DP-200 Q&As

Implementing an Azure Data Solution

Pass Microsoft DP-200 Exam with 100% Guarantee

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

https://www.passapply.com/dp-200.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

SIMULATION Use the following login credentials as needed:

Azure Username: xxxxx Azure Password: xxxxx The following information is for technical support purposes only:

Lab Instance: 10277521

You plan to create large data sets on db2.

You need to ensure that missing indexes are created automatically by Azure in db2. The solution must apply ONLY to db2.

To complete this task, sign in to the Azure portal.

Correct Answer: See the below.

Explanation:

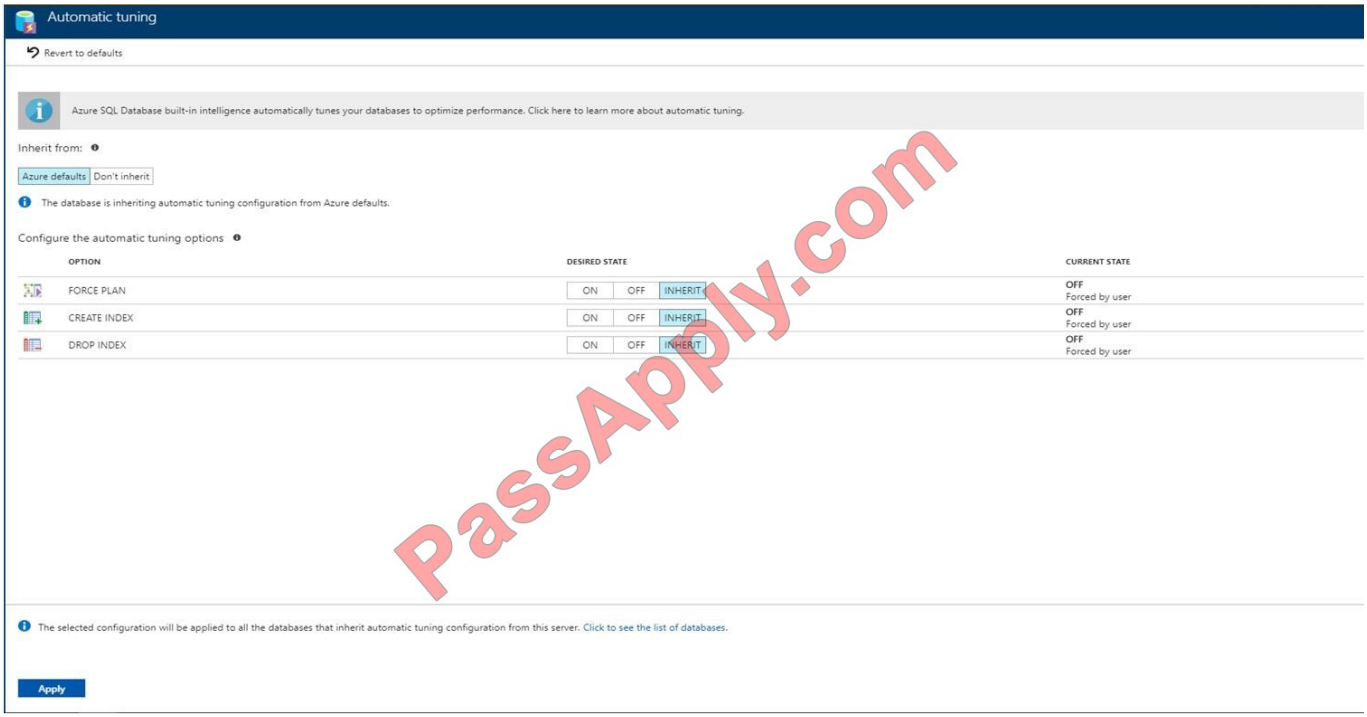

1. To enable automatic tuning on Azure SQL Database logical server, navigate to the server in Azure portal and then select Automatic tuning in the menu.

2.

Select database db2

3.

Click the Apply button

Reference: https://docs.microsoft.com/en-us/azure/sql-database/sql-database-automatic-tuning-enable

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 questions sets might have more than one correct solution,

while others might not have a correct solution.

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

You have an Azure subscription that contains an Azure Storage account.

You plan to implement changes to a data storage solution to meet regulatory and compliance standards.

Every day, Azure needs to identify and delete blobs that were NOT modified during the last 100 days.

Solution: You schedule an Azure Data Factory pipeline.

Does this meet the goal?

A. Yes

B. No

Correct Answer: B

Instead apply an Azure Blob storage lifecycle policy.

Reference: https://docs.microsoft.com/en-us/azure/storage/blobs/storage-lifecycle-management-concepts?tabs=azure-portal

QUESTION 3

HOTSPOT

You have the following Azure Stream Analytics query.

```
WITH

step1 AS (SELECT *
      FROM input1
      PARTITION BY StateID
      INTO 10),
step1 AS (SELECT *
      FROM input2
      PARTITION BY StateID
      INTO 10)

SELECT *
INTO output
FROM step1
PARTITION BY StateID
UNION step2
  BY StateID
```


For each of the following statements, select Yes if the statement is true. Otherwise, select No.

NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Statements	Yes	No
The query joins two streams of partitioned data.	○	○
The stream scheme key and count must match the output scheme.	○	○
Providing 60 streaming units will optimize the performance of the query.	○	○

Correct Answer:

Answer Area

Statements	Yes	No
The query joins two streams of partitioned data.	○	○
The stream scheme key and count must match the output scheme.	○	○
Providing 60 streaming units will optimize the performance of the query.	○	○

Box 1: Yes

You can now use a new extension of Azure Stream Analytics SQL to specify the number of partitions of a stream when reshuffling the data.

The outcome is a stream that has the same partition scheme. Please see below for an example:

WITH step1 AS (SELECT * FROM [input1] PARTITION BY DeviceID INTO 10), step2 AS (SELECT * FROM [input2] PARTITION BY DeviceID INTO 10) SELECT * INTO [output] FROM step1 PARTITION BY DeviceID UNION step2 PARTITION BY DeviceID

Note: The new extension of Azure Stream Analytics SQL includes a keyword INTO that allows you to specify the number of partitions for a stream when performing reshuffling using a PARTITION BY statement.

Box 2: Yes

When joining two streams of data explicitly repartitioned, these streams must have the same partition key and partition count.

Box 3: Yes

10 partitions x six SUs = 60 SUs is fine.

Note: Remember, Streaming Unit (SU) count, which is the unit of scale for Azure Stream Analytics, must be adjusted so the number of physical resources available to the job can fit the partitioned flow. In general, six SUs is a good number to

assign to each partition. In case there are insufficient resources assigned to the job, the system will only apply the repartition if it benefits the job.

Reference: https://azure.microsoft.com/en-in/blog/maximize-throughput-with-repartitioning-in-azure-stream-analytics/

QUESTION 4

HOTSPOT

You are implementing mapping data flows in Azure Data Factory to convert daily logs of taxi records into aggregated datasets.

You configure a data flow and receive the error shown in the following exhibit.

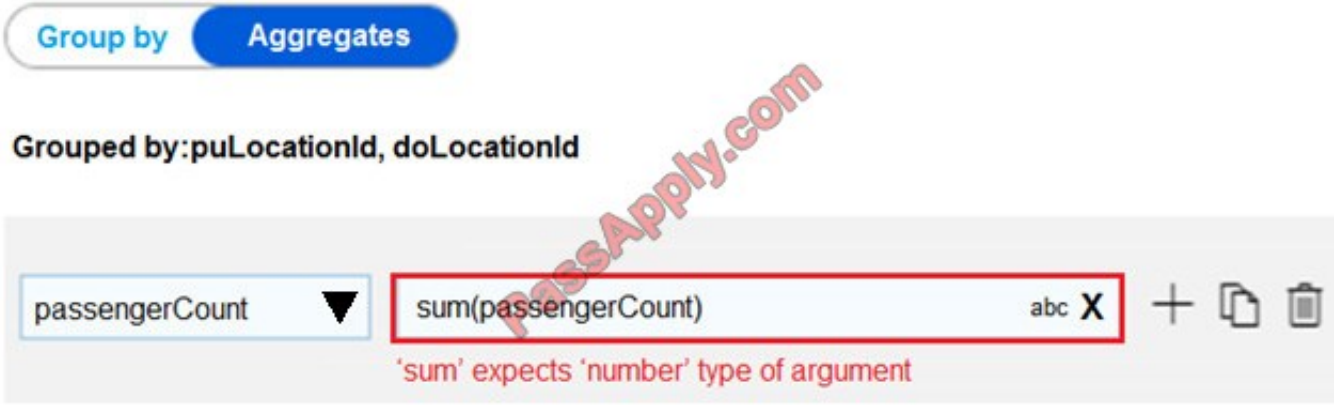

You need to resolve the error.

Which setting should you configure? To answer, select the appropriate setting in the answer area.

Hot Area:

Answer Area

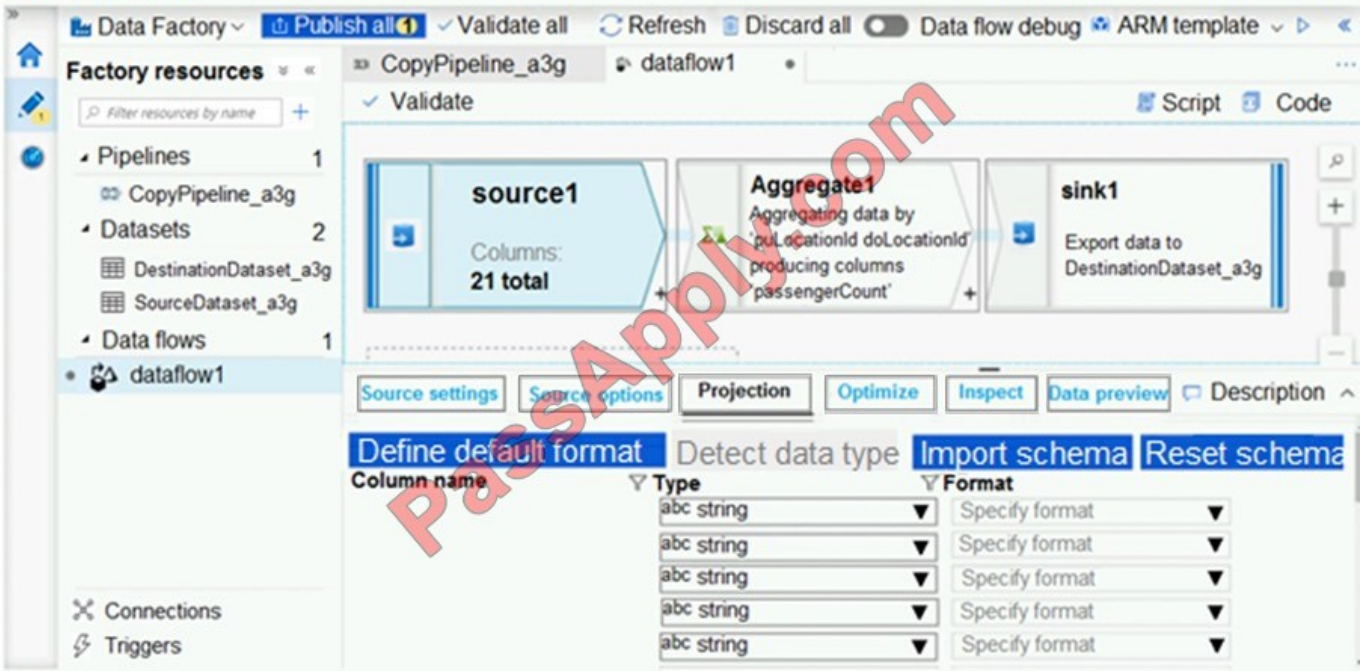

Correct Answer:

Answer Area

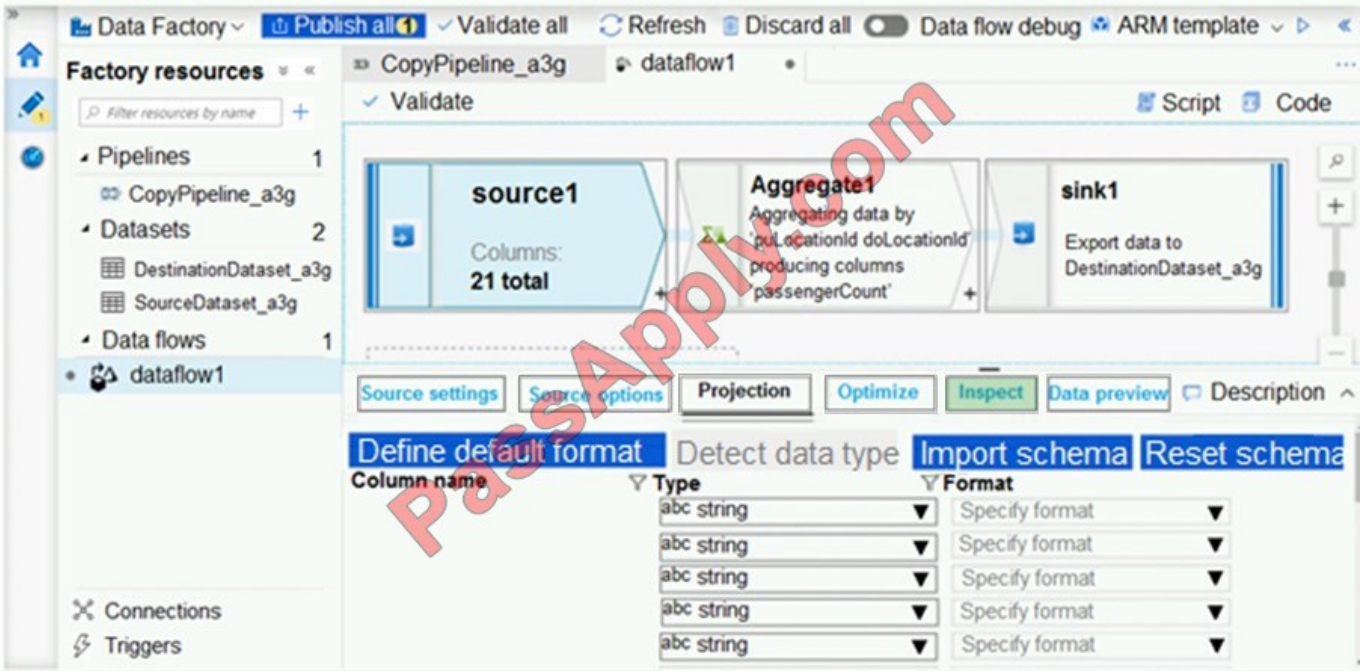

The Inspect tab provides a view into the metadata of the data stream that you\\'re transforming. You can see column counts, the columns changed, the columns added, data types, the column order, and column references. Inspect is a read-only view of your metadata. You don\\'t need to have debug mode enabled to see metadata in the Inspect pane.

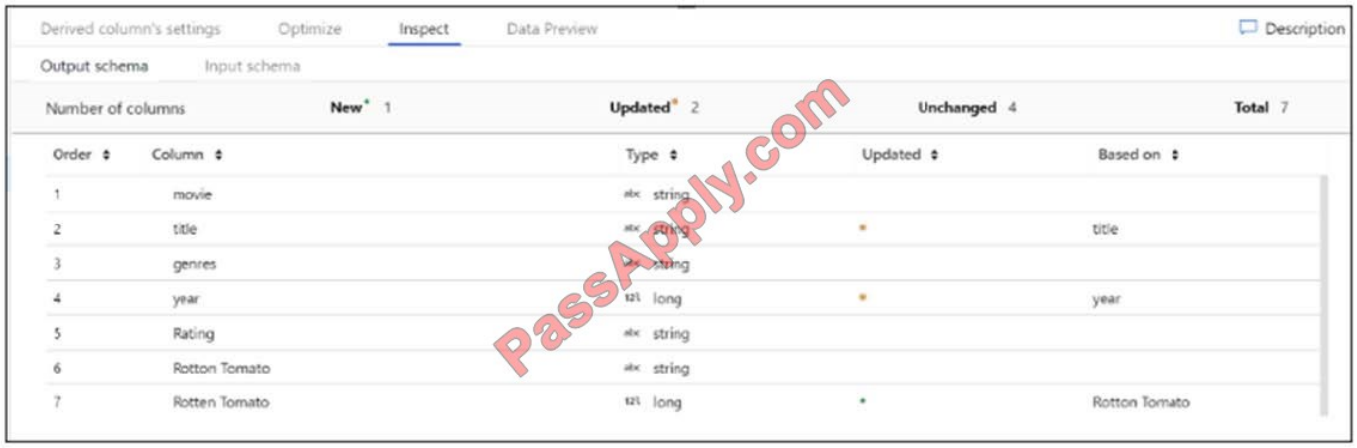

Reference: https://docs.microsoft.com/en-us/azure/data-factory/concepts-data-flow-overview

QUESTION 5

You have an Azure subscription that contains an Azure Data Factory version 2 (V2) data factory named df1. Df1 contains a linked service.

You have an Azure Key vault named vault1 that contains an encryption key named key1.

You need to encrypt df1 by using key1.

What should you do first?

A. Disable purge protection on vault1.

B. Create a self-hosted integration runtime.

C. Disable soft delete on vault1.

D. Remove the linked service from df1.

Correct Answer: D

Linked services are much like connection strings, which define the connection information needed for Data Factory to connect to external resources.

Incorrect Answers:

A, C: Data Factory requires two properties to be set on the Key Vault, Soft Delete and Do Not Purge

B: A self-hosted integration runtime copies data between an on-premises store and cloud storage.

Reference: https://docs.microsoft.com/en-us/azure/data-factory/enable-customer-managed-key https://docs.microsoft.com/en-us/azure/data-factory/concepts-linked-services https://docs.microsoft.com/en-us/azure/data-factory/create-self-hosted-integration-runtime

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.passapply.com/allproducts

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket: