Exam Number/Code:70-767

Exam Name: Implementing a Data
Warehouse using SQL

Version: Demo

Q1

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

Q2

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.
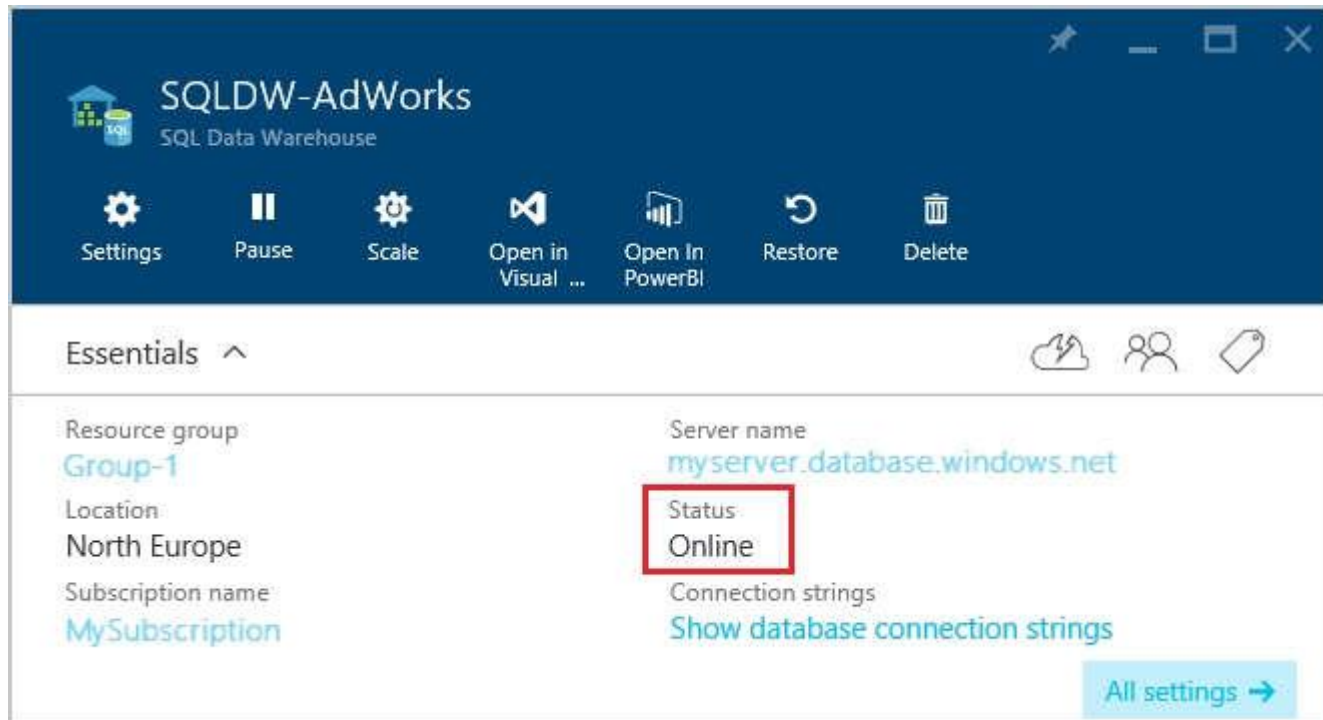
Pause compute with Azure portal

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-compute-bk

Q3

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

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

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

Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
To pause a SQL Data Warehouse database, use any of these individual methods.
Pause compute with Azure portal
Pause compute with PowerShell
Pause compute with REST APIs
References:
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-manage
-compute- overview

Q4

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

Q5

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 0 PERCENT
```

Does the solution meet the goal?

A. Yes
B. No

Answer: B

Explanation:
Microsoft recommend against specifying 0 PERCENT or 0 ROWS in a CREATE
STATISTICS..WITH SAMPLE statement. When 0 PERCENT or ROWS is specified, the
statistics object is created but does not contain statistics data.
References:
https://docs.microsoft.com/en-us/sql/t-sql/statements/create-statistics-transact-sql

Q6

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

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

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

You need to optimize performance.

Solution: You run the following Transact-SQL statement:

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

Does the solution meet the goal?

A. Yes
B. No

Answer: A

Explanation:
You can specify the sample size as a percent. A 5% statistics sample size would be helpful.
References:
https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-tables-statistics

Q7

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/

Q8

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/

Q9

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 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.

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/

Q10

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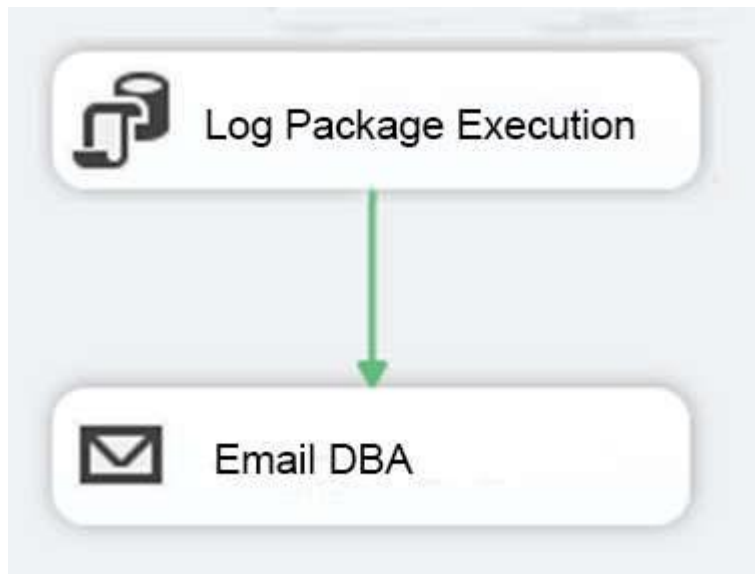
Solution: You add the control flow to a script task. You add an instance of the script task to the storage account in Microsoft Azure.

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