



70-475^{Q&As}

Designing and Implementing Big Data Analytics Solutions

Pass Microsoft 70-475 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/70-475.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Microsoft
Official Exam Center

- ⚙ **Instant Download** After Purchase
- ⚙ **100% Money Back** Guarantee
- ⚙ **365 Days** Free Update
- ⚙ **800,000+** Satisfied Customers



**QUESTION 1**

You have a Microsoft Azure Stream Analytics job that contains several pipelines.

The Stream Analytics job is configured to trigger an alert when the sale of products in specific categories exceeds a specified threshold.

You plan to change the product-to-category mappings next month to meet future business requirements.

You need to create the new product-to-category mappings to prepare for the planned change. The solution must ensure that the Stream Analytics job only uses the new product-to-category mappings when the mappings are ready to be activated.

Which naming structure should you use for the file that contains the product-to-category mappings?

- A. Use any date after the day the file becomes active.
- B. Use any date before the day the categories become active.
- C. Use the date and hour that the categories are to become active.
- D. Use the current date and time.

Correct Answer: C

QUESTION 2

References: <https://docs.microsoft.com/en-us/azure/batch/batch-technical-overview>

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.

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

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

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.

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:

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

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

Relecloud plans to implement a data warehouse named DB2.

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.

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.

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.

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.

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.

Sample code and data:

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 = suser_sname()
    and table1.customerid = @CustomerId
union all
select 1 as result
from dbo.table1
join dbo.table3 on table1.salespersonid = table3.EmployeeId
where table3.UserName = suser_sname()
    and table1.salespersonid = @SalesPersonId
go
```

You use the following code to create Table1.

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

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



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.

Hot Area:

Answer area

Table: ▼

Change Data Capture
Change tracking
Temporal table

History table storage: ▼

Clustered columnstore
In-Memory OLTP
Row store

Correct Answer:

Answer area

Table: ▼

Change Data Capture
Change tracking
Temporal table

History table storage: ▼

Clustered columnstore
In-Memory OLTP
Row store

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.

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

Box 2: Clustered Columnstore

From Scenario:

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.

Relecloud identifies the following requirements for DB2:

Data from DB1 must be archived in DB2 for long-term storage.

How SQL Database deletes aged rows?

The cleanup process depends on the index layout of the history table. It is important to notice that only history tables with a clustered index (B-tree or columnstore) can have finite retention policy configured.

Reference: <https://docs.microsoft.com/en-us/sql/relational-databases/tables/manage-retention-of-historical-data-in-system-versioned-temporal-tables>

QUESTION 3

You need to recommend a data analysis solution for 20,000 Internet of Things (IoT) devices. The solution must meet the following requirements:

Each device must be identified by using its own credentials.

Each device must be able to route data to multiple endpoints.

The solution must require the minimum amount of customized code.

What should you recommend?

- A. Microsoft Azure Notification Hubs
- B. Microsoft Azure IoT Hub
- C. Microsoft Azure Service Bus
- D. Microsoft Azure Event Hubs



Correct Answer: B

IoT Hub gives you a secure communication channel for your devices to send data.

Per-device authentication enables each device to connect securely to IoT Hub and for each device to be managed securely.

Azure IoT Hub has built-in message routing functionality that gives you flexibility to set up automatic rules-based message fan-out:

Use message routing to control where your hub sends device telemetry.

There is no additional cost to route messages to multiple endpoints.

No-code routing rules take the place of custom message dispatcher code.

References:

<https://docs.microsoft.com/en-us/azure/iot-hub/about-iot-hub>

QUESTION 4

Note: This question is part of a series of questions that present the same scenario. Each question in the series contains a unique solution that might meet the stated goals. Some question sets might have more than one correct solution, while

the 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 Azure SQL data warehouse and a web application.

The data warehouse will ingest 5 TB of data from an on-premises Microsoft SQL Server database daily. The web application will query the data warehouse.

You need to design a solution to ingest data into the data warehouse.

Solution: You use the bcp utility to export CSV files from SQL Server and then to import the files to Azure SQL Data Warehouse.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

If you need the best performance, then use PolyBase to import data into Azure SQL warehouse. References: <https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-migrate-data>

QUESTION 5



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

After you answer a question in this section, you will NOT be able to return to it. As a result, these questions will not appear in the review screen. You have a Microsoft Azure deployment that contains the following services: Azure Data Lake Azure Cosmos DB Azure Data Factory Azure SQL Database

You load several types of data to Azure Data Lake.

You need to load data from Azure SQL Database to Azure Data Lake.

Solution: You use the Azure Import/Export service.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Azure Import/Export service is used to securely import large amounts of data to Azure Blob storage and Azure Files by shipping disk drives to an Azure datacenter. This service can also be used to transfer data from Azure Blob storage to disk

drives and ship to your on-premises sites.

Note:

You can use the Copy Activity in Azure Data Factory to copy data to and from Azure Data Lake Storage Gen1 (previously known as Azure Data Lake Store). Azure SQL database is supported as source.

References: <https://docs.microsoft.com/en-us/azure/data-factory/connector-azure-data-lake-store>

[Latest 70-475 Dumps](#)

[70-475 Study Guide](#)

[70-475 Braindumps](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.pass4itsure.com/allproducts>

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 One Year Free Update <p>Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 Money Back Guarantee <p>To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 Security & Privacy <p>We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © pass4itsure, All Rights Reserved.