



SQL SCHOOL
QUALITY TRAINING ASSURED

NEW COURSE



MICROSOFT WITH AI CLOUD DATA ENGINEER

Build • Engineer • Migrate • Innovate with Microsoft Cloud & AI

MODULE 01		MSSQL & TSQL CONCEPTS MSSQL, TSQL, Database Concepts, Queries, Functions, Views & more.	 3 WEEKS
MODULE 02		AZURE DATA ENGINEERING ADF, Databricks, Spark, PySpark, ADLS, IoT, Event Hubs, Stream Analytics & more.	 7 WEEKS
MODULE 03		FABRIC DATA ENGINEERING FDF, LakeHouse, DataFlow, PySpark, SparkSQL, KQL Database & more.	 6 WEEKS
MODULE 04		REALTIME PROJECTS & MIGRATIONS End to End Realtime Projects & Cloud Migrations (On-Prem to Cloud)	 1 WEEK


100% PRACTICAL
Hands-on Training


REAL-TIME PROJECTS
End to End Implementation


CLOUD MIGRATIONS
On-Prem to Cloud


CERTIFICATION GUIDANCE
Microsoft & Industry Certs


REACH US FOR FREE DEMO


+91 99514 40801


www.sqlschool.com


20+ YEARS OF TRUST

Learn from
Mr. Sai Phanindra Tholeti
20+ Years of Experience
on Data & Cloud Technologies

I am **Mr. Sai Phanindra**, trainer for this Step by step, 100% Practical Job Oriented Training Program on **Microsoft Cloud Data Engineer** Course.

With over 20 Years of experience in Database Technologies, Azure, Fabric, Power BI, Python and Big Data. I am strongly committed to give KT sessions with 100% practical, step-by-step approach and real-world projects. Let's Connect on LinkedIn: [linkedin.com/in/saiphanindra/](https://www.linkedin.com/in/saiphanindra/)

This Microsoft Cloud Engineer Course is divided into FOUR Comprehensive Modules:

Module 1: MSSQL & TSQL Concepts

Module 2: Azure Data Engineering (ADF, Databricks, Spark, PySpark, ADLS, IoT, more..)

Module 3: Fabric Data Engineering (FDF, LakeHouse, DataFlow, PySpark, SparkSQL, AI, more..)

Module 4: Realtime Projects & Migrations (End to End)

This course includes Complete Guidance for:

- **Microsoft DP-750 Exam** (Azure Databricks)
- **Databricks Data Engineer Associate Exam** (Open Source Databricks)
- **DP-700 Exam** (Microsoft Fabric Data Engineer)

Training Details:

- ✓ **100% Practical & Realtime**
- ✓ **Training Modes:**
 - **LIVE Online Training**
 - **Self Paced Videos** (Detailed, step by Step, Latest Videos with Notes, Project Solution)

www.sqlschool.com For Free Demo: +91 99514 40801, +91 96664 40801

★ COURSE PLAN		MICROSOFT CLOUD ENGINEER Build • Engineer • Migrate • Innovate on Microsoft Cloud	
MODULE	COURSE MODULE	KEY TECHNOLOGIES / TOPICS	DURATION
01 MODULE 1	 MSSQL & TSQL Concepts	MSSQL, TSQL, Database Concepts	 3 Weeks
02 MODULE 2	 Azure Data Engineering	ADF, Databricks, Spark, PySpark, ADLS, IoT, more..	 7 Weeks
03 MODULE 3	 Fabric Data Engineering	FDF, LakeHouse, DataFlow, PySpark, SparkSQL, more..	 6 Weeks
04 MODULE 4	 Realtime Projects & Migrations	End to End Projects & Cloud Migrations	 1 Week
 TOTAL COURSE DURATION		  Complete Microsoft Cloud Engineer Program	 17 Weeks

Who Should Join?

- Freshers
- SQL Developers
- BI Developers
- Data Analysts
- SDEs & SREs

Prerequisites

- No pre-requisites needed.
- We start from the basics and build your skills step by step.

NEW COURSE

MICROSOFT CLOUD ENGINEER

Build • Engineer • Migrate • Innovate on Microsoft Cloud

100% PRACTICAL
REAL-TIME LEARNING

MODULE 1



MSSQL & TSQL CONCEPTS

- Database Fundamentals
- MSSQL Architecture
- T-SQL Queries & Functions
- Joins, Subqueries, Views
- Stored Procedures, Triggers
- Indexes & Performance

MODULE 2



AZURE DATA ENGINEERING

(ADF, DATABRICKS, SPARK, PYSPARK, ADLS, IOT, MORE..)

- Azure Data Factory (ADF)
- Databricks, Spark, PySpark
- Azure Data Lake Storage (ADLS)
- IoT Hub & Event Hubs
- Stream Analytics & Logic Apps
- End-to-End Data Engineering

MODULE 3



FABRIC DATA ENGINEERING

(FDF, LAKEHOUSE, DATAFLOW, PYSPARK, SPARKSQL, MORE..)

- Fabric Data Factory (FDF)
- Lakehouse Architecture
- Dataflow Gen2
- PySpark & Spark SQL
- KQL Database & Notebooks
- Integrated Analytics & AI

MODULE 4



REALTIME PROJECTS & MIGRATIONS

(END TO END)

- Real-Time Data Pipelines
- End-to-End Project Implementation
- Cloud Migrations (On-Prem to Cloud)
- Performance & Cost Optimization
- Best Practices & Governance
- Production-Ready Solutions

WHY CHOOSE THIS COURSE?

100% PRACTICAL

REAL-TIME PROJECTS

CLOUD MIGRATIONS

CERTIFICATION GUIDANCE

INTERVIEW & JOB SUPPORT


SQL SCHOOL
QUALITY TRAINING ASSURED

REACH US FOR
FREE DEMO

 **+91 99514 40801**
 **www.sqlschool.com**


20+
YEARS OF TRUST

Detailed Course Content

Module 1: SQL Server (MSSQL), TSQL

Ch 1: SQL Database Job Roles

- ✓ Database Intro
- ✓ OLTP, DWH, OLAP
- ✓ DBMS Basics
- ✓ Data Engineer Job Roles

Ch 2: Database Intro & Installations

- ✓ SQL Server Installations
- ✓ Instance & Collations
- ✓ SSMS Tool Installation
- ✓ Connections, Authentications

Ch 3: SQL Basics V1 (Commands)

- ✓ SQL Basics (DDL, DML, etc..)
- ✓ Creating Databases, Tables
- ✓ Data Inserts (GUI, SQL)
- ✓ Basic SELECT Queries

Ch 4: SQL Basics V2 (Commands, Operators)

- ✓ DDL: Create, Alter, Drop
- ✓ DML: Insert, Update, Delete
- ✓ DQL: Select, Fetch
- ✓ Add, Truncate Statements
- ✓ SQL Operators

Ch 5: Data Types & Variables

- ✓ Integer Data Types
- ✓ Character, MAX Data Types
- ✓ Decimal & Boolean Data Types
- ✓ Date and Time Data Types
- ✓ SQL_Variant Type

Ch 6: Data Imports From Files

- ✓ Data Imports with Excel
- ✓ Data Imports with CSV
- ✓ Auto Detection of DataTypes
- ✓ OLE-DB Connections
- ✓ Order By, TOP, OFFSET

Ch 7: Sub Queries

- ✓ Basic Sub Queries
- ✓ Aggregations
- ✓ Combining Queries
- ✓ UNION, UNION ALL

Ch 8: Schemas & Batches

- ✓ Schemas: Creation, Usage
- ✓ Schemas & Table Grouping
- ✓ Real-world Banking Database
- ✓ 2 Part, 3 Part & 4 Part Naming
- ✓ Batch Concept & “Go” Command

Ch 9: Constraints, Keys & RDBMS

- ✓ Null, Not Null Constraints
- ✓ Unique Key & Check
- ✓ Primary Key Constraint
- ✓ Foreign Keys, Default
- ✓ DB Diagrams & ER Models

Ch 10: Normal Forms & RDBMS

- ✓ Normal Forms: 1 NF, 2 NF
- ✓ 3 NF, BCNF and 4 NF
- ✓ 1:1, 1:M, M:1 Cardinality
- ✓ Cascading Keys
- ✓ Self Referencing Keys

Ch 11: Joins & Queries

- ✓ Joins: Table Comparisons
- ✓ Inner Joins & Matching Data
- ✓ Outer Joins: LEFT, RIGHT
- ✓ Full Outer Joins & Aliases
- ✓ Self Joins & Aliases

Ch 12: Views & RLS

- ✓ Views: Realtime Usage
- ✓ DML, SELECT with Views
- ✓ Excel Analytics with Views
- ✓ Important System Views

Ch 13: Stored Procedures - 1

- ✓ Stored Procedures: Realtime Use
- ✓ Parameters Concept with SPs
- ✓ Procedures with SELECT
- ✓ System Stored Procedures
- ✓ Stored Procedures, Tuning

Ch 14: Stored Procedures - 2

- ✓ Merge Statement (Upsert)
- ✓ Merge with OLTP & DWH
- ✓ Matched and Not Matched
- ✓ Merge Statement inside SPs
- ✓ SP Recompilations

Ch 15: User Defined Functions - 1

- ✓ Scalar Functions in Real-world
- ✓ Inline & Multiline Functions
- ✓ Parameterized Queries
- ✓ Variables & Parameters
- ✓ Function Executions

Ch 16: User Defined Functions - 2

- ✓ Date & Time Functions
- ✓ String Functions & Queries
- ✓ Aggregated Functions & Usage
- ✓ Window Functions (Rank)
- ✓ Row_Number, DenseRank
- ✓ Partition By & Order By

Ch 17: Triggers & Automations

- ✓ Need for Triggers in Real-world
- ✓ DDL & DML Triggers
- ✓ For / After Triggers
- ✓ Instead Of Triggers
- ✓ Memory Tables with Triggers
- ✓ Disabling DMLs Triggers

Ch 18: Group By Queries

- ✓ Group By, Distinct
- ✓ GROUP BY, HAVING
- ✓ Cube() and Rollup()
- ✓ Sub Totals & Grand Totals
- ✓ Grouping() & Usage

Ch 19: Joins with Group By

- ✓ 3 Table, 4 Table Joins
- ✓ Join Queries & WHERE
- ✓ Join Queries & Group By
- ✓ IIF(), CASE Statement
- ✓ Query Execution Order

Ch 20: Transactions & ACID

- ✓ Auto Commit Transaction
- ✓ Explicit Transactions
- ✓ COMMIT, ROLLBACK
- ✓ Checkpoint & Query Blocking
- ✓ READPAST, LOCKHINT

Ch 21: Indexes Basics, Tuning

- ✓ Clustered Index, Primary Key
- ✓ Non Clustered Index

- ✓ Query Optimizer
- ✓ Tuning Join Queries
- ✓ Tuning Group By Queries

Ch 22: CTEs & Tuning

- ✓ Common Table Expression
- ✓ CTEs for Data Retrieval
- ✓ CTEs for DML Operations
- ✓ CTEs for Data Cleansing
- ✓ Using CTEs with Row Number

Ch 23: Cursors & Temp Tables

- ✓ Cursors & Fetch
- ✓ Scroll, Absolute Types
- ✓ Local Temp Tables
- ✓ Global Temp Tables
- ✓ SELECT..INTO Statement

Ch 24: SQL Server Architecture

- ✓ Network Protocols
- ✓ Storage Engine
- ✓ Query Processing Engine
- ✓ Query Execution Order
- ✓ SQL OS Components

Ch 25: Capstone Project (HealthCare Management System)

- ✓ ECommerce Database
- ✓ Data Validations
- ✓ Query Writing
- ✓ Query Tuning
- ✓ Excel Analytics

Module 2: Azure Data Engineer

Part 1: Azure Data Factory & Synapse

Ch 1: Azure Fundamentals

- ✓ Cloud Introduction, Azure
- ✓ Cloud Implementations
- ✓ Azure Account, Subscription
- ✓ Azure ETL & DWH Resources
- ✓ Azure Storage, IoT Resources

Ch 2: Azure Storage & ADLS

- ✓ Azure Storage Account
- ✓ Azure Data Lake Storage
- ✓ Azure BLOB Containers

- ✓ Blob File Uploads
- ✓ Azure Tables

Ch 3: Azure Deployments, Azure SQL

- ✓ Azure SQL Server, SQL DB
- ✓ Azure SQL Database (OLTP)
- ✓ Azure Firewall Rules
- ✓ Connections from SSMS Tool

Ch 4: Azure SQL DB Migrations

- ✓ On-Premise SQL DB, bacpac
- ✓ Azure SQL Deployment
- ✓ Azure Storage from SSMS
- ✓ Azure SQL DB Migration
- ✓ Migration Verifications

Ch 5: Azure Synapse (DWH)

- ✓ Synapse Pool Architecture
- ✓ Control Node, Compute Node
- ✓ DMS (Data Movement Service)
- ✓ Azure DWH Tables: Partitions
- ✓ Distributions & MAXDOP

Ch 6: Azure Data Factory (ADF)

- ✓ Need for ADF & Pipelines
- ✓ Data Orchestration with IR
- ✓ Integration Runtime Engine
- ✓ Linked Services, Datasets
- ✓ Pipelines: Copy Data Activity
- ✓ Data Flow Activity with IR

Ch 7: Azure SQL DB Loads

- ✓ ADF: Author, Azure SQL DB Reads
- ✓ Azure SQL Pool Writes
- ✓ Synapse Analytics with IR
- ✓ Pipeline Design, Validation
- ✓ Pipeline Runs, Monitoring

Ch 8: File Incremental Loads

- ✓ File Incremental Loads
- ✓ Storage Account, Data Lake
- ✓ Binary Copy, Schema Drift
- ✓ Staging Concept in ADF
- ✓ Initial, Incremental Loads

- ✓ Schema & Data Changes

Ch 9: Pipeline Settings

- ✓ ADF Pipeline Settings
- ✓ Staging: Advantages
- ✓ Reliable Logging
- ✓ Best Effort Logging
- ✓ DIU & DOCP with IR
- ✓ Compressions, Health Check

Ch 10: Table Incremental Loads

- ✓ Implement SCD with ADF
- ✓ Self Hosted IR: Realtime Use
- ✓ On-premise Data: Incr Loads
- ✓ Copy Method: Upsert, Keys
- ✓ Staging & ADF Optimizations
- ✓ Pipeline Runs, Activity IDs

Ch 11: ADF Data Flow - 1

- ✓ Creating Data Flow Items
- ✓ Using Multiple Sinks
- ✓ Conditional Split Transformation
- ✓ Select, Sort, Union, Loops
- ✓ ADF Debug Options

Ch 12: ADF Data Flow - 2

- ✓ Working with Multiple Tables
- ✓ Join Transform, Broadcast
- ✓ Surrogate Keys, Derived Cols
- ✓ ETL Loads Dates, Sink Options
- ✓ Aggregated Data Loads

Ch 13: ADF Data Flow - 3

- ✓ Pivot Transformation
- ✓ Group By & Pivot Keys
- ✓ Column Pattern, Deduplicate
- ✓ Lookup, Cached Lookup
- ✓ Tuning Transformations
- ✓ Tuning Data Flow, Spark

Ch 14: ADF Data Flow - 4

- ✓ Get Metadata
- ✓ IF & Lookup Transformation

- ✓ Cache Lookup
- ✓ Data Validations
- ✓ Lookup Versus Joins

Ch 15: ADF Metrics, Alerts

- ✓ Azure Insights
- ✓ Azure Metrics for ADF
- ✓ Azure Metrics for Synapse
- ✓ CPU, Memory Metrics
- ✓ Alerts and Notifications
- ✓ Action Groups, Tuning Options

Ch 16: ADF with Azure Functions

- ✓ Azure Functions
- ✓ Function Activity in ADF
- ✓ Linked Services
- ✓ Pipeline Debug
- ✓ ADF Activity Controls

Ch 17: ADF Optimizations

- ✓ Synapse SQL Pool Partitions
- ✓ ADF Partitions
- ✓ Broadcast Options
- ✓ Staging, Logging
- ✓ DIU, DOCP
- ✓ Spar Cluster Optimizations

Ch 18: ADF Parameters, Security

- ✓ Linked Service Parameters
- ✓ Creating Logins & Users
- ✓ Schemas, ETL Permissions
- ✓ Logins Parameters in ADF
- ✓ Dynamic Linked Services

Ch 19: SCD & ETL with Control Tables

- ✓ ADF Templates in Realtime
- ✓ Incremental Loads (SCD)
- ✓ Control Tables, Watermarks
- ✓ Lookup Activity, Delta Queries
- ✓ SP Activity & Parameters
- ✓ Pipeline Parameters, SPs

Ch 20: Azure Key Vaults

- ✓ Azure Key Vaults

- ✓ Access Policies
- ✓ Secret Management
- ✓ Managed Identity
- ✓ Key Vault Integrations

Ch 21: Synapse Analytics

- ✓ Azure Synapse Analytics
- ✓ Synapse Deployments
- ✓ Synapse Configurations
- ✓ ADLS Containers
- ✓ Workspace Server Setup
- ✓ Synapse Studio (GUI)

Ch 22: Synapse: Dedicated SQL Pools

- ✓ Creating Dedicated SQL Pools
- ✓ BLOB Data Imports
- ✓ TSQL Queries, Data Imports
- ✓ Big Data Analytics

Ch 23: Synapse: Serverless Pools

- ✓ Serverless Pools, TSQL
- ✓ Serverless Architecture
- ✓ OPENROWSET Operations
- ✓ Big Data Analytics

Ch 24: Synapse: Apache Spark Pools

- ✓ Apache Spark Pools
- ✓ Nodes and Executors
- ✓ Big Data Analytics
- ✓ Pipeline Integrations

Part 2: Databricks (Spark, PySpark, Big Data, Genie AI)

Ch 1: Databricks Introduction

- ✓ Cloud ETL, DWH
- ✓ Cloud Computing
- ✓ Databricks Concepts
- ✓ Databricks Account
- ✓ Big Data in Cloud

Ch 2: Databricks Architecture

- ✓ Unity Catalog, Volume
- ✓ Spark Clusters
- ✓ Apache Spark and Databricks
- ✓ Apache Spark Ecosystem

- ✓ Hadoop, MapReduce, Apache Spark

Ch 3: Unity Catalog

- ✓ Unity Catalog Concepts
- ✓ Databricks Workspace UI
- ✓ Organizing Workspace Objects
- ✓ Volumes, File Uploads
- ✓ Managed & External Tables
- ✓ Spark UI: Limitations

Ch 4: Spark SQL: Basics

- ✓ Spark SQL Notebooks
- ✓ Creating Catalog
- ✓ Creating Schemas
- ✓ Spark Data Types
- ✓ PySpark API: SQL Queries
- ✓ Notebooks: Exports, Clone

Ch 5: Spark SQL: Functions

- ✓ Math, Sort Functions
- ✓ String, DateTime Functions
- ✓ SQL Expressions with `expr()`
- ✓ Volume for our Data Assets
- ✓ File Formats, Schema Inference
- ✓ Spark SQL Aggregations

Ch 6: Spark SQL: Time Travel

- ✓ Time Travel Concepts
- ✓ Spark DB: Logical Architecture
- ✓ Spark DB: Physical Store
- ✓ Time Travel, History
- ✓ `DESCRIBE`, `EXTENDED`
- ✓ Version Numbers

Ch 7: Python Introduction, Print

- ✓ Python Introduction
- ✓ Python Versions
- ✓ `print()` & `display()`
- ✓ Single, Multiline Statements

Ch 8: Python Variables

- ✓ Python Variables
- ✓ Declarations, Values
- ✓ Multi Variable Values

- ✓ Common Variable Values
- ✓ Realtime use of Variables

Ch 9: Python Operators

- ✓ Need for Operators
- ✓ Arithmetic Operators
- ✓ Assignment Operators
- ✓ Comparison Operators
- ✓ Operator Precedence
- ✓ Operands in Python

Ch 10: Python Control Statements

- ✓ Python Control Structures
- ✓ If ... Else Statement
- ✓ Short Hand If
- ✓ ELIF & ELSE IF Statements
- ✓ OR, AND Concepts
- ✓ Python Loops

Ch 11: Python Data Types

- ✓ Python Data Types
- ✓ Integer / Int Data Types
- ✓ Float, String Data Types
- ✓ List Data Type
- ✓ Dictionary Data Type
- ✓ Tuple Data Type

Ch 12: Python Modules & Dataframes

- ✓ Pandas
- ✓ NumPy
- ✓ Dataframe Concepts
- ✓ Handling Nulls
- ✓ Data Cleansing Concepts
- ✓ Pandas Series, arrays
- ✓ Indexes, Indexed Lists

Ch 13: PySpark Concepts

- ✓ Constructing Dataframes
- ✓ List Dataframes
- ✓ Pandas Dataframes
- ✓ Contact & Union
- ✓ Merge
- ✓ Join Options with Dataframes

Ch 14: Medallion Architecture - 1

- ✓ Medallion Architecture
- ✓ Aggregated Data Loads
- ✓ Bronze, Silver and Gold
- ✓ Temp Views
- ✓ Spark Tables (Parquet)
- ✓ Work with File Sources

Ch 15: Medallion Architecture - 2

- ✓ Medallion Architecture
- ✓ Azure SQL DB Connections
- ✓ Joining Source Tables
- ✓ Dataframes, Temp Views
- ✓ Aggregated Data Loads
- ✓ Gold Data Consumption

Ch 16: Delta Lake

- ✓ Databricks DeltaLake
- ✓ Schema Evolution
- ✓ Dataframes, Temp Views
- ✓ Delta Table API
- ✓ Update, Delete Records
- ✓ Merging Records
- ✓ Old History Retention
- ✓ Delta Transaction Log

Ch 17: PySpark: Widgets

- ✓ PySpark Parameters
- ✓ Text Widgets
- ✓ User Parameters
- ✓ Manual Executions
- ✓ Automations
- ✓ UI & JSON For Widgets

Ch 18: Lake Flow Jobs

- ✓ Workflows & CRON
- ✓ Job Compute, Running Tasks
- ✓ Parameters into Notebook Tasks
- ✓ Parameters into Python Script Tasks
- ✓ Concurrent Executions, Dependencies
- ✓ Branching Control with the If-Else Task

Ch 19: Pyspark: Auto Loader - 1

- ✓ AutoLoader Concept

- ✓ Cloudfiles Architecture
- ✓ Checkpoint Configurations
- ✓ Creating Directories
- ✓ Reading Databricks Cloud Sources
- ✓ Initial Loads

Ch 20: PySpark: Auto Loader - 2

- ✓ Reading Streams with Auto Loader
- ✓ Reading a Data Stream
- ✓ Manually Cancel your Data Streams
- ✓ Writing to a Data Stream
- ✓ Schema Evaluation Modes
- ✓ Adding New Columns
- ✓ Workspace Modules

Ch 21: Lake Flow Declarative Pipelines

- ✓ SDP: Spark Declarative Pipelines
- ✓ Delta LIVE Tables
- ✓ Streaming Data Loads
- ✓ Materialized Views
- ✓ Pipeline Clusters
- ✓ Databricks CLI
- ✓ Data Quality Checks

Ch 22: Databricks Optimizations

- ✓ Lazy Evaluation
- ✓ Data Shuffling
- ✓ Broadcast Joins
- ✓ Data Skipping
- ✓ Z Ordering
- ✓ Liquid Clustering
- ✓ VACUUM
- ✓ OPTIMIZE

Ch 23: Databricks Security, AI

- ✓ Overview of ACLs
- ✓ Adding a New User to Workspace
- ✓ Workspace Access Control
- ✓ Cluster Access Control
- ✓ Groups & LakeBridge
- ✓ Access Keys (Tokens), Genie AI

Ch 24: Azure Databricks

- ✓ Databricks Deployment Modes

- ✓ Classic Deployments
- ✓ Azure Databricks Workspace
- ✓ Databricks Compute
- ✓ Scaling & Tuning & AI
- ✓ Open Source Databricks Vs Azure Databricks

Ch 25: GitHub Concepts

- ✓ Creating Github Account
- ✓ GIT Project Concept
- ✓ GIT Project Creation
- ✓ GIT: Main, Branches
- ✓ GIT Credentials
- ✓ Connecting with ADF
- ✓ Connecting with Databricks

Part 3: Data Engineer Project 1 (For Resume)

Project 1: ECommerce Platform

Project Objective:

Build an end-to-end Azure Data Engineering solution to process, transform, and analyze e-commerce business data from multiple sources.

Skills Gained:

- Data Ingestion & ETL Development
- Azure Data Factory Pipelines
- Data Orchestration (End to End)
- Data Lake Architecture, RBAC
- Real-Time Industry Experience
- Azure IoT, Stream Analytics
- Azure Databricks with Data Factory

Components For Project (From Resume Perspective):

- ADF Pipelines
- Databricks Notebooks
- Synapse Analytics
- Apache Spark Schemas
- Power BI Reporting
- CI/CD Deployment
- End to End Integrations
- IAM & Managed Identity

Module 3: Fabric Data Engineering

Part 1: Fabric Concepts, DWH & Fabric Data Factory

Ch 1: Fabric Introduction

- ✓ Need for Fabric, Big Data

- ✓ Fabric Data Engineering Model
- ✓ Fabric Components (Items)
- ✓ Microsoft Fabric: Advantages
- ✓ Cloud Warehouse & AI
- ✓ AI with CoPilot
- ✓ Azure Versus Fabric DWH

Ch 2: Fabric Account, Workspace

- ✓ Need for Fabric Workspace
- ✓ Workspace Creation Process
- ✓ ETL, Storage, Analytical
- ✓ Streaming, Monitoring
- ✓ Compute & Separation

Ch 3: Fabric OneLake Architecture

- ✓ Intelligent Data Foundation
- ✓ Polaris Distributed Engine
- ✓ Stateless & Stateful
- ✓ Cache, Metadata, Xact & Data
- ✓ Fabric Tasks, Inputs & DAG
- ✓ State Machine & Statistics
- ✓ Hot Spot Recovery

Ch 4: Fabric Warehouse

- ✓ Fabric Warehouse Creation
- ✓ Fabric Warehouse Features
- ✓ Fabric Warehouse Properties
- ✓ Fabric Warehouse Limitations
- ✓ DWH Internal Operations
- ✓ Default Schemas & Objects
- ✓ SSMS Connections

Ch 5: Fabric Data Types

- ✓ Realtime use of Fabric Houses
- ✓ Exact, Approximate Numbers
- ✓ Date and Time Data Types
- ✓ Fixed & Variable Length
- ✓ Binary & String Data Types
- ✓ Fabric Type Limitations
- ✓ Save As table, Save As View

Ch 6: Fabric Caching

- ✓ Fabric Caching Process
- ✓ In-memory Cache, Disk Cache
- ✓ Cache Types: LRU /MRU
- ✓ Cold Cache / Cold Run

- ✓ Realtime use of Caching
- ✓ Performance Advantages
- ✓ Warehouse Optimizations

Ch 7: Fabric Statistics

- ✓ Query Engine Options
- ✓ Statistics Types
- ✓ Leverage Statistics
- ✓ Auto, Manual Statistics
- ✓ Update Statistics
- ✓ Statistics Consistency
- ✓ Statistics Lists & Reports

Ch 8: Time Travel

- ✓ Continuous Data Protection
- ✓ Data Storage, Retention
- ✓ FOR TIMESTAMP AS OF
- ✓ Time Travel Scenarios
- ✓ Time Travel Implementation
- ✓ Time Travel on Queries
- ✓ Time Travel Limitations

Ch 9: Zero Copy Cloning

- ✓ User Layer, Storage Layer
- ✓ Cloning & Parquet Files
- ✓ Synapse Data Warehouse
- ✓ Data History Retention
- ✓ Point In Time , Schema Level
- ✓ Zero Copy Cloning Limitations

Ch 10: Fabric Security

- ✓ Workspace Security
- ✓ Warehouse Security
- ✓ Item Security & Roles
- ✓ Adding AD Users
- ✓ Item Security Limitations
- ✓ MFA & Client Security

Ch 11: Fabric Copy Job

- ✓ ETL Implementation Options
- ✓ Copy Job Item
- ✓ Data Loads with Copy Job
- ✓ Full Loads
- ✓ Testing Copy Jobs

Ch 12: Fabric Copy Job

- ✓ Incremental Loads with Copy Job

- ✓ Business Key Concept
- ✓ DWH: Data Storage
- ✓ Testing Initial Loads
- ✓ Testing Incremental Loads
- ✓ Copy Job Limitations

Ch 13: Fabric Data Factory

- ✓ Need for Fabric Data Factory
- ✓ ETL Operations in FDF
- ✓ Data Sources, Transformations
- ✓ Activities and Connections
- ✓ Data Destinations (Sinks)
- ✓ Creating Pipelines

Ch 14: Fabric Pipelines Design

- ✓ Creation Options for Pipelines
- ✓ Azure SQL DB Data Loads
- ✓ Creating Data Sets
- ✓ Copy Command Usage
- ✓ Run ID & Monitoring
- ✓ Pipeline Creation, Verification

Ch 15: Data Loads with Azure

- ✓ Azure Data Lake Storage (ADLS)
- ✓ Azure BLOB Containers
- ✓ Fabric Data Loads From Azure Files
- ✓ Fabric Warehouse with Azure
- ✓ Run IDs and Activity
- ✓ Compressions & Advantages

Ch 16: ETL Staging

- ✓ Staging: Advantages
- ✓ Caching & Storing Concept
- ✓ Staging Types in Fabric
- ✓ Workspace & External
- ✓ External Stages in Pipelines
- ✓ Pipeline Trigger, Monitor

Ch 17: Fabric Aggr Data Loads

- ✓ Aggregation Scenarios
- ✓ Creating Views in TSQL
- ✓ Using Views in FDF Pipelines
- ✓ Using Pipeline Editor
- ✓ Data Loads to Warehouse
- ✓ Pipeline Verifications

Ch 18: Fabric Incremental Loads - 1

- ✓ Upsert (Incremental Loads)
- ✓ Business Key Concept
- ✓ SCD: Slowly Changing Dimension
- ✓ Full Loads Vs Incr Loads
- ✓ Testing Incrementing Loads
- ✓ Pipeline Execution Tests

Ch 19: Fabric Incremental Loads - 2

- ✓ Control Tables
- ✓ Watermark Columns
- ✓ Lookup Activity
- ✓ Stored Procedure Activity
- ✓ Parameters & Runs
- ✓ Concurrency & Batch Count

Ch 20: OnPrem Gateways

- ✓ Need for On-Premise Gateway
- ✓ Installing & Configuring
- ✓ Authentication, Usage
- ✓ On-Premise Connections
- ✓ Pipelines for Data Loads
- ✓ Warehouse Data Storage
- ✓ Data Refresh with Gateways

Ch 21: Data Factory Pipeline Tuning

- ✓ Intelligent Throughput
- ✓ DOCP & Optimizations
- ✓ Staging & In-Memory
- ✓ Spark Compute Options
- ✓ Concurrent Connections
- ✓ ETL Partitions in Real-world

Part 2: Fabric Data Flow, Lake House

Ch 1: Fabric Lakehouse

- ✓ Fabric Lakehouse Architecture
- ✓ Files and Tables Storage
- ✓ Direct Lake & AI
- ✓ Creating Lakehouse
- ✓ Azure SQL Database Source
- ✓ UI: Limitations

Ch 2: Lakehouse File Loads

- ✓ Creating Lakehouse
- ✓ Incremental Refresh
- ✓ Computed Tables
- ✓ Reusable Transformations

- ✓ Scheduling
- ✓ Monitoring
- ✓ Best Practices

Ch 3: Power Query Level 1

- ✓ Power Query Concept
- ✓ Fabric Lake House
- ✓ ETL, ELT Process with AI
- ✓ Data Combine: Union, Append
- ✓ Duplicate / Reference Queries
- ✓ Warehouse Data Loads

Ch 4: Power Query Level 2

- ✓ Table Transformations
- ✓ Group By, Transpose
- ✓ Header Row Promotion
- ✓ Reverse Rows, Count Rows
- ✓ Any Column Transformations
- ✓ Data Type, Fill & Pivot

Ch 5: Power Query Level 3

- ✓ Text Transformations
- ✓ Format, SubString
- ✓ Number Transformations
- ✓ Date Time Transformations
- ✓ Add Column Transformation

Ch 6: Power Query Level 4

- ✓ Column From Examples
- ✓ Conditional Column
- ✓ Index Transformation
- ✓ Duplicate Rows, Errors
- ✓ Advanced Editor

Ch 7: Power Query Level 5

- ✓ ETL Parameters
- ✓ Big Data Access
- ✓ Static Parameters
- ✓ Dynamic Parameters
- ✓ List Queries
- ✓ M Language Expressions

Ch 8: Stream House, KQL

- ✓ Need for Stream House
- ✓ Auto creation of KQL
- ✓ Manual KQL Databases
- ✓ Differences with Warehouse
- ✓ Differences with Lakehouse

Ch 9: KQL Query Sets

- ✓ KQL Database Extraction
- ✓ File Imports - on Premises
- ✓ Metadata Edit Options
- ✓ Query Analytics
- ✓ Exports, Visualizations
- ✓ Query Sets Versus Notebooks

Ch 10: Fabric Data Activator

- ✓ Need for Alerts, Notifications
- ✓ Fabric Data Activator Options
- ✓ Alert Conditions, Thresholds
- ✓ Email Notifications
- ✓ Events & Notifications
- ✓ Edit / Enable / Disable

Ch 11: Mirror Database

- ✓ Need for Mirror Databases
- ✓ Configure Mirror Databases
- ✓ Data Replication
- ✓ Schema Replication
- ✓ Connections & Usage
- ✓ Mirror DB Practical uses

Part 3: Python, PySpark, DWH

Ch 1: Fabric Notebooks

- ✓ Need for Notebooks
- ✓ Fabric Notebook Types
- ✓ Creating Environment
- ✓ Creating Spark Clusters
- ✓ Standard, High Concurrency
- ✓ Magic Command
- ✓ Freeze Cells

Ch 2: Spark SQL - 1

- ✓ Spark SQL Notebooks
- ✓ Creating Schemas
- ✓ Delta Tables
- ✓ Parquet Tables
- ✓ Spark Joins
- ✓ Data Partitioning
- ✓ Union, Views in Spark

Ch 3: Spark SQL - 2

- ✓ Math, Sort Functions
- ✓ String, Date Time Functions
- ✓ Conditional Statements

- ✓ Data Recovery & Undo
- ✓ Version Number
- ✓ Describe Extended

Ch 4: Python Intro & Print

- ✓ Python Introduction
- ✓ Python Versions
- ✓ Python in Spark (PySpark)
- ✓ Python Print()
- ✓ Single, Multiline Statements

Ch 5: Python Variables

- ✓ Defining Variables
- ✓ Printing Variables
- ✓ Display Variables
- ✓ Variable Types
- ✓ Multi Value Variables
- ✓ If ... Else Statement

Ch 6: Python Operators

- ✓ Integer Operators
- ✓ String Operators
- ✓ Arithmetic Operators
- ✓ Comparison Operators
- ✓ Formatted Strings
- ✓ Indexing Operators
- ✓ ELIF, ELSE IF Statements

Ch 7: Python Data Types

- ✓ Python Data Types
- ✓ Integer / Int Data Types
- ✓ Float, String Data Types
- ✓ List Items, Indexes
- ✓ Tuple Data Type
- ✓ Dictionary Data Type

Ch 8: Python Dataframes

- ✓ Pandas Module (Python)
- ✓ Dataframes from Lists
- ✓ Dataframe from Dict
- ✓ Pandas Dataframes
- ✓ Dataframe print, display

Ch 9: Python Dataframes Transformations

- ✓ Append
- ✓ Append with NoIndex
- ✓ Merge with ON, KIND
- ✓ spark.read.csv()
- ✓ spark.read.format()

Ch 10: Medallion Architecture

- ✓ Bronze, Gold and Silver
- ✓ Raw Data
- ✓ Data Preparation (Prepping)
- ✓ Temporary Views
- ✓ Big Data Analytics

Ch 11: PySpark: Medallion Loads 1

- ✓ Data Prep (Silver)
- ✓ Filtering DataFrame Records
- ✓ Removing Duplicate Records
- ✓ Sorting and Limiting Records
- ✓ Spark SQL Dataframes
- ✓ Gold Layer Implementation
- ✓ Testing Aggregated Loads

Ch 12: PySpark: Medallion Loads 2

- ✓ Azure SQL DB Connections
- ✓ SQL Queries in PySpark
- ✓ Data Prep (Silver)
- ✓ Filtering Null Values
- ✓ Grouping and Aggregating
- ✓ Spark SQL Dataframes
- ✓ Gold Layer Implementation
- ✓ Notebook Utilities

Ch 13: PySpark: SCD

- ✓ Slowly Changing Dimension
- ✓ Merge Into Statement
- ✓ Error Handling
- ✓ Merge with OLTP Data Sources

Ch 14: PySpark: Widgets

- ✓ Notebook Parameters
- ✓ Text Widgets
- ✓ Parameters & JSON
- ✓ Notebook Schedules, Retry
- ✓ Modular Notebook Design
- ✓ Notebook Chaining

Ch 15: LakeHouse Architecture Optimizations

- ✓ Delta Log
- ✓ Delta Versioning
- ✓ Partition Strategy
- ✓ Small File Problem
- ✓ Adaptive Query Execution
- ✓ Explain Plan, Spark UI
- ✓ mssparkutils

Ch 16: LakeHouse Optimizations

- ✓ **VACUUM, OPTIMIZE**
- ✓ **ZORDER**
- ✓ **Broadcast Join**
- ✓ **Cache, Persist**
- ✓ **Shuffle Optimization**

Ch 17: Semantic Models

- ✓ **Creating Semantic Model**
- ✓ **Spark SQL: DDL, DML**
- ✓ **Adding References, Keys**
- ✓ **Using Model Layouts**

Ch 18: Fabric Security

- ✓ **Workspace Security**
- ✓ **Lakehouse Security**
- ✓ **Notebook Security**
- ✓ **Security Principals**
- ✓ **Authentication Options**
- ✓ **MFA (Multi Factor Authentication)**

Realtime Project on ECommerce Platform using Microsoft Fabric

Business Scenario

A multinational e-commerce company wants to centralize sales, customers, products, orders, inventory, logistics and payment data into Microsoft Fabric.

Goal:

Create a scalable Fabric Data Platform capable of processing millions of records daily.

Skills Gained:

- Fabric Pipelines
- Fabric Notebooks
- Fabric Data Flow
- Semantic Models
- Medallion Architecture (Bronze/Silver/Gold)
- Real-Time Industry Experience
- Fabric APIs (REST, Workspace)

Components For Project (From Resume Perspective):

- Bronze
- Silver
- Gold
- PySpark
- CI/CD
- Deployment
- End to End Integrations

- DP-700 Exam: Complete Guidance

Learning Outcomes

After completing this project, you can confidently showcase experience in:

- Power BI Reporting
- Azure Synapse with Power BI
- Spark Schemas with AI
- Performance Optimization
- End-to-End Azure Data Engineering Solutions

Career Guidance

- ✓ ATS-Friendly Resume
- ✓ Resume Optimization
- ✓ Job Search Platforms

SQL SCHOOL
QUALITY TRAINING ASSURED

THANK YOU!
Your interest in becoming a
MICROSOFT CLOUD ENGINEER
is confirmed.

We have received your details and our team will connect with you soon.

You will receive a confirmation email with details about the course and next steps.

WHAT'S NEXT?

- Our team will contact you soon.
- We'll help you choose the best batch.
- Start your journey towards a Cloud Career!

YOUR CHOSEN COURSE

MICROSOFT CLOUD ENGINEER

01 MODULE 1		MSSQL & TSQL CONCEPTS Database Fundamentals, T-SQL Queries, Joins, Functions, Views, Indexes & more.
02 MODULE 2		AZURE DATA ENGINEERING ADF, Databricks, Spark, PySpark, ADLS, IoT, Event Hubs, Stream Analytics & more.
03 MODULE 3		FABRIC DATA ENGINEERING FDF, LakeHouse, DataFlow Gen2, PySpark, SparkSQL, KQL Database & more.
04 MODULE 4		REALTIME PROJECTS & MIGRATIONS End-to-End Real-Time Projects & Cloud Migrations (End to End)

WHAT TO EXPECT

- Expert Trainers
- Hands-on Training
- Real-time Projects
- Certification Guidance
- Job Support

SQL SCHOOL
QUALITY TRAINING ASSURED

CALL / WHATSAPP
+91 99514 40801

VISIT OUR WEBSITE
www.sqlschool.com

Let's build your **CLOUD CAREER** together!