DevOpsSchool

Course curriculum · November 2026

AZURE · TRAINING PROGRAM

# Data Engineering on Microsoft Azure Course

A practical, hands-on Data Engineering on Microsoft Azure Course curriculum: 129 topics over 5 days. Learn live online in a small batch, as a private corporate batch, or self-paced with one year of LMS access.

| | | |
|---|---|---|
| LIVE ONLINE<br>**Max 8-10 hours** | CORPORATE<br>**Private team batch, on demand** | SELF-PACED VIDEO<br>**1-year LMS access** |
| BATCH SIZE<br>**Max 10 learners** | NEXT BATCHES<br>**1st of every month · weekends: first Saturday** | CERTIFICATE<br>**Verifiable certificate of completion** |

## INSIDE THIS CURRICULUM


| | |
|---|---|
| Who should attend | Objectives, audience, prerequisites |
| Agenda | 129 topics over 5 days |
| Formats and fees | Live, 1:1, self-paced, corporate |
| Trainer and reviews | Rajesh Kumar |
| Enrolment | Batches, payment, contacts |


## WHY DEVOPSSCHOOL

- Rated 4.6 on Google
- 1,500+ verifiable certificates issued in the public registry
- Trainer Rajesh Kumar: 20 years in industry
- Small batches: max 10 learners, every session recorded

**Course page:** www.devopsschool.com/courses/azure/data-engineering-on-microsoft-azure-training.html
**Talk to us:** WhatsApp +91 99057 40781 · +1 (469) 756-6329 · contact@DevOpsSchool.com · reply within 1 hour

COURSE FIT

# Who this program is for

Use this page to decide whether Data Engineering on Microsoft Azure Course matches your role and experience.

1

### Objectives

Understand the core concepts and usage model, apply practical workflows through examples and guided tasks, and prepare for implementation, certification or team adoption.

2

### Who should attend

Developers, DevOps engineers, administrators, SREs, security engineers, platform teams, release teams, technical leads and corporate teams working with this topic.

3

### Before you start

Basic computer literacy and awareness of software, operations, cloud, DevOps or security workflows is helpful. Deep programming experience is not needed unless the topic requires it.

AGENDA AT A GLANCE · 129 TOPICS OVER 5 DAYS

**Day 1**

36 topics · Module 1: Explore compute and storage options for data engineering workloads

**Day 2**

35 topics · Module 5: Explore, transform, and load data into the Data Warehouse using Apache Spark

**Day 3**

15 topics · Module 9: Orchestrate data movement and transformation in Azure Synapse Pipelines

**Day 4**

24 topics · Module 12: Support Hybrid Transactional Analytical Processing (HTAP) with Azure Synapse Li

**Day 5**

19 topics · Module 15: Create a Stream Processing Solution with Event Hubs and Azure Databricks

AGENDA · 129 TOPICS OVER 5 DAYS

# Data Engineering on Microsoft Azure Course agenda

Days are for corporate (day-wise) delivery; live online batches cover the agenda in max 8-10 hours.

## DAY 1

36 topics

1. Module 1: Explore compute and storage options for data engineering workloads
2. Introduction to Azure Synapse Analytics
3. Describe Azure Databricks
4. Introduction to Azure Data Lake storage
5. Describe Delta Lake architecture
6. Work with data streams by using Azure Stream Analytics
7. Combine streaming and batch processing with a single pipeline
8. Organize the data lake into levels of file transformation
9. Index data lake storage for query and workload acceleration
10. Describe Azure Synapse Analytics
11. Describe Azure Data Lake storage
12. Describe Azure Stream Analytics
13. Module 2: Design and implement the serving layer
14. Design a multidimensional schema to optimize analytical workloads
15. Code-free transformation at scale with Azure Data Factory
16. Populate slowly changing dimensions in Azure Synapse Analytics pipelines
17. Design a star schema for analytical workloads
18. Populate slowly changing dimensions with Azure Data Factory and mapping data flows
19. Populate a slowly changing dimensions with Azure Data Factory and mapping data flows
20. Module 3: Data engineering considerations for source files
21. Design a Modern Data Warehouse using Azure Synapse Analytics
22. Secure a data warehouse in Azure Synapse Analytics
23. Managing files in an Azure data lake
24. Securing files stored in an Azure data lake
25. Module 4: Run interactive queries using Azure Synapse Analytics serverless SQL pools
26. Explore Azure Synapse serverless SQL pools capabilities
27. Query data in the lake using Azure Synapse serverless SQL pools
28. Create metadata objects in Azure Synapse serverless SQL pools
29. Secure data and manage users in Azure Synapse serverless SQL pools
30. Query Parquet data with serverless SQL pools
31. Create external tables for Parquet and CSV files
32. Create views with serverless SQL pools
33. Secure access to data in a data lake when using serverless SQL pools
34. Configure data lake security using Role-Based Access Control (RBAC) and
35. Access Control List
36. Understand Azure Synapse serverless SQL pools capabilities

**DAY 2** — 35 topics

37. Module 5: Explore, transform, and load data into the Data Warehouse using Apache Spark
38. Understand big data engineering with Apache Spark in Azure Synapse Analytics
39. Ingest data with Apache Spark notebooks in Azure Synapse Analytics
40. Transform data with DataFrames in Apache Spark Pools in Azure Synapse
41. Analytics
42. Integrate SQL and Apache Spark pools in Azure Synapse Analytics
43. Perform Data Exploration in Synapse Studio
44. Ingest data with Spark notebooks in Azure Synapse Analytics
45. Transform data with DataFrames in Spark pools in Azure Synapse Analytics
46. Integrate SQL and Spark pools in Azure Synapse Analytics
47. Describe big data engineering with Apache Spark in Azure Synapse Analytics
48. Module 6: Data exploration and transformation in Azure Databricks
49. Describe Azure Databricks
50. Read and write data in Azure Databricks
51. Work with DataFrames in Azure Databricks
52. Work with DataFrames advanced methods in Azure Databricks
53. Use DataFrames in Azure Databricks to explore and filter data
54. Cache a DataFrame for faster subsequent queries
55. Remove duplicate data
56. Manipulate date/time values
57. Remove and rename DataFrame columns
58. Aggregate data stored in a DataFrame
59. Module 7: Ingest and load data into the data warehouse
60. Use data loading best practices in Azure Synapse Analytics
61. Petabyte-scale ingestion with Azure Data Factory
62. Perform petabyte-scale ingestion with Azure Synapse Pipelines
63. Import data with PolyBase and COPY using T-SQL
64. Module 8: Transform data with Azure Data Factory or Azure Synapse Pipelines
65. Data integration with Azure Data Factory or Azure Synapse Pipelines
66. Code-free transformation at scale with Azure Data Factory or Azure Synapse Pipelines
67. Execute code-free transformations at scale with Azure Synapse Pipelines
68. Create data pipeline to import poorly formatted CSV files
69. Create Mapping Data Flows
70. Perform data integration with Azure Data Factory
71. Perform code-free transformation at scale with Azure Data Factory

**DAY 3** — 15 topics

72. Module 9: Orchestrate data movement and transformation in Azure Synapse Pipelines
73. Orchestrate data movement and transformation in Azure Data Factory

74. Integrate Data from Notebooks with Azure Data Factory or Azure Synapse Pipelines
75. Orchestrate data movement and transformation in Azure Synapse Pipelines
76. Module 10: Optimize query performance with dedicated SQL pools in Azure Synapse
77. Optimize data warehouse query performance in Azure Synapse Analytics
78. Understand data warehouse developer features of Azure Synapse Analytics
79. Understand developer features of Azure Synapse Analytics
80. Improve query performance
81. Module 11: Analyze and Optimize Data Warehouse Storage
82. Analyze and optimize data warehouse storage in Azure Synapse Analytics
83. Check for skewed data and space usage
84. Understand column store storage details
85. Study the impact of materialized views
86. Explore rules for minimally logged operations

**DAY 4** 24 topics

87. Module 12: Support Hybrid Transactional Analytical Processing (HTAP) with Azure Synapse Link
88. Design hybrid transactional and analytical processing using Azure Synapse Analytics
89. Configure Azure Synapse Link with Azure Cosmos DB
90. Query Azure Cosmos DB with Apache Spark pools
91. Query Azure Cosmos DB with serverless SQL pools
92. Query Azure Cosmos DB with Apache Spark for Synapse Analytics
93. Query Azure Cosmos DB with serverless SQL pool for Azure Synapse Analytics
94. Query Azure Cosmos DB with Apache Spark for Azure Synapse Analytics
95. Query Azure Cosmos DB with SQL serverless for Azure Synapse Analytics
96. Module 13: End-to-end security with Azure Synapse Analytics
97. Secure a data warehouse in Azure Synapse Analytics
98. Configure and manage secrets in Azure Key Vault
99. Implement compliance controls for sensitive data
100. Secure Azure Synapse Analytics supporting infrastructure
101. Secure the Azure Synapse Analytics workspace and managed services
102. Secure Azure Synapse Analytics workspace data
103. Module 14: Real-time Stream Processing with Stream Analytics
104. Enable reliable messaging for Big Data applications using Azure Event Hubs
105. Work with data streams by using Azure Stream Analytics
106. Ingest data streams with Azure Stream Analytics
107. Use Stream Analytics to process real-time data from Event Hubs
108. Use Stream Analytics windowing functions to build aggregates and output to Synapse Analytics
109. Scale the Azure Stream Analytics job to increase throughput through partitioning
110. Repartition the stream input to optimize parallelization

**DAY 5** 19 topics

111. Module 15: Create a Stream Processing Solution with Event Hubs and Azure Databricks
112. Process streaming data with Azure Databricks structured streaming
113. Explore key features and uses of Structured Streaming
114. Stream data from a file and write it out to a distributed file system
115. Use sliding windows to aggregate over chunks of data rather than all data
116. Apply watermarking to remove stale data
117. Connect to Event Hubs read and write streams
118. Module 16: Build reports using Power BI integration with Azure Synapase Analytics
119. Create reports with Power BI using its integration with Azure Synapse Analytics
120. Integrate an Azure Synapse workspace and Power BI
121. Optimize integration with Power BI
122. Improve query performance with materialized views and result-set caching
123. Visualize data with SQL serverless and create a Power BI report
124. Module 17: Perform Integrated Machine Learning Processes in Azure Synapse Analytics
125. Use the integrated machine learning process in Azure Synapse Analytics
126. Create an Azure Machine Learning linked service
127. Trigger an Auto ML experiment using data from a Spark table
128. Enrich data using trained models
129. Serve prediction results using Power BI

HOW YOU CAN LEARN

## Formats and fees

The same agenda in four formats. Fees are in Indian rupees as of November 2026; the course page shows them in your currency.

| FORMAT | DURATION | WHAT YOU GET | FEE |
|---|---|---|---|
| **Live & interactive (online)** | Max 8-10 hours | Instructor-led public batch of max 10 learners with demos, labs, Q&A and assignment review. Weekday batches start on the 1st of every month, weekend batches on the first Saturday of every month. | **INR 34,999** |
| **1:1 live training (online)** | Max 8-10 hours, one-to-one | The same agenda, taught one-to-one by the trainer. Starts immediately, without waiting for the next batch; sessions planned around your working hours, at your pace. | **INR 99,999** |
| **Self-paced video** | 1-year LMS access | This program plus 20+ more courses and 50+ tools, at your own pace. | **INR 833/month** billed yearly at INR 9,996 |
| **Corporate** | Planned with your team | A private batch for your team, on demand, with examples mapped to your platforms. | **Quote on request** |

ASSESSMENT AND CERTIFICATE

Your progress is checked through hands-on tasks and assignments. On completion you receive a DevOpsSchool certificate of completion with a unique credential ID that anyone can verify at devopsschool.com/certificates.

INCLUDED WITH EVERY LIVE BATCH

- Every live session is recorded and stays in your LMS for a year.
- Can't attend your batch? Move to next month's batch.
- Free demo class before you pay.
- Free 10-15 minute call with an advisor or the trainer to check the course fits you.

YOUR TRAINER

## Learn from Rajesh Kumar

Principal Architect, trainer and mentor in DevOps, SRE, DevSecOps and platform engineering, with 20 years across 8+ software MNCs, including Staff DevOps/SRE Engineer for a leading Japanese mobile-payment platform and Principal DevOps Architect and Manager in Europe.

Profile: www.rajeshkumar.xyz · Videos: youtube.com/@TheDevOpsSchool

| | |
|---|---|
| **20** years in industry | **200+** corporate trainings |
| **10,000+** engineers mentored | **70+** organisations coached |

REVIEWS

## What learners say

Rated 4.6 on Google · more reviews at devopsschool.com/testimonials and in our video testimonials.

> “Rajesh is very good trainer. Rajesh was able to resolve our queries and question effectively. We really liked the hands-on examples covered during this training program.”
>
> Indrayani, India

> “The training was very useful and interactive. Rajesh helped develop the confidence of all.”
>
> Abhinav Gupta, Pune

START YOUR TRAINING

## Join the next batch

Weekday batches start on the 1st of every month and weekend batches on the first Saturday of every month, with max 10 learners each. 1:1 training can start immediately; corporate batches run on demand. We reply within 1 hour.

Pay by UPI, credit and debit cards, net banking or PayPal. Before you pay, read our refund policy. All courses: www.devopsschool.com/courses/

**ENROL OR RESERVE A SEAT**
Data Engineering on Microsoft Azure Course

**WHATSAPP · INDIA**
+91 99057 40781

**WHATSAPP · USA**
+1 (469) 756-6329

**EMAIL**
contact@DevOpsSchool.com

**More Azure training: www.devopsschool.com/courses/azure/**