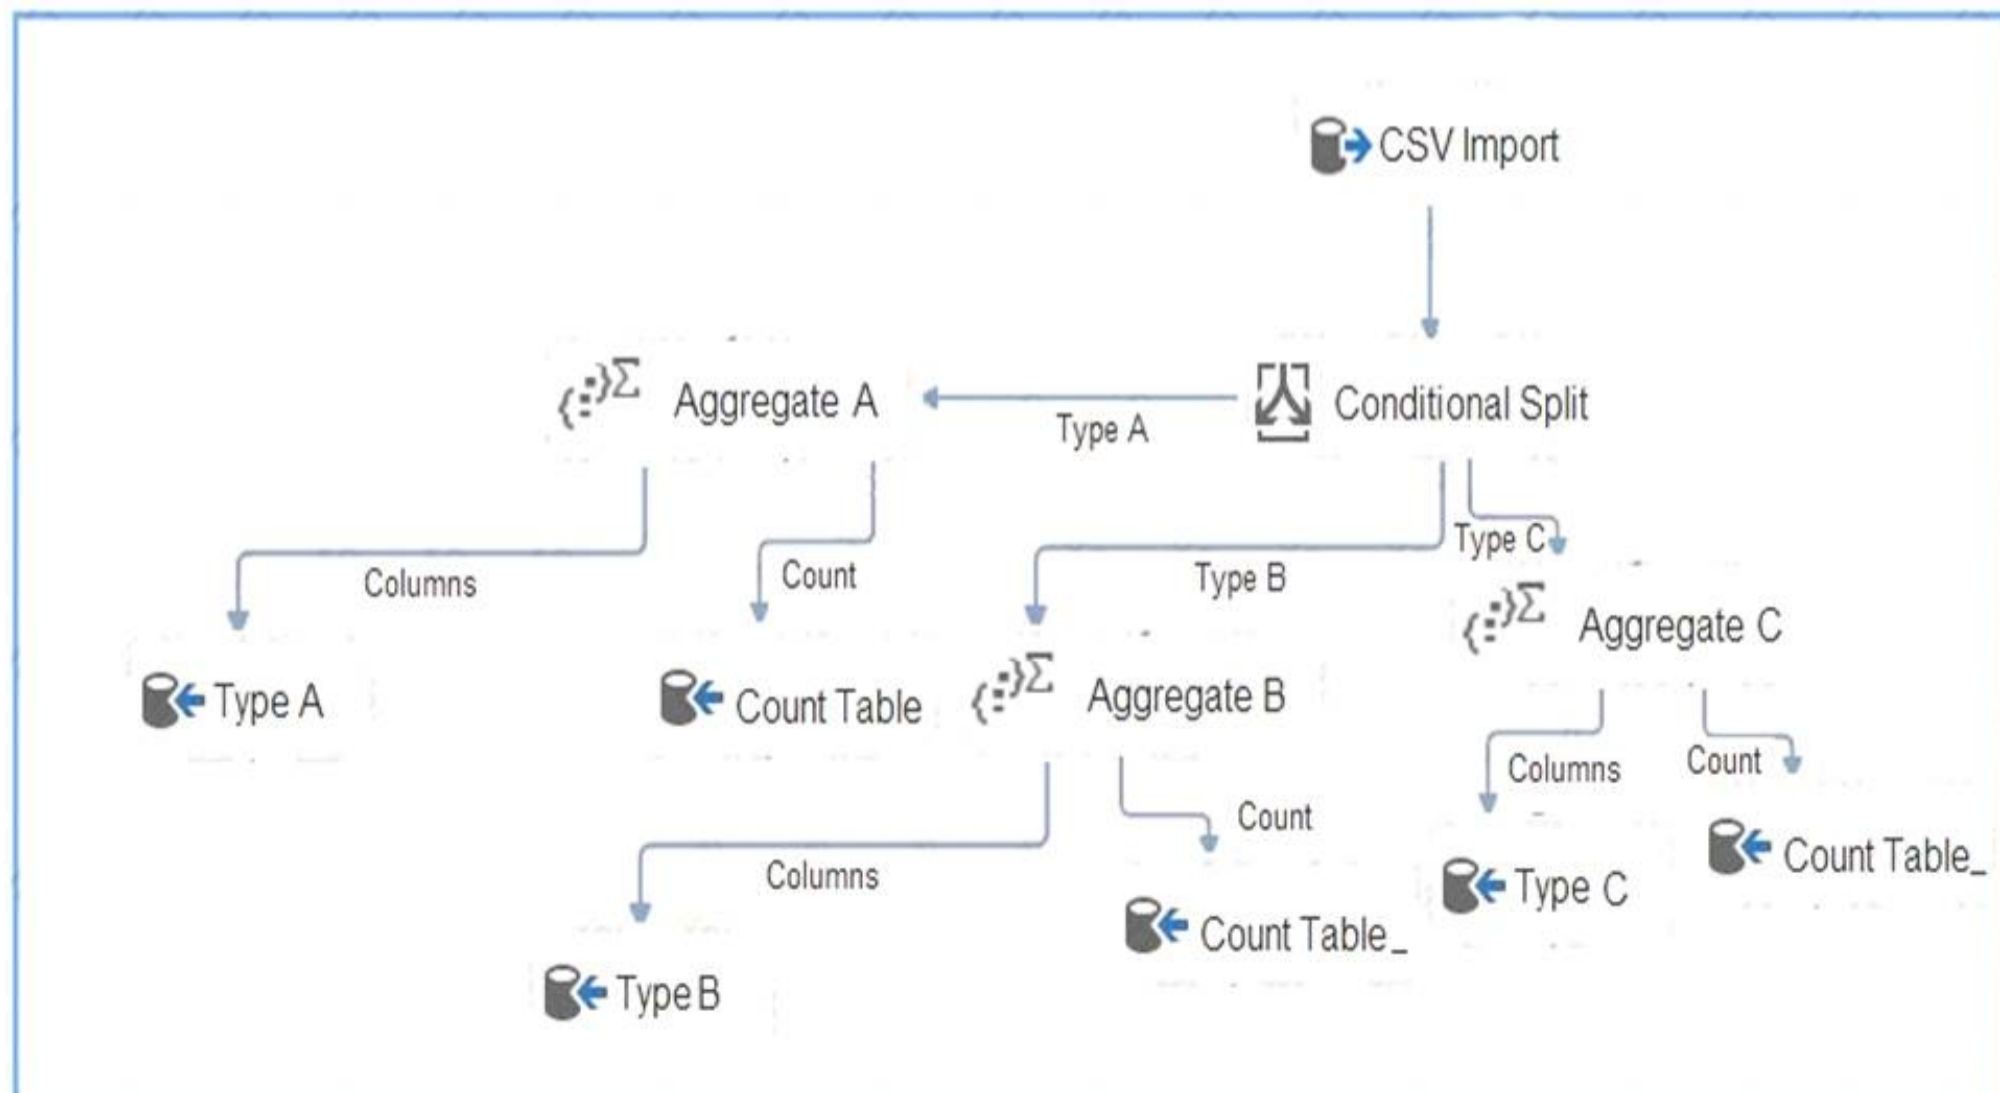
Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while others might not have a correct solution.

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

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

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

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



Does the solution meet the goal?

- A. Yes
- B. NO

Answer: A

Explanation:

The conditional split is correctly placed before the count.

NEW QUESTION 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

NEW QUESTION 61

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].Table1 to DW.Table1.
- All rows from [source].Table1 must be captured in DW.Table1 for error.Table1.
- The table error.Table1 must accept rows that fail upon insertion into DW.Table1 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 [source].Table1 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 63

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

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

You have the following line-of-business solutions:

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

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

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

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

- ▶ If a change is made to the ReferenceNr column in any of the sources, set the value of IsDisabled to True and create a new row in the Products table.
- ▶ If a row is deleted in any of the sources, set the value of IsDisabled to True in the data warehouse. Solution: Perform the following actions:
- ▶ Enable the Change Tracking feature for the Products table in the three source databases.
- ▶ Query the CHANGETABLE function from the sources for the deleted rows.
- ▶ Set the IsDisabled column to True on the data warehouse Products table for the listed rows. Does the solution meet the goal?

- A. Yes
- B. No

Answer: B

Explanation:

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

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

NEW QUESTION 68

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 segments

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

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

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

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

Answer area

>
<

^
v

- A. Mastered
B. Not Mastered

Answer: A

Explanation:

Transact-SQL segments

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

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

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

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

Answer area

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

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

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

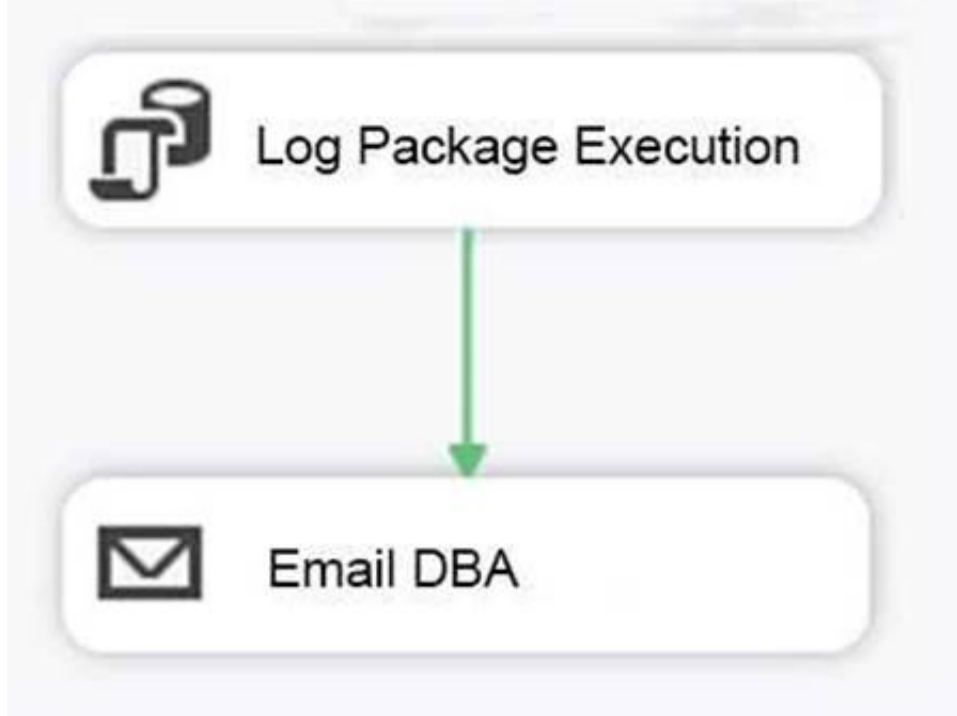
>
<

^
v

NEW QUESTION 71

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 72

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 76

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 80

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

Start of repeated scenario

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

Dimension.SalesTerritory, Dimension.Customer,

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

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

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

and is considered historical.

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

End of repeated scenario

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

Which technology should you use for each table?

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

Answer area

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

Table	Technology
Dimension.SalesTerritory	<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 81

You are designing a data transformation process using Microsoft SQL Server Integration Services (SSIS). You need to ensure that every row is compared with every other row during transformation.
What should you configure? To answer, select the appropriate options in the answer area.
NOTE: Each correct selection is worth one point.

SSIS package option	SSIS package configuration
Transformation type	<div><div></div><div>Fuzzy Grouping Merge Merge Join Multicast</div></div>
Transformation property	<div><div></div><div>Exhaustive Precision SpecialFlags</div></div>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

When you configure the Fuzzy Grouping transformation, you can specify the comparison algorithm that the transformation uses to compare rows in the transformation input. If you set the Exhaustive property to true, the transformation compares every row in the input to every other row in the input. This comparison algorithm may produce more accurate results, but it is likely to make the transformation perform more slowly unless the number of rows in the input is small.
References:
<https://docs.microsoft.com/en-us/sql/integration-services/data-flow/transformations/fuzzy-grouping-transformati>

NEW QUESTION 82

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

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

* One year free update

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

* Trusted by Millions

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

* Shop Securely

All transactions are protected by VeriSign!

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

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