

# Valid DP-203 Exam Notes, DP-203 Test Assessment

Exam DP-203: Data Engineering on Microsoft Azure

## Study Guide

### Exam DP-203: Data Engineering on Microsoft Azure

#### Purpose of this document

This study guide should help you understand what to expect on the exam and includes a summary of the topics the exam might cover and links to additional resources. The information and materials in this document should help you focus your studies as you prepare for the exam.

Useful links	Description
<a href="#">How to earn the certification</a>	Some certifications only require one exam, while others require more. On the details page, you'll find information about what skills are measured and links to registration. Each exam also has its own details page covering exam specifics.
<a href="#">Certification renewal</a>	Once you earn your certification, don't let it expire. When you have an active certification that's expiring within six months, you should renew it—at no cost—by passing a renewal assessment on Microsoft Learn. Remember to renew your certification annually if you want to retain it.
<a href="#">Your Microsoft Learn profile</a>	Connecting your certification profile to Learn brings all your learning activities together. You'll be able to schedule and renew exams, share and print certificates, badges and transcripts, and review your learning statistics inside your Learn profile.
<a href="#">Passing score</a>	All technical exam scores are reported on a scale of 1 to 1,000. A passing score is 700 or greater. As this is a scaled score, it may not equal 70% of the points. A passing score is based on the knowledge and skills needed to demonstrate competence as well as the difficulty of the questions.
<a href="#">Exam sandbox</a>	Are you new to Microsoft certification exams? You can explore the exam environment by visiting our exam sandbox. We created the sandbox as an opportunity for you to experience an exam before you take it. In the sandbox, you can interact with different question types, such as build list, case studies.



1

2026 Latest Exams-boost DP-203 PDF Dumps and DP-203 Exam Engine Free Share: [https://drive.google.com/open?id=12yDtc83Tv0BOCvysfYE-D0cqc9eh\\_VWP](https://drive.google.com/open?id=12yDtc83Tv0BOCvysfYE-D0cqc9eh_VWP)

Unlike many other learning materials, our DP-203 study materials are specially designed to help people pass the exam in a more productive and time-saving way, and such an efficient feature makes it a wonderful assistant in personal achievement as people have less spare time nowadays. On the other hand, DP-203 Study Materials are aimed to help users make best use of their sporadic time by adopting flexible and safe study access.

To take the DP-203 exam, candidates need to have a good understanding of data engineering concepts and experience working with Azure services. It is recommended that candidates have experience with data storage solutions, data processing, and data security. Candidates can prepare for the exam by taking the official Microsoft training course, DP-203T00-A: Data Engineering on Microsoft Azure, or by using study materials available online.

Earning the Microsoft DP-203 certification demonstrates that a candidate has the skills and knowledge needed to design, implement, and manage data solutions in a cloud-based environment. Data Engineering on Microsoft Azure certification is highly valued by employers and can help data professionals advance their careers by opening up new job opportunities and higher salaries. Overall, the DP-203 Certification Exam is an excellent way for data professionals to validate their skills and knowledge in data engineering on Microsoft Azure.

>> Valid DP-203 Exam Notes <<

## Quiz 2026 Newest DP-203: Valid Data Engineering on Microsoft Azure Exam Notes

In order to make sure your whole experience of buying our DP-203 prep guide more comfortable, our company will provide all people with 24 hours online service. The experts and professors from our company designed the online service system for all customers. If you decide to buy the DP-203 study braindumps from our company, we can make sure that you will have the opportunity to enjoy the best online service provided by our excellent online workers. If you purchasing the DP-203 Test Practice files designed by many experts and professors from our company, we can promise that our online workers are going to serve you day and night during your learning period. If you have any questions about our study materials, you can send an email to us, and then the online workers from our company will help you solve your problem in the shortest time. So do not hesitate to buy our DP-203 prep guide.

## Why it's worth investing in a certification like Microsoft DP-203 Exam

If you are interested in passing Microsoft DP-203 Exam and getting Data Engineering on Microsoft Azure Certification, you should know that this certification will offer you the credibility that employers seek. Data Engineering on Microsoft Azure Certification is an award-winning certification that will allow you to prove your proficiency in cloud computing. It is designed for IT professionals who have knowledge about data engineering and cloud infrastructure services. Data Engineering on Microsoft Azure Certification will help you to demonstrate that you have mastered the skills required to deploy, configure and manage data solutions in the cloud. You will also be able to prove that you have a high level of expertise in the area of data warehousing with SQL Server. Data Engineering on Microsoft Azure Certification has been created to provide individuals with a strong set of skills. **Microsoft DP-203 Dumps** will enable you to acquire the skills you need to work as a data engineer. It can really help you to get ahead in your career by proving your abilities to potential employers. While it is true that there are many other options available when it comes to certifications, this one offers something special because it provides an opportunity for job seekers as well as working professionals to get a competitive advantage over others when it comes to hiring opportunities. Data Engineering on Microsoft Azure Certificate training has been designed by the best industry experts and they have ensured that all students clear the exam successfully with ease.

## Microsoft Data Engineering on Microsoft Azure Sample Questions (Q194-Q199):

### NEW QUESTION # 194

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 Azure Stream Analytics solution that will analyze Twitter data.

You need to count the tweets in each 10-second window. The solution must ensure that each tweet is counted only once.

Solution: You use a tumbling window, and you set the window size to 10 seconds.

Does this meet the goal?

- A. No
- B. Yes


### Answer: B

Explanation:

Explanation

Tumbling windows are a series of fixed-sized, non-overlapping and contiguous time intervals. The following diagram illustrates a stream with a series of events and how they are mapped into 10-second tumbling windows.

Tell me the count of tweets per time zone every 10 seconds



```
SELECT TimeZone, COUNT(*) AS Count
FROM TwitterStream TIMESTAMP BY CreatedAt
GROUP BY TimeZone, TumblingWindow(second,10)
```

Reference:

<https://docs.microsoft.com/en-us/stream-analytics-query/tumbling-window-azure-stream-analytics>

#### NEW QUESTION # 195

You have an Azure Synapse Analytics dedicated SQL pool named pool1.

You plan to implement a star schema in pool1 and create a new table named DimCustomer by using the following code.

```
CREATE TABLE dbo.[DimCustomer](
    [CustomerKey] int NOT NULL,
    [CustomerSourceID] [int] NOT NULL,
    [Title] [nvarchar](8) NULL,
    [FirstName] [nvarchar](50) NOT NULL,
    [MiddleName] [nvarchar](50) NULL,
    [LastName] [nvarchar](50) NOT NULL,
    [Suffix] [nvarchar](10) NULL,
    [CompanyName] [nvarchar](128) NULL,
    [SalesPerson] [nvarchar](256) NULL,
    [EmailAddress] [nvarchar](50) NULL,
    [Phone] [nvarchar](25) NULL,
    [InsertedDate] [datetime] NOT NULL,
    [ModifiedDate] [datetime] NOT NULL,
    [HashKey] [varchar](100) NOT NULL,
    [IsCurrentRow] [bit] NOT NULL
)
WITH
(
    DISTRIBUTION = REPLICATE,
    CLUSTERED COLUMNSTORE INDEX
);
GO
```

You need to ensure that DimCustomer has the necessary columns to support a Type 2 slowly changing dimension (SCD). Which two columns should you add? Each correct answer presents part of the solution.

NOTE: Each correct selection is worth one point.

- A. [EffectiveEndDate] [datetime] NOT NULL
- B. [PreviousModifiedDate] [datetime] NOT NULL
- C. [RowID] [bigint] NOT NULL
- D. [HistoricalSalesPerson] [nvarchar] (256) NOT NULL
- E. [EffectiveStartDate] [datetime] NOT NULL

**Answer: A,D**

#### NEW QUESTION # 196

You have an Azure Data Lake Storage Gen2 container.

Data is ingested into the container, and then transformed by a data integration application. The data is NOT modified after that. Users can read files in the container but cannot modify the files.

You need to design a data archiving solution that meets the following requirements:

- \* New data is accessed frequently and must be available as quickly as possible.
- \* Data that is older than five years is accessed infrequently but must be available within one second when requested.
- \* Data that is older than seven years is NOT accessed. After seven years, the data must be persisted at the lowest cost possible.
- \* Costs must be minimized while maintaining the required availability.

How should you manage the data? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point

Five-year-old data:

- Microsoft
- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Seven-year-old data:

- Microsoft
- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

**Answer:**

Explanation:

Five-year-old data:

- Microsoft
- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Seven-year-old data:

- Microsoft
- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Explanation

Five-year-old data:

- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Seven-year-old data:

- Delete the blob.
- Move to archive storage.
- Move to cool storage.
- Move to hot storage.

Box 1: Move to cool storage

Box 2: Move to archive storage

Archive - Optimized for storing data that is rarely accessed and stored for at least 180 days with flexible latency requirements, on the order of hours.

The following table shows a comparison of premium performance block blob storage, and the hot, cool, and archive access tiers.

	Premium performance	Hot tier	Cool tier	Archive tier
Availability	99.9%	99.9%	99%	Offline
Availability (RA-GRS reads)	N/A	99.99%	99.9%	Offline
Usage charges	Higher storage costs, lower access, and transaction cost	Higher storage costs, lower access, and transaction costs	Lower storage costs, higher access, and transaction costs	Lowest storage costs, highest access, and transaction costs
Minimum storage duration	N/A	N/A	30 days <sup>1</sup>	180 days
Latency (Time to first byte)	Single-digit milliseconds	milliseconds	milliseconds	hours <sup>2</sup>

Reference:

<https://docs.microsoft.com/en-us/azure/storage/blobs/storage-blob-storage-tiers>

Box 1: Replicated

Replicated tables are ideal for small star-schema dimension tables, because the fact table is often distributed on a column that is not

compatible with the connected dimension tables. If this case applies to your schema, consider changing small dimension tables currently implemented as round-robin to replicated.

Box 2: Replicated

Box 3: Replicated

Box 4: Hash-distributed

For Fact tables use hash-distribution with clustered columnstore index. Performance improves when two hash tables are joined on the same distribution column.

Reference:

<https://azure.microsoft.com/en-us/updates/reduce-data-movement-and-make-your-queries-more-efficient-with-th>

<https://azure.microsoft.com/en-us/blog/replicated-tables-now-generally-available-in-azure-sql-data-warehouse/>

### NEW QUESTION # 197

You have an Azure Data Lake Storage Gen2 account named adls2 that is protected by a virtual network.

You are designing a SQL pool in Azure Synapse that will use adls2 as a source.

What should you use to authenticate to adls2?

- A. an Azure Active Directory (Azure AD) user
- B. a managed identity
- C. a shared access signature (SAS)
- D. a shared key

**Answer: B**

### NEW QUESTION # 198

You build an Azure Data Factory pipeline to move data from an Azure Data Lake Storage Gen2 container to a database in an Azure Synapse Analytics dedicated SQL pool.

Data in the container is stored in the following folder structure.

/in/{YYYY}/{MM}/{DD}/{HH}/{mm}

The earliest folder is /in/2021/01/01/00/00. The latest folder is /in/2021/01/15/01/45.

You need to configure a pipeline trigger to meet the following requirements:

Existing data must be loaded.

Data must be loaded every 30 minutes.

Late-arriving data of up to two minutes must be included in the load for the time at which the data should have arrived.

How should you configure the pipeline trigger? To answer, select the appropriate options in the answer area.

NOTE: Each correct selection is worth one point.

Type:

Event
On-demand
Schedule
Tumbling window

Additional properties:

Prefix: /in/, Event: Blob created
Recurrence: 30 minutes, Start time: 2021-01-01T00:00
Recurrence: 30 minutes, Start time: 2021-01-01T00:00, Delay: 2 minutes
Recurrence: 32 minutes, Start time: 2021-01-15T01:45

**Answer:**

Explanation:



Type:

- Event
- On-demand
- Schedule
- Tumbling window**

Additional properties:

Prefix: /in/, Event: Blob created  
Recurrence: 30 minutes, Start time: 2021-01-01T00:00  
Recurrence: 30 minutes, Start time: 2021-01-01T00:00, Delay: 2 minutes  
Recurrence: 32 minutes, Start time: 2021-01-15T01:45

### Reference:

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

## NEW QUESTION # 199

• • • • •

**DP-203 Test Assessment:** <https://www.exams-boost.com/DP-203-valid-materials.html>

2026 Latest Exams-boost DP-203 PDF Dumps and DP-203 Exam Engine Free Share: <https://drive.google.com/open?id=12yDtc83Tv0BOCvysfYE-D0cqce9eh> VWP