



DP-201^{Q&As}

Designing an Azure Data Solution

Pass Microsoft DP-201 Exam with 100% Guarantee

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

<https://www.lead4pass.com/dp-201.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

HOTSPOT

Which Azure service and feature should you recommend using to manage the transient data for Data Lake Storage? To answer, select the appropriate options in the answer area. NOTE: Each correct selection is worth one point.

Hot Area:

Answer Area

Service:

- Azure Data Factory
- Azure Storage
- Azure SQL Data Warehouse

Feature:

- Delete Activity
- DROP EXTERNAL TABLE
- Lifecycle management rule

Correct Answer:



Answer Area

Service:

Feature:

Azure Data Factory
Azure Storage
Azure SQL Data Warehouse

Delete Activity
DROP EXTERNAL TABLE
Lifecycle management rule

Scenario: Stage inventory data in Azure Data Lake Storage Gen2 before loading the data into the analytical data store. Litware wants to remove transient data from Data Lake Storage once the data is no longer in use. Files that have a modified date that is older than 14 days must be removed.

Service: Azure Data Factory

Clean up files by built-in delete activity in Azure Data Factory (ADF).

ADF built-in delete activity, which can be part of your ETL workflow to deletes undesired files without writing code. You can use ADF to delete folder or files from Azure Blob Storage, Azure Data Lake Storage Gen1, Azure Data Lake Storage

Gen2, File System, FTP Server, sFTP Server, and Amazon S3.

You can delete expired files only rather than deleting all the files in one folder. For example, you may want to only delete the files which were last modified more than 13 days ago.

Feature: Delete Activity

Reference:

<https://azure.microsoft.com/sv-se/blog/clean-up-files-by-built-in-delete-activity-in-azure-data-factory/>

QUESTION 2

You need to recommend a solution for storing the image tagging data. What should you recommend?



- A. Azure File Storage
- B. Azure Cosmos DB
- C. Azure Blob Storage
- D. Azure SQL Database
- E. Azure SQL Data Warehouse

Correct Answer: C

Image data must be stored in a single data store at minimum cost.

Note: Azure Blob storage is Microsoft's object storage solution for the cloud. Blob storage is optimized for storing massive amounts of unstructured data. Unstructured data is data that does not adhere to a particular data model or definition,

such as text or binary data.

Blob storage is designed for:

1.
Serving images or documents directly to a browser.
2.
Storing files for distributed access.
3.
Streaming video and audio.
4.
Writing to log files.
5.
Storing data for backup and restore, disaster recovery, and archiving.
6.
Storing data for analysis by an on-premises or Azure-hosted service.

References: <https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blobs-introduction>

QUESTION 3

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 streaming data that is received by Azure Event Hubs and stored in Azure Blob storage. The data contains social media posts that relate to a keyword of Contoso.



You need to count how many times the Contoso keyword and a keyword of Litware appear in the same post every 30 seconds. The data must be available to Microsoft Power BI in near real-time.

Solution: You use Azure Data Factory and an event trigger to detect when new blobs are created. You use mapping data flows in Azure Data Factory to aggregate and filter the data, and then send the data to an Azure SQL database. You

consume the data in Power BI by using DirectQuery mode.

Does the solution meet the goal?

A. Yes

B. No

Correct Answer: B

QUESTION 4

You manage a process that performs analysis of daily web traffic logs on an HDInsight cluster. Each of the 250 web servers generates approximately 10 megabytes (MB) of log data each day. All log data is stored in a single folder in Microsoft

Azure Data Lake Storage Gen 2.

You need to improve the performance of the process.

Which two changes should you make? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

A. Combine the daily log files for all servers into one file

B. Increase the value of the mapreduce.map.memory parameter

C. Move the log files into folders so that each day's logs are in their own folder

D. Increase the number of worker nodes

E. Increase the value of the hive.tez.container.size parameter

Correct Answer: AC

A: Typically, analytics engines such as HDInsight and Azure Data Lake Analytics has a per-file overhead. If you store your data as many small files, this can negatively affect performance. In general, organize your data into larger sized files

for better performance (256MB to 100GB in size). Some engines and applications might have trouble efficiently processing files that are greater than 100GB in size.

C: For Hive workloads, partition pruning of time-series data can help some queries read only a subset of the data which improves performance. Those pipelines that ingest time-series data, often place their files with a very structured naming



for files and folders. Below is a very common example we see for data is structured by date:

```
\\DataSet\YYYY\MM\DD\datafile_YYYY_MM_DD.tsv
```

Notice that the datetime information appears both as folders and in the filename.

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-performance-tuning-guidance>

QUESTION 5

You need to recommend an Azure SQL Database service tier.

What should you recommend?

- A. Business Critical
- B. General Purpose
- C. Premium
- D. Standard
- E. Basic

Correct Answer: C

The data engineers must set the SQL Data Warehouse compute resources to consume 300 DWUs.

Note: There are three architectural models that are used in Azure SQL Database:

1.
General Purpose/Standard
2.
Business Critical/Premium
3.
Hyperscale

Incorrect Answers:

A: Business Critical service tier is designed for the applications that require low-latency responses from the underlying SSD storage (1-2 ms in average), fast recovery if the underlying infrastructure fails, or need to off-load reports, analytics, and read-only queries to the free of charge readable secondary replica of the primary database.

References: <https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tier-business-critical>



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.