



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.lead4pass.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

You manage an enterprise data warehouse in Azure Synapse Analytics.

Users report slow performance when they run commonly used queries. Users do not report performance changes for infrequently used queries.

You need to monitor resource utilization to determine the source of the performance issues.

Which metric should you monitor?

- A. Cache used percentage
- B. Local tempdb percentage
- C. DWU percentage
- D. CPU percentage
- E. Data IO percentage

Correct Answer: A

The Azure Synapse Analytics storage architecture automatically tiers your most frequently queried columnstore segments in a cache residing on NVMe based SSDs designed for Gen2 data warehouses. Greater performance is realized when your queries retrieve segments that are residing in the cache. You can monitor and troubleshoot slow query performance by determining whether your workload is optimally leveraging the Gen2 cache.

Note: As of November 2019, Azure SQL Data Warehouse is now Azure Synapse Analytics References:
<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/sql-data-warehouse-how-to-monitor-cache>

<https://docs.microsoft.com/bs-latn-ba/azure/sql-data-warehouse/sql-data-warehouse-concept-resource-utilization-query-activity>

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 apply an Azure Blob storage lifecycle policy.

Does this meet the goal?



A. Yes

B. No

Correct Answer: A

Azure Blob storage lifecycle management offers a rich, rule-based policy for GPv2 and Blob storage accounts. Use the policy to transition your data to the appropriate access tiers or expire at the end of the data's lifecycle. The lifecycle management policy lets you:

1.

Transition blobs to a cooler storage tier (hot to cool, hot to archive, or cool to archive) to optimize for performance and cost

2.

Delete blobs at the end of their lifecycles

3.

Define rules to be run once per day at the storage account level

4.

Apply rules to containers or a subset of blobs (using prefixes as filters)

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

QUESTION 3

You plan to build a structured streaming solution in Azure Databricks. The solution will count new events in five-minute intervals and report only events that arrive during the interval. The output will be sent to a Delta Lake table.

Which output mode should you use?

A. complete

B. update

C. append

Correct Answer: C

Append Mode: Only new rows appended in the result table since the last trigger are written to external storage. This is applicable only for the queries where existing rows in the Result Table are not expected to change.

Incorrect Answers:

A: Complete Mode: The entire updated result table is written to external storage. It is up to the storage connector to decide how to handle the writing of the entire table.

B: Update Mode: Only the rows that were updated in the result table since the last trigger are written to external storage. This is different from Complete Mode in that Update Mode outputs only the rows that have changed since the last



trigger.

If the query doesn't contain aggregations, it is equivalent to Append mode.

Reference: <https://docs.databricks.com/getting-started/spark/streaming.html>

QUESTION 4

HOTSPOT

You have the following Azure Stream Analytics query.

WITH

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

```
step2 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.	<input type="radio"/>	<input type="radio"/>
The stream scheme key and count must match the output scheme.	<input type="radio"/>	<input type="radio"/>
Providing 60 streaming units will optimize the performance of the query.	<input type="radio"/>	<input type="radio"/>

Correct Answer:

Answer Area

Statements	Yes	No
The query joins two streams of partitioned data.	<input checked="" type="radio"/>	<input type="radio"/>
The stream scheme key and count must match the output scheme.	<input checked="" type="radio"/>	<input type="radio"/>
Providing 60 streaming units will optimize the performance of the query.	<input checked="" type="radio"/>	<input type="radio"/>

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 5

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 multiple pipelines in a new Azure Data Factory V2.

You need to create the data factory, and then create a scheduled trigger for the planned pipelines. The trigger must execute every two hours starting at 24:00:00.

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

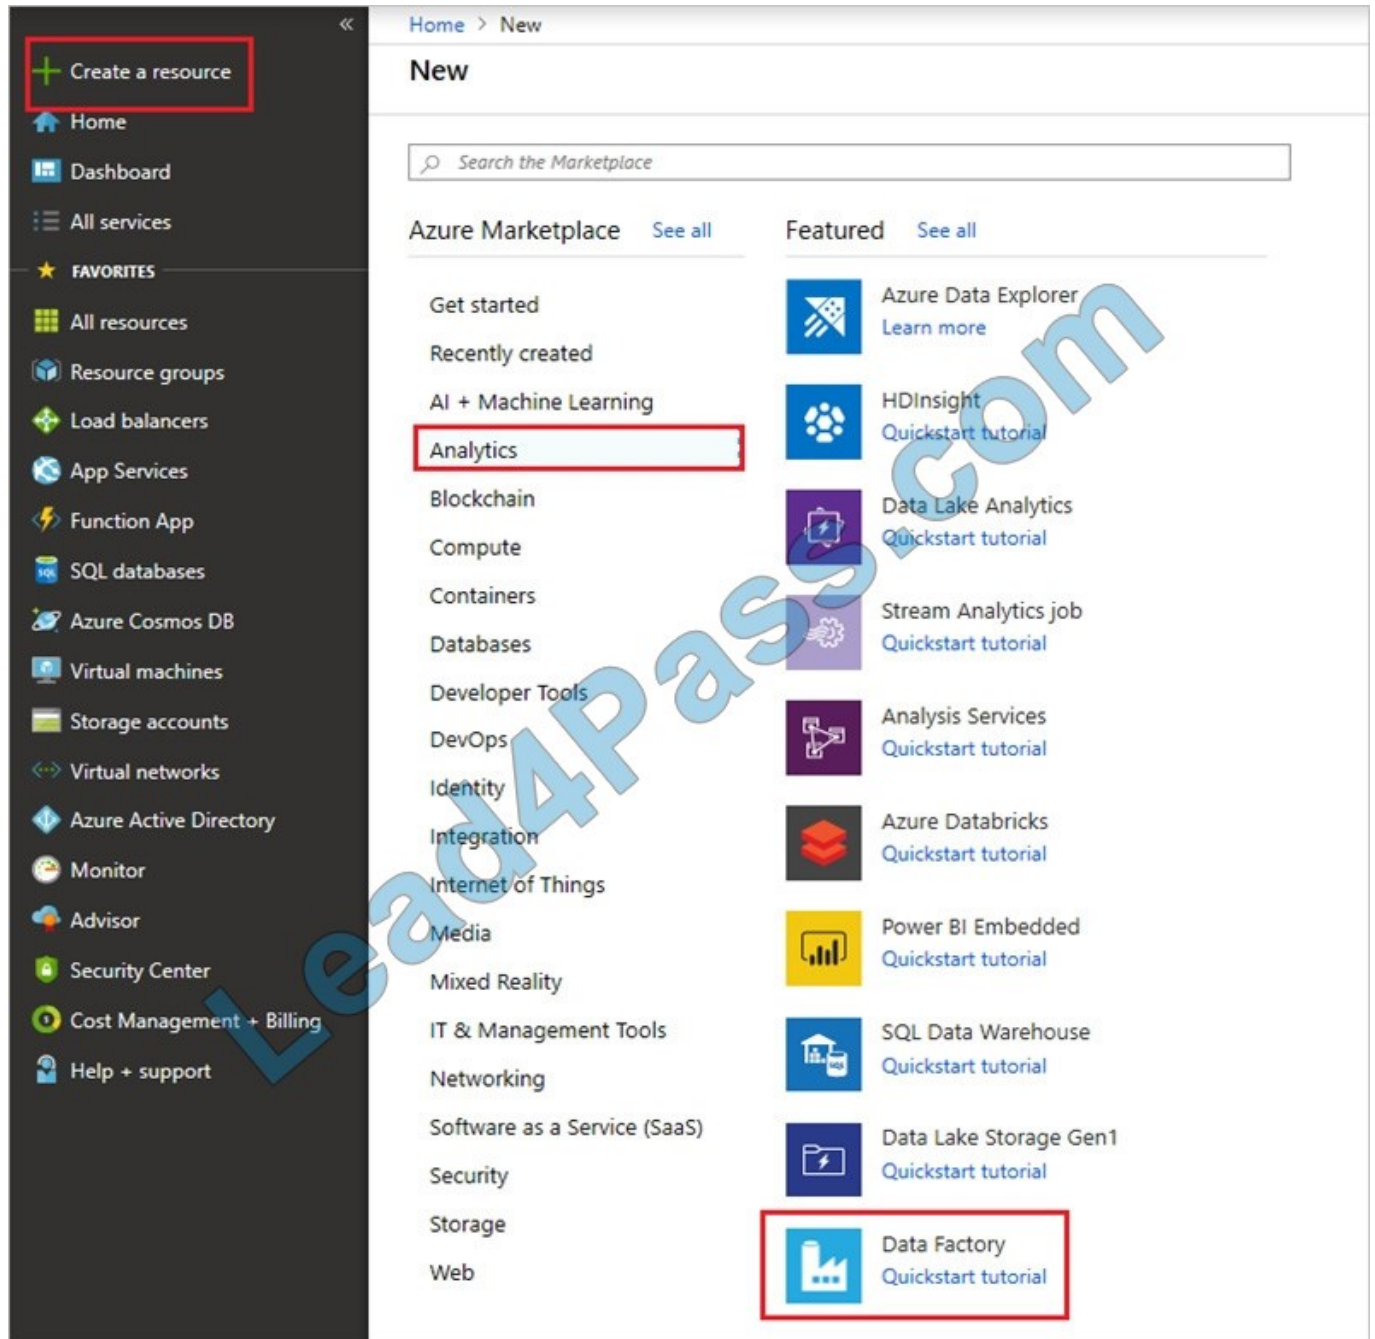
Correct Answer: See the below.



Explanation:

Step 1: Create a new Azure Data Factory V2

1. Go to the Azure portal.



2. Select Create a resource on the left menu, select Analytics, and then select Data Factory. 4. On the New data factory page, enter a name.

5.

For Subscription, select your Azure subscription in which you want to create the data factory.

6.



For Resource Group, use one of the following steps:

Select Use existing, and select an existing resource group from the list.

Select Create new, and enter the name of a resource group.

7.

For Version, select V2.

8.

For Location, select the location for the data factory.

9.

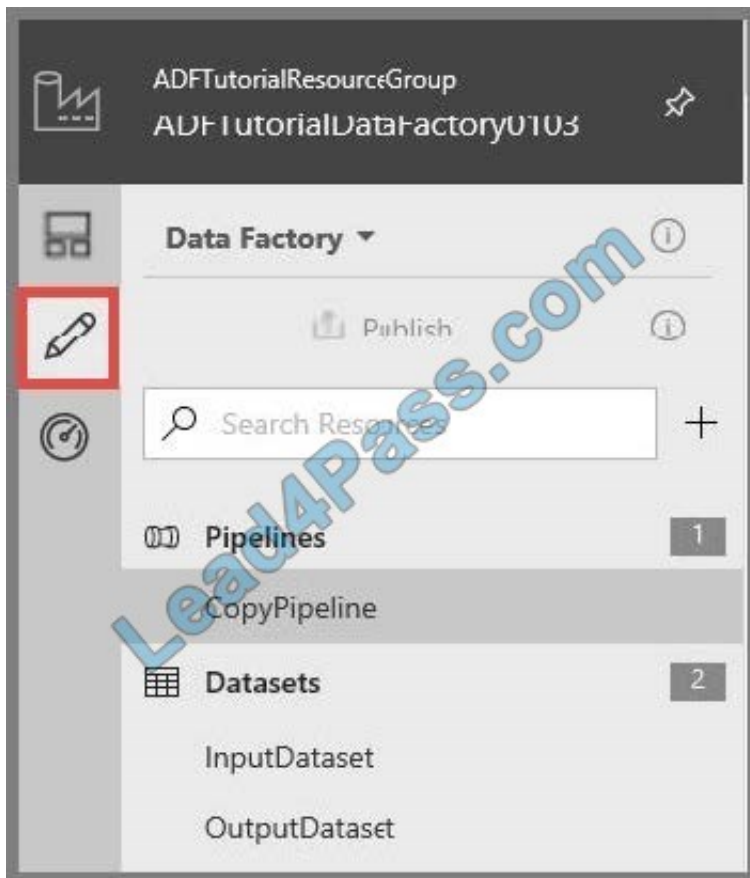
Select Create.

10.

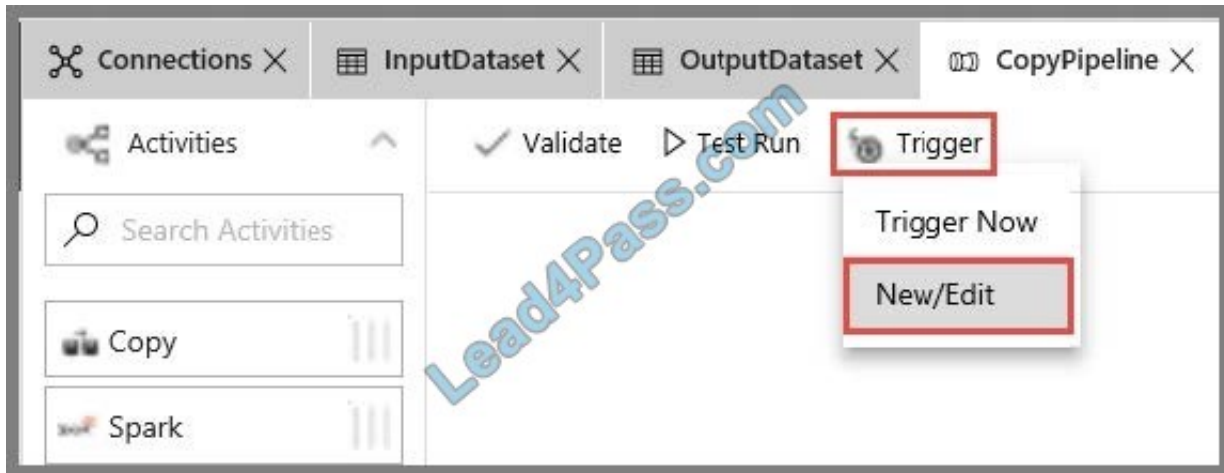
After the creation is complete, you see the Data Factory page.

Step 2: Create a schedule trigger for the Data Factory

1. Select the Data Factory you created, and switch to the Edit tab.



2. Click Trigger on the menu, and click New/Edit.



3.

In the Add Triggers page, click Choose trigger..., and click New.

4.

In the New Trigger page, do the following steps:



a.

Confirm that Schedule is selected for Type.

b.

Specify the start datetime of the trigger for Start Date (UTC) to: 24:00:00

c.

Specify Recurrence for the trigger. Select Every Hour, and enter 2 in the text box.



Type *

Schedule Tumbling Window

Start Date (UTC) * (i)

01/03/2018, 10:54 PM

Recurrence * (i)

Every Minute Every 1 Minute(s)

End * (i)

1 On Date

End On (UTC) * (i)

01/03/2018, 11:10 PM

January 2018

Sun	Mon	Tue	Wed	Thu	Fri	Sat
31	1	2	3	4	5	6
7	8	9	10	11	12	13
14	15	16	17	18	19	20
21	22	23	24	25	26	27
28	29	30	31	1	2	3
4	5	6	7	8	9	10

2 11 : 10 PM

3 Can Apply

Next

5.

In the New Trigger window, check the Activated option, and click Next.


6.



In the New Trigger page, review the warning message, and click Finish.

7.

Click Publish to publish changes to Data Factory. Until you publish changes to Data Factory, the trigger does not start triggering the pipeline runs.



References: <https://docs.microsoft.com/en-us/azure/data-factory/quickstart-create-data-factory-portal>

<https://docs.microsoft.com/en-us/azure/data-factory/how-to-create-schedule-trigger>

[Latest DP-200 Dumps](#)

[DP-200 Practice Test](#)

[DP-200 Exam Questions](#)



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.lead4pass.com/allproducts>

Need Help

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



 <p>One Year Free Update 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>	 <p>Money Back Guarantee To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p>Security & Privacy 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 © lead4pass, All Rights Reserved.