

Microsoft

70-475 Exam

**Microsoft Designing and Implementing Big Data Analytics
Solutions Exam**

**Questions & Answers
Demo**

Version: 9.0

Case Study: 1

Relecloud

General Overview

Relecloud is a social media company that processes hundreds of millions of social media posts per day and sells advertisements to several hundred companies.

Relecloud has a Microsoft SQL Server database named DB1 that stores information about the advertisers. DB1 is hosted on a Microsoft Azure virtual machine.

Physical locations

Relecloud has two main offices. The offices are located in San Francisco and New York City.

The offices are connected to each other by using a site-to-site VPN. Each office connects directly to the Internet.

Business model

Relecloud modifies the pricing of its advertisements based on trending topics. Topics are considered to be trending if they generate many mentions in a specific country during a 15-minute time frame. The highest trending topics generate the highest advertising revenue.

CTO statement

Relecloud wants to deliver reports to the advertisers by using Microsoft Power BI. The reports will provide real-time data on trending topics, current advertising rates, and advertising costs for a given month.

Relecloud will analyze the trending topics data, and then store the data in a new data warehouse for ad-hoc analysis. The data warehouse is expected to grow at a rate of 1 GB per hour or 8.7 terabytes (TB) per year. The data will be retained for five years for the purpose of long term trending.

Requirements

Business goals

Management at Relecloud must be able to view which topics are trending to adjust advertising rates in near real-time.

Planned changes

Relecloud plans to implement a new streaming analytics platform that will report on trending topics.

Relecloud plans to implement a data warehouse named DB2.

General technical requirements

Relecloud identifies the following technical requirements:

- Social media data must be analyzed to identify trending topics in real time.
- The use of Infrastructure as a Service (IaaS) platforms must be minimized, whenever possible.
- The real-time solution used to analyze the social media data must support scaling up and down without service interruption.

Technical requirements for advertisers

Relecloud identifies the following technical requirements for the advertisers

- The advertisers must be able to see only their own data in the Power BI reports.
- The advertisers must authenticate to Power BI by using Azure Active Directory (Azure AD) credentials.
- The advertisers must be able to leverage existing Transact-SQL language knowledge when developing the real-time streaming solution.
- Members of the internal advertising sales team at Relecloud must be able to see only the sales data of the advertisers to which they are assigned.
- The Internal Relecloud advertising sales team must be prevented from inserting, updating, and deleting rows for the advertisers to which they are not assigned.
- The internal Relecloud advertising sales team must be able to use a text file to update the list of advertisers, and then to upload the file to Azure Blob storage.

DB1 requirements

Relecloud identifies the following requirements for DB1:

- Data generated by the streaming analytics platform must be stored in DB1.
- The user names of the advertisers must be mapped to CustomerID in a table named Table2.
- The advertisers in DB1 must be stored in a table named Table1 and must be refreshed nightly.
- The user names of the employees at Relecloud must be mapped to EmployeeID in a table named Table3.

DB2 requirements

Relecloud identifies the following requirements for DB2:

- DB2 must have minimal storage costs.
- DB2 must run load processes in parallel.

- DB2 must support massive parallel processing.
- DB2 must be able to store more than 40 TB of data.
- DB2 must support scaling up and down, as required.
- Data from DB1 must be archived in DB2 for long-term storage.
- All of the reports that are executed from DB2 must use aggregation.
- Users must be able to pause DB2 when the data warehouse is not in use.
- Users must be able to view previous versions of the data in DB2 by using aggregates.

ETL requirements

Relecloud identifies the following requirements for extract, transformation, and load (ETL):

- Data movement between DB1 and DB2 must occur each hour.
- An email alert must be generated when a failure of any type occurs during ETL processing.

rls_table1

You execute the following code for a table named rls_table1.

```
create function rls_table1 (@CustomerId int, @SalesPersonId int)
    returns table
    with schemabinding
as
return
select 1 as result
from dbo.table1
join dbo.table2 on table1.customerid = Table2.CustomerId
where table2.UserName = user_name()
    and table1.customerid = @CustomerId
union all
select 1 as result
from dbo.table1
join dbo.table3 on table1.salespersonid = table3.employeeid
where table3.UserName = user_name()
    and table1.salespersonid = @SalesPersonId
go
```

dbo.table1

You use the following code to create Table1.

```
create table table1
(customerid int,
salespersonid int
...
)
go
```

Streaming data

The following is a sample of the Streaming data.

User	Country	Topic	Time
user1	USA	Topic1	2017-01-01T00:00:01.0000000Z

user1	USA	Topic3	2017-01-01T00:02:01.0000000Z
user2	Canada	Topic2	2017-01-01T00:01:11.0000000Z
user3	India	Topic1	2017-01-01T00:03:14.0000000Z

Question: 1

Which technology should you recommend to meet the technical requirement for analyzing the social media data?

- A. Azure Stream Analytics
- B. Azure Data Lake Analytics
- C. Azure Machine Learning
- D. Azure HDInsight Storm clusters

Answer: A

Azure Stream Analytics is a fully managed event-processing engine that lets you set up real-time analytic computations on streaming data.

Scalability

Stream Analytics can handle up to 1 GB of incoming data per second. Integration with Azure Event Hubs and Azure IoT Hub allows jobs to ingest millions of events per second coming from connected devices, clickstreams, and log files, to name a few. Using the partition feature of event hubs, you can partition computations into logical steps, each with the ability to be further partitioned to increase scalability.

Question: 2

DRAG DROP

You need to implement a solution that meets the data refresh requirement for DB1.

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

Answer Area

In DB1, create external objects.

From the Azure portal, export the storage account key.

In DB1, create a stored procedure that imports data from an external table to Table1.

From the Azure portal, create and schedule an Azure Automation job that executes the stored procedure.

In DB1, create a staging table.



Answer:

Answer Area

In DB1, create a stored procedure that imports data from an external table to Table1.

From the Azure portal, export the storage account key.

From the Azure portal, create and schedule an Azure Automation job that executes the stored procedure.

Azure Data Factory can be used to orchestrate the execution of stored procedures. This allows more complex pipelines to be created and extends Azure Data Factory's ability to leverage the computational power of SQL Data Warehouse.

From scenario:

Relecloud has a Microsoft SQL Server database named DB1 that stores information about the advertisers. DB1 is hosted on a Microsoft Azure virtual machine.

Relecloud identifies the following requirements for DB1:

Data generated by the streaming analytics platform must be stored in DB1.

The advertisers in DB1 must be stored in a table named Table1 and must be refreshed nightly.

Question: 3

DRAG DROP

You need to create a query that identifies the trending topics.

How should you complete the query? To answer, drag the appropriate values to the correct targets. Each value 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.

Values

- DATETIME
- GROUP BY
- HoppingWindow
- ORDER BY
- SlidingWindow
- TIMESTAMP
- VIEW BY

Answer Area

SELECT Country, Topic, count(*)

FROM Input BY Time

Country, Topic, (minute, 15)

Answer:

Answer Area

SELECT Country, Topic, count(*)

FROM Input BY Time

Country, Topic, (minute, 15)

From scenario: Topics are considered to be trending if they generate many mentions in a specific country during a 15-minute time frame.

Box 1: TimeStamp

Azure Stream Analytics (ASA) is a cloud service that enables real-time processing over streams of data flowing in from devices, sensors, websites and other live systems. The stream-processing logic in ASA is expressed in a SQL-like query language with some added extensions such as windowing for performing temporal calculations.

ASA is a temporal system, so every event that flows through it has a timestamp. A timestamp is assigned automatically based on the event's arrival time to the input source but you can also access a timestamp in your event payload explicitly using **TIMESTAMP BY**:

```
SELECT * FROM SensorReadings TIMESTAMP BY time
```

Box 2: **GROUP BY**

Example: Generate an output event if the temperature is above 75 for a total of 5 seconds

```
SELECT sensorId, MIN(temp) as temp
```

```
FROM SensorReadings
```

```
TIMESTAMP BY time
```

```
GROUP BY sensorId, SlidingWindow(second, 5)
```

```
HAVING MIN(temp) > 75
```

Box 3: **SlidingWindow**

Windowing is a core requirement for stream processing applications to perform set-based operations like counts or aggregations over events that arrive within a specified period of time. ASA supports three types of windows: Tumbling, Hopping, and Sliding.

With a Sliding Window, the system is asked to logically consider all possible windows of a given length and output events for cases when the content of the window actually changes – that is, when an event entered or existed the window.

Question: 4

HOTSPOT

You implement DB2.

You need to configure the tables in DB2 to host the data from DB1. The solution must meet the requirements for DB2.

Which type of table and history table storage should you use for the tables? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Answer area

Table: ▼

Change Data Capture

Change tracking

Temporal table

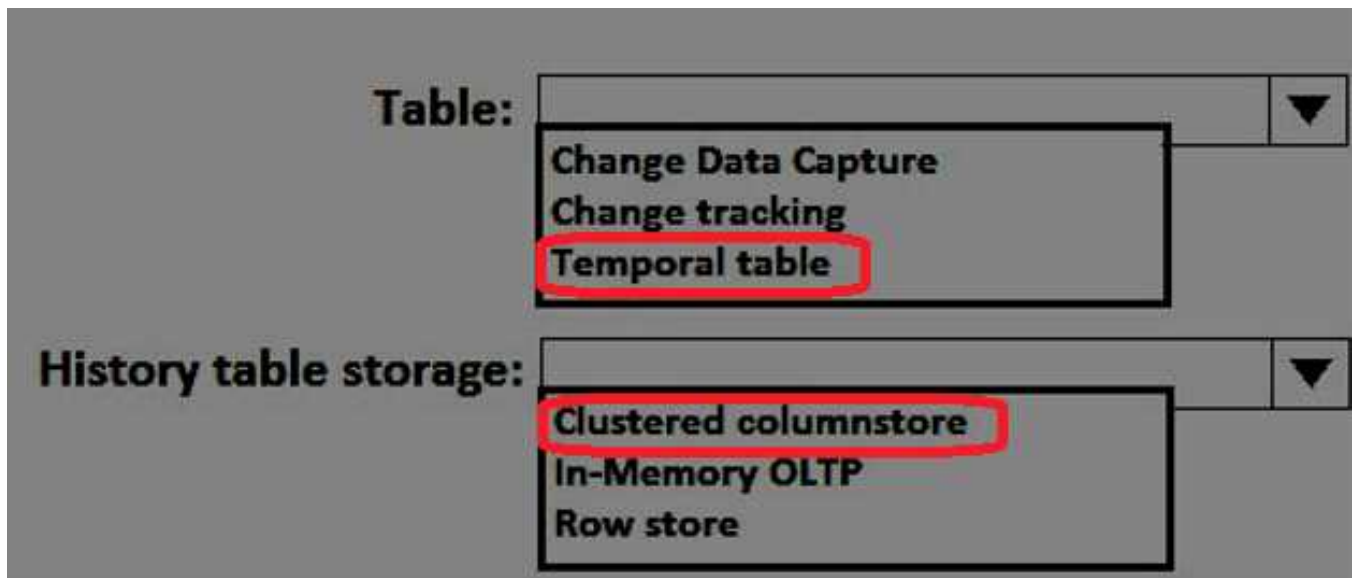
History table storage: ▼

Clustered columnstore

In-Memory OLTP

Row store

Answer:



From Scenario: Relecloud plans to implement a data warehouse named DB2.

Box 1: Temporal table

From Scenario:

Relecloud identifies the following requirements for DB2:

Users must be able to view previous versions of the data in DB2 by using aggregates.

DB2 must be able to store more than 40 TB of data.

A system-versioned temporal table is a new type of user table in SQL Server 2017, designed to keep a full history of data changes and allow easy point in time analysis. A temporal table also contains a reference to another table with a mirrored schema. The system uses this table to automatically store the previous version of the row each time a row in the temporal table gets updated or deleted. This additional table is referred to as the history table, while the main table that stores current (actual) row versions is referred to as the current table or simply as the temporal table.