



# DP-201<sup>Q&As</sup>

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.passapply.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**

You plan to create a real-time monitoring app that alerts users when a device travels more than 200 meters away from a designated location.

You need to design an Azure Stream Analytics job to process the data for the planned app. The solution must minimize the amount of code developed and the number of technologies used.

What should you include in the Stream Analytics job? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:

**Answer Area**

Input type:   
Stream  
Reference

Input source:   
Azure IoT Hub  
Azure Event Hubs  
Azure Blob storage

Function:   
Aggregate  
Geospatial  
Windowing

Correct Answer:



## Answer Area

Input type:

Input source:

Function:

Input type: Stream

You can process real-time IoT data streams with Azure Stream Analytics.

Input source: Azure IoT Hub

In a real-world scenario, you could have hundreds of these sensors generating events as a stream. Ideally, a gateway device would run code to push these events to Azure Event Hubs or Azure IoT Hubs.

Function: Geospatial

With built-in geospatial functions, you can use Azure Stream Analytics to build applications for scenarios such as fleet management, ride sharing, connected cars, and asset tracking.

Reference:

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-get-started-with-azure-stream-analytics-to-process-data-from-iot-devices>

<https://docs.microsoft.com/en-us/azure/stream-analytics/geospatial-scenarios>



## QUESTION 2

You plan to store 100 GB of data used by a line-of-business (LOB) app.

You need to recommend a data storage solution for the data. The solution must meet the following requirements:

1.  
Minimize storage costs.
2.  
Natively support relational queries.
3.  
Provide a recovery time objective (RTO) of less than one minute. What should you include in the recommendation?

- A. Azure Cosmos DB
- B. Azure SQL Database
- C. Azure SQL Data Warehouse
- D. Azure Blob storage

Correct Answer: D

Incorrect Answers:

A: Azure Cosmos DB would require an SQL API.

---

## QUESTION 3

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

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 are designing an HDInsight/Hadoop cluster solution that uses Azure Data Lake Gen1 Storage.

The solution requires POSIX permissions and enables diagnostics logging for auditing.

You need to recommend solutions that optimize storage.

Proposed Solution: Ensure that files stored are larger than 250MB.

Does the solution meet the goal?

- A. Yes



B. No

Correct Answer: A

Depending on what services and workloads are using the data, a good size to consider for files is 256 MB or greater. If the file sizes cannot be batched when landing in Data Lake Storage Gen1, you can have a separate compaction job that combines these files into larger ones.

Note: POSIX permissions and auditing in Data Lake Storage Gen1 comes with an overhead that becomes apparent when working with numerous small files. As a best practice, you must batch your data into larger files versus writing thousands or millions of small files to Data Lake Storage Gen1. Avoiding small file sizes can have multiple benefits, such as:

1.

Lowering the authentication checks across multiple files

2.

Reduced open file connections

3.

Faster copying/replication

4.

Fewer files to process when updating Data Lake Storage Gen1 POSIX permissions

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

---

#### QUESTION 4

##### HOTSPOT

You design data engineering solutions for a company.

You must integrate on-premises SQL Server data into an Azure solution that performs Extract-Transform-Load (ETL) operations have the following requirements:

1. Develop a pipeline that can integrate data and run notebooks.
2. Develop notebooks to transform the data.
3. Load the data into a massively parallel processing database for later analysis.

You need to recommend a solution.

What should you recommend? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:



## Answer Area

### Requirement

### Service

Integrate the on-premises data into the cloud.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	

Develop notebooks to transform the data.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	

Run notebooks.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	

Load the data.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	

Store the transformed data.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	

Correct Answer:



## Answer Area

### Requirement

### Service

Integrate the on-premises data into the cloud.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	

Develop notebooks to transform the data.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	

Run notebooks.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	

Load the data.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	

Store the transformed data.

	▼
Azure Databricks	
Azure Data Factory	
Azure SQL Data Warehouse	
Azure Batch	



### QUESTION 5

You are designing an application. You plan to use Azure SQL Database to support the application.

The application will extract data from the Azure SQL Database and create text documents. The text documents will be placed into a cloud-based storage solution. The text storage solution must be accessible from an SMB network share.

You need to recommend a data storage solution for the text documents.

Which Azure data storage type should you recommend?

- A. Queue
- B. Files
- C. Blob
- D. Table

Correct Answer: B

Azure Files enables you to set up highly available network file shares that can be accessed by using the standard Server Message Block (SMB) protocol. Incorrect Answers:

A: The Azure Queue service is used to store and retrieve messages. It is generally used to store lists of messages to be processed asynchronously.

C: Blob storage is optimized for storing massive amounts of unstructured data, such as text or binary data. Blob storage can be accessed via HTTP or HTTPS but not via SMB.

D: Azure Table storage is used to store large amounts of structured data. Azure tables are ideal for storing structured, non-relational data.

References: <https://docs.microsoft.com/en-us/azure/storage/common/storage-introduction> <https://docs.microsoft.com/en-us/azure/storage/tables/table-storage-overview>

---

### QUESTION 6

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 create an Azure Stream Analytics job that uses an input from Event Hubs to count the posts that have the specified keywords, then and 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: A

Reference: <https://docs.microsoft.com/en-us/power-bi/service-real-time-streaming>

<https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-twitter-sentiment-analysis-trends>

---

### QUESTION 7

You need to optimize storage for CONT\_SQL3.

What should you recommend?

- A. AlwaysOn
- B. Transactional processing
- C. General
- D. Data warehousing

Correct Answer: B

CONT\_SQL3 with the SQL Server role, 100 GB database size, Hyper-VM to be migrated to Azure VM. The storage should be configured to optimized storage for database OLTP workloads.

Azure SQL Database provides three basic in-memory based capabilities (built into the underlying database engine) that can contribute in a meaningful way to performance improvements:

In-Memory Online Transactional Processing (OLTP) Clustered columnstore indexes intended primarily for Online Analytical Processing (OLAP) workloads Nonclustered columnstore indexes geared towards Hybrid Transactional/Analytical Processing (HTAP) workloads

References: <https://www.databasejournal.com/features/mssql/overview-of-in-memory-technologies-of-azure-sql-database.html>

---

### QUESTION 8

HOTSPOT

You need to design storage for the solution.

Which storage services should you recommend? To answer, select the appropriate configuration in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:



## Answer Area

**Solution component**

**Storage service**

**Images**

	▼
Azure Blob Storage	
Azure Data Lake Storage	
Azure SQL Database	
Azure SQL Data Warehouse	

**Customer data**

	▼
Azure Blob Storage	
Azure Cosmos DB	
Azure SQL Database	
Azure SQL Data Warehouse	

Correct Answer:



## Answer Area

**Solution component**

**Storage service**

**Images**

	▼
Azure Blob Storage	
Azure Data Lake Storage	
Azure SQL Database	
Azure SQL Data Warehouse	

**Customer data**

	▼
Azure Blob Storage	
Azure Cosmos DB	
Azure SQL Database	
Azure SQL Data Warehouse	

Images: Azure Data Lake Storage

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

Customer data: Azure Blob Storage

Scenario: Customer data must be analyzed using managed Spark clusters.

Spark clusters in HDInsight are compatible with Azure Storage and Azure Data Lake Storage.

Azure Storage includes these data services: Azure Blob, Azure Files, Azure Queues, and Azure Tables.

References:

<https://docs.microsoft.com/en-us/azure/hdinsight/spark/apache-spark-overview>

### QUESTION 9

You need to recommend a solution to quickly identify all the columns in Health Review that contain sensitive health data. What should you include in the recommendation?

- A. classifications
- B. data masking



- C. SQL Server auditing
- D. Azure tags

Correct Answer: A

Data Discovery and Classification introduces a set of advanced capabilities aimed at protecting data and not just the data warehouse itself. Classification/Labeling

---

#### QUESTION 10

You plan to use an Azure SQL data warehouse to store the customer data. You need to recommend a disaster recovery solution for the data warehouse. What should you include in the recommendation?

- A. AzCopy
- B. Read-only replicas
- C. AdlCopy
- D. Geo-Redundant backups

Correct Answer: D

<https://docs.microsoft.com/en-us/azure/sql-data-warehouse/backup-and-restore>

---

#### QUESTION 11

You need to design the image processing and storage solutions.

What should you recommend? To answer, select the appropriate configuration in the answer area.

NOTE: Each correct selection is worth one point.

Hot Area:



Image processing

- Azure HDInsight
- Azure Databricks
- Azure Batch
- Azure Cognitive Services

data storage for tagging data

- Azure Blob Storage
- Azure Table Storage
- Azure Cosmos DB
- Azure SQL Database

Correct Answer:

Image processing

- Azure HDInsight
- Azure Databricks
- Azure Batch
- Azure Cognitive Services

data storage for tagging data

- Azure Blob Storage
- Azure Table Storage
- Azure Cosmos DB
- Azure SQL Database

From the scenario:



The company identifies the following business requirements:

1.

You must transfer all images and customer data to cloud storage and remove on-premises servers.

2.

You must develop an image object and color tagging solution.

The solution has the following technical requirements:

1.

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

2.

All data must be backed up in case disaster recovery is required.

All cloud data must be encrypted at rest and in transit. The solution must support: hyper-scale storage of images

References: <https://docs.microsoft.com/en-us/azure/architecture/data-guide/technology-choices/batch-processing>  
<https://docs.microsoft.com/en-us/azure/sql-database/sql-database-service-tier-hyperscale>

---

## QUESTION 12

What should you do to improve high availability of the real-time data processing solution?

- A. Deploy identical Azure Stream Analytics jobs to paired regions in Azure.
- B. Deploy a High Concurrency Databricks cluster.
- C. Deploy an Azure Stream Analytics job and use an Azure Automation runbook to check the status of the job and to start the job if it stops.
- D. Set Data Lake Storage to use geo-redundant storage (GRS).

Correct Answer: A

Guarantee Stream Analytics job reliability during service updates Part of being a fully managed service is the capability to introduce new service functionality and improvements at a rapid pace. As a result, Stream Analytics can have a service update deploy on a weekly (or more frequent) basis. No matter how much testing is done there is still a risk that an existing, running job may break due to the introduction of a bug. If you are running mission critical jobs, these risks need to be avoided. You can reduce this risk by following Azure's paired region model.

Scenario: The application development team will create an Azure event hub to receive real-time sales data, including store number, date, time, product ID, customer loyalty number, price, and discount amount, from the point of sale (POS) system and output the data to data storage in Azure

Reference: <https://docs.microsoft.com/en-us/azure/stream-analytics/stream-analytics-job-reliability>



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:



	<b>One Year Free Update</b> 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.		<b>Money Back Guarantee</b> To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.		<b>Security &amp; Privacy</b> We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.
--	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--	-----------------------------------------------------------------------------------------------------------------------------------------------------------------	--	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Any charges made through this site will appear as Global Simulators Limited.  
All trademarks are the property of their respective owners.  
Copyright © passapply, All Rights Reserved.