



---

# MICROSOFT DP-203

---

**Microsoft Data Engineering on Microsoft Azure Certification Questions  
& Answers**

---

Exam Summary – Syllabus – Questions

---

**DP-203**

**[Microsoft Certified - Azure Data Engineer Associate](#)**

**40-60 Questions Exam - 700/1000 Cut Score - Duration of 150 minutes**

## Table of Contents:

Know Your DP-203 Certification Well: .....	2
Microsoft DP-203 Data Engineering on Microsoft Azure Certification Details: .....	2
DP-203 Syllabus:.....	3
Design and Implement Data Storage (40-45%) .....	3
Design and Develop Data Processing (25-30%) .....	4
Design and Implement Data Security (10-15%) .....	5
Monitor and Optimize Data Storage and Data Processing (10-15%) .....	6
Microsoft DP-203 Sample Questions: .....	7
Study Guide to Crack Microsoft Data Engineering on Microsoft Azure DP-203 Exam: .....	11

## Know Your DP-203 Certification Well:

The DP-203 is best suitable for candidates who want to gain knowledge in the Microsoft Azure. Before you start your DP-203 preparation you may struggle to get all the crucial Data Engineering on Microsoft Azure materials like DP-203 syllabus, sample questions, study guide.

But don't worry the DP-203 PDF is here to help you prepare in a stress free manner.

The PDF is a combination of all your queries like-

- What is in the DP-203 syllabus?
- How many questions are there in the DP-203 exam?
- Which Practice test would help me to pass the DP-203 exam at the first attempt?

Passing the DP-203 exam makes you Microsoft Certified - Azure Data Engineer Associate. Having the Data Engineering on Microsoft Azure certification opens multiple opportunities for you. You can grab a new job, get a higher salary or simply get recognition within your current organization.

## Microsoft DP-203 Data Engineering on Microsoft Azure Certification Details:

Exam Name	Microsoft Certified - Azure Data Engineer Associate
Exam Code	DP-203
Exam Price	\$165 (USD)
Duration	150 mins
Number of Questions	40-60
Passing Score	700 / 1000
Books / Training	<a href="#">DP-203T00: Data Engineering on Microsoft Azure</a>
Schedule Exam	<a href="#">Pearson VUE</a>
Sample Questions	<a href="#">Data Engineering on Microsoft Azure Sample Questions</a>
Practice Exam	<a href="#">Microsoft DP-203 Certification Practice Exam</a>

## DP-203 Syllabus:

Topic	Details
<b>Design and Implement Data Storage (40-45%)</b>	
Design a data storage structure	<ul style="list-style-type: none"> <li>- design an Azure Data Lake solution</li> <li>- recommend file types for storage</li> <li>- recommend file types for analytical queries</li> <li>- design for efficient querying</li> <li>- design for data pruning</li> <li>- design a folder structure that represents the levels of data transformation</li> <li>- design a distribution strategy</li> <li>- design a data archiving solution</li> </ul>
Design a partition strategy	<ul style="list-style-type: none"> <li>- design a partition strategy for files</li> <li>- design a partition strategy for analytical workloads</li> <li>- design a partition strategy for efficiency/performance</li> <li>- design a partition strategy for Azure Synapse Analytics</li> <li>- identify when partitioning is needed in Azure Data Lake Storage Gen2</li> </ul>
Design the serving layer	<ul style="list-style-type: none"> <li>- design star schemas</li> <li>- design slowly changing dimensions</li> <li>- design a dimensional hierarchy</li> <li>- design a solution for temporal data</li> <li>- design for incremental loading</li> <li>- design analytical stores</li> <li>- design metastores in Azure Synapse Analytics and Azure Databricks</li> </ul>
Implement physical data storage structures	<ul style="list-style-type: none"> <li>- implement compression</li> <li>- implement partitioning</li> <li>- implement sharding</li> <li>- implement different table geometries with Azure Synapse Analytics pools</li> <li>- implement data redundancy</li> <li>- implement distributions</li> <li>- implement data archiving</li> </ul>
Implement logical data structures	<ul style="list-style-type: none"> <li>- build a temporal data solution</li> <li>- build a slowly changing dimension</li> <li>- build a logical folder structure</li> <li>- build external tables</li> <li>- implement file and folder structures for efficient querying and data pruning</li> </ul>

Topic	Details
Implement the serving layer	<ul style="list-style-type: none"> <li>- deliver data in a relational star schema</li> <li>- deliver data in Parquet files</li> <li>- maintain metadata</li> <li>- implement a dimensional hierarchy</li> </ul>
<p><b>Design and Develop Data Processing (25-30%)</b></p>	
Ingest and transform data	<ul style="list-style-type: none"> <li>- transform data by using Apache Spark</li> <li>- transform data by using Transact-SQL</li> <li>- transform data by using Data Factory</li> <li>- transform data by using Azure Synapse Pipelines</li> <li>- transform data by using Stream Analytics</li> <li>- cleanse data</li> <li>- split data</li> <li>- shred JSON</li> <li>- encode and decode data</li> <li>- configure error handling for the transformation</li> <li>- normalize and denormalize values</li> <li>- transform data by using Scala</li> <li>- perform data exploratory analysis</li> </ul>
Design and develop a batch processing solution	<ul style="list-style-type: none"> <li>- develop batch processing solutions by using Data Factory, Data Lake, Spark, Azure Synapse Pipelines, PolyBase, and Azure Databricks</li> <li>- create data pipelines</li> <li>- design and implement incremental data loads</li> <li>- design and develop slowly changing dimensions</li> <li>- handle security and compliance requirements</li> <li>- scale resources</li> <li>- configure the batch size</li> <li>- design and create tests for data pipelines</li> <li>- integrate Jupyter/Python notebooks into a data pipeline</li> <li>- handle duplicate data</li> <li>- handle missing data</li> <li>- handle late-arriving data</li> <li>- upsert data</li> <li>- regress to a previous state</li> <li>- design and configure exception handling</li> <li>- configure batch retention</li> <li>- design a batch processing solution</li> <li>- debug Spark jobs by using the Spark UI</li> </ul>

Topic	Details
Design and develop a stream processing solution	<ul style="list-style-type: none"> <li>- develop a stream processing solution by using Stream Analytics, Azure Databricks, and Azure Event Hubs</li> <li>- process data by using Spark structured streaming</li> <li>- monitor for performance and functional regressions</li> <li>- design and create windowed aggregates</li> <li>- handle schema drift</li> <li>- process time series data</li> <li>- process across partitions</li> <li>- process within one partition</li> <li>- configure checkpoints/watermarking during processing</li> <li>- scale resources</li> <li>- design and create tests for data pipelines</li> <li>- optimize pipelines for analytical or transactional purposes</li> <li>- handle interruptions</li> <li>- design and configure exception handling</li> <li>- upsert data</li> <li>- replay archived stream data</li> <li>- design a stream processing solution</li> </ul>
Manage batches and pipelines	<ul style="list-style-type: none"> <li>- trigger batches</li> <li>- handle failed batch loads</li> <li>- validate batch loads</li> <li>- manage data pipelines in Data Factory/Synapse Pipelines</li> <li>- schedule data pipelines in Data Factory/Synapse Pipelines</li> <li>- implement version control for pipeline artifacts</li> <li>- manage Spark jobs in a pipeline</li> </ul>
<b>Design and Implement Data Security (10-15%)</b>	
Design security for data policies and standards	<ul style="list-style-type: none"> <li>- design data encryption for data at rest and in transit</li> <li>- design a data auditing strategy</li> <li>- design a data masking strategy</li> <li>- design for data privacy</li> <li>- design a data retention policy</li> <li>- design to purge data based on business requirements</li> <li>- design Azure role-based access control (Azure RBAC) and POSIX-like Access Control List (ACL) for Data Lake Storage Gen2</li> <li>- design row-level and column-level security</li> </ul>
Implement data security	<ul style="list-style-type: none"> <li>- implement data masking</li> <li>- encrypt data at rest and in motion</li> <li>- implement row-level and column-level security</li> <li>- implement Azure RBAC</li> </ul>

Topic	Details
	<ul style="list-style-type: none"> <li>- implement POSIX-like ACLs for Data Lake Storage Gen2</li> <li>- implement a data retention policy</li> <li>- implement a data auditing strategy</li> <li>- manage identities, keys, and secrets across different data platform technologies</li> <li>- implement secure endpoints (private and public)</li> <li>- implement resource tokens in Azure Databricks</li> <li>- load a DataFrame with sensitive information</li> <li>- write encrypted data to tables or Parquet files</li> <li>- manage sensitive information</li> </ul>
<p><b>Monitor and Optimize Data Storage and Data Processing (10-15%)</b></p>	
<p>Monitor data storage and data processing</p>	<ul style="list-style-type: none"> <li>- implement logging used by Azure Monitor</li> <li>- configure monitoring services</li> <li>- measure performance of data movement</li> <li>- monitor and update statistics about data across a system</li> <li>- monitor data pipeline performance</li> <li>- measure query performance</li> <li>- monitor cluster performance</li> <li>- understand custom logging options</li> <li>- schedule and monitor pipeline tests</li> <li>- interpret Azure Monitor metrics and logs</li> <li>- interpret a Spark directed acyclic graph (DAG)</li> </ul>
<p>Optimize and troubleshoot data storage and data processing</p>	<ul style="list-style-type: none"> <li>- compact small files</li> <li>- rewrite user-defined functions (UDFs)</li> <li>- handle skew in data</li> <li>- handle data spill</li> <li>- tune shuffle partitions</li> <li>- find shuffling in a pipeline</li> <li>- optimize resource management</li> <li>- tune queries by using indexers</li> <li>- tune queries by using cache</li> <li>- optimize pipelines for analytical or transactional purposes</li> <li>- optimize pipeline for descriptive versus analytical workloads</li> <li>- troubleshoot a failed spark job</li> <li>- troubleshoot a failed pipeline run</li> </ul>

## Microsoft DP-203 Sample Questions:

### Question: 1

A company is planning on creating an Azure SQL database to support a mission critical application. The application needs to be highly available and not have any performance degradation during maintenance windows.

Which of the following technologies can be used to implement this solution?

(Choose 3)

- a) Premium Service Tier
- b) Virtual Machine Scale Sets
- c) Basic Service Tier
- d) SQL Data Sync
- e) Always On Availability Groups
- f) Zone-redundant configuration

**Answer: a, e, f**

### Question: 2

An in-house team is developing a new application. The design document specifies that data should be represented using nodes and relationships in graph structures. Individual data elements are relatively small.

You need to recommend an appropriate data storage solution. Which solution should you recommend?

- a) Azure Storage Blobs
- b) Cosmos DB
- c) Azure Data Lake Store
- d) HBase in HDInsight

**Answer: b**

### Question: 3

Which offering provides scale-out parallel processing and dramatically accelerates performance of analytics clusters when integrated with the IBM Flash System?

- a) IBM Cloud Object Storage
- b) IBM Spectrum Accelerate
- c) IBM Spectrum Scale
- d) IBM Spectrum Connect

**Answer: c**



**Question: 4**

The data engineering team manages Azure HDInsight clusters. The team spends a large amount of time creating and destroying clusters daily because most of the data pipeline process runs in minutes.

You need to implement a solution that deploys multiple HDInsight clusters with minimal effort. What should you implement?

- a) Azure Databricks
- b) Azure Traffic Manager
- c) Azure Resource Manager templates
- d) Ambari web user interface

**Answer: c**

**Question: 5**

A company purchases IoT devices to monitor manufacturing machinery. The company uses an Azure IoT Hub to communicate with the IoT devices. The company must be able to monitor the devices in real-time.

You need to design the solution. What should you recommend?

- a) Azure Data Factory instance using Azure Portal
- b) Azure Analysis Services using Microsoft Visual Studio
- c) Azure Stream Analytics Edge application using Microsoft Visual Studio
- d) Azure Data Factory instance using Microsoft Visual Studio

**Answer: c**

**Question: 6**

A company is planning on creating an Azure SQL database to support a mission critical application. The application needs to be highly available and not have any performance degradation during maintenance windows.

Which of the following technologies can be used to implement this solution?

(Choose 3)

- a) Premium Service Tier
- b) Virtual Machine Scale Sets
- c) Basic Service Tier
- d) SQL Data Sync
- e) Always On Availability Groups
- f) Zone-redundant configuration

**Answer: a, e, f**

**Question: 7**

A company manages several on-premises Microsoft SQL Server databases. You need to migrate the databases to Microsoft Azure by using a backup process of Microsoft SQL Server. Which data technology should you use?

- a) Azure SQL Database single database
- b) Azure SQL Data Warehouse
- c) Azure Cosmos DB
- d) Azure SQL Database Managed Instance
- e) HDInsight Spark cluster

**Answer: d**

**Question: 8**

You are a data engineer for an Azure SQL Database. You write the following SQL statements:

```
CREATE TABLE Customer (  
CustomerID int IDENTITY PRIMARY KEY,  
GivenName varchar(100) MASKED WITH (FUNCTION = 'partial(2,"XX",0)') NULL,  
SurName varchar(100) NOT NULL,  
Phone varchar(12) MASKED WITH (FUNCTION = 'default()')  
INSERT Customer (GivenName, SurName, Phone) VALUES ('Sammy', 'Jack', '555.111.2222');  
SELECT * FROM Customer;
```

You need to determine what is returned by the SELECT query. What data is returned?

- a) 1 SaXX Jack XXX.XXX.2222
- b) 1 XXXX Jack XXX.XXX.XXXX
- c) 1 xx Jack XXX.XXX.2222
- d) 1 SaXX Jack xxxx

**Answer: d**

Question: 9

A company has an Azure SQL data warehouse. They want to use PolyBase to retrieve data from an Azure Blob storage account and ingest into the Azure SQL data warehouse. The files are stored in parquet format. The data needs to be loaded into a table called lead2pass\_sales. Which of the following actions need to be performed to implement this requirement?

(Choose 4)

- a) Create an external file format that would map to the parquet-based files
- b) Load the data into a staging table
- c) Create an external table called lead2pass\_sales\_details
- d) Create an external data source for the Azure Blob storage account
- e) Create a master key on the database
- f) Configure Polybase to use the Azure Blob storage account

**Answer: b, c, d, e**

Question: 10

You are designing a big data streaming solution. You need to choose the most appropriate resource for different scenarios. Which Azure resource should you choose?

To answer, drag the appropriate resource to each scenario. A resource may be used once, more than once, or not at all.

	You want to run parallel processing on 4-terabytes (TB) of data.								
	You want to analyze telemetry data sent to an IoT Hub.								
<table border="1" style="width: 100%; border-collapse: collapse;"> <tr> <td style="padding: 5px;">Stream Analytics</td> <td style="padding: 5px;">SQL Data Warehouse</td> <td style="padding: 5px;">Databricks</td> <td style="padding: 5px;">IoT Hub</td> </tr> <tr> <td style="padding: 5px;">Data Lake</td> <td style="padding: 5px;">Data Factory</td> <td colspan="2"></td> </tr> </table>		Stream Analytics	SQL Data Warehouse	Databricks	IoT Hub	Data Lake	Data Factory		
Stream Analytics	SQL Data Warehouse	Databricks	IoT Hub						
Data Lake	Data Factory								

- a) SQL Data Warehouse  
Stream Analytics
- b) SQL Data Warehouse  
Databricks
- c) Stream Analytics  
Data Lake
- d) IoT Hub  
Data Factory

**Answer: a**

## Study Guide to Crack Microsoft Data Engineering on Microsoft Azure DP-203 Exam:

- Getting details of the DP-203 syllabus, is the first step of a study plan. This pdf is going to be of ultimate help. Completion of the syllabus is must to pass the DP-203 exam.
- Making a schedule is vital. A structured method of preparation leads to success. A candidate must plan his schedule and follow it rigorously to attain success.
- Joining the Microsoft provided training for DP-203 exam could be of much help. If there is specific training for the exam, you can discover it from the link above.
- Read from the DP-203 sample questions to gain your idea about the actual exam questions. In this PDF useful sample questions are provided to make your exam preparation easy.
- Practicing on DP-203 practice tests is must. Continuous practice will make you an expert in all syllabus areas.

### Reliable Online Practice Test for DP-203 Certification

Make EduSum.com your best friend during your Data Engineering on Microsoft Azure exam preparation. We provide authentic practice tests for the DP-203 exam. Experts design these online practice tests, so we can offer you an exclusive experience of taking the actual DP-203 exam. We guarantee you 100% success in your first exam attempt if you continue practicing regularly. Don't bother if you don't get 100% marks in initial practice exam attempts. Just utilize the result section to know your strengths and weaknesses and prepare according to that until you get 100% with our practice tests. Our evaluation makes you confident, and you can score high in the DP-203 exam.

**Start Online practice of DP-203 Exam by visiting URL**

**<https://www.edusum.com/microsoft/dp-203-data-engineering-microsoft-azure>**