SQL School™

Quality Training Assured

Complete Practical; Real-time Job Oriented Training

Azure Cosmos DB Training

	PLAN A	PLAN B	PLAN C
Applicable For (Resume Plan)	**Azure Cosmos DB**	**Azure Cosmos DB Azure Databricks**	**Azure Data Engineering**
Course Curriculum	**Chapters 1 to 10**	**Chapters 1 to 20**	**Chapters 1 to 40**
Azure Cosmos DB Configuration	✓	✓	✓
Table API, Import, Migrate Data	✓	✓	✓
Query Data, Notebooks	✓	✓	✓
Distribute Data in Cosmos	✓	✓	✓
Query Tuning in CosmosDb	✓	✓	✓
Monitoring, Partitions	✓	✓	✓
Data Bricks Architecture	X	✓	✓
Run Spark Jobs, ADB Workspace	X	✓	✓
Virtual Network, File System	X	✓	✓
Azure Synapse with ADB	X	✓	✓
Azure Data Bricks Pools	X	✓	✓
Azure Structured Streaming	X	✓	✓
Azure Delta Lake, delta engine	X	✓	✓
Azure Synapse Architecture	X	X	✓
Azure Analytics Workspace	X	X	✓
Basic ETL Concepts with SSMS	X	X	✓
Bulk Loads with Synapse	X	X	✓
Data Bricks Architecture	X	X	✓
Run Spark Jobs, ADB Workspace	X	X	✓
Data Warehouse Units (DWUs)	X	X	✓
Azure Synapse Management	X	X	✓
Data Factory, Data Engineering	X	X	✓
TOTAL DURATION	2 Weeks	4 Weeks	8 Weeks

Trainer : **Mr. Sai Phanindra T** [15+ Yrs of Real-time Exp]. Profile @ linkedin.com/in/saiphanindra

Azure Cosmos DB & Data Engineering

Training Module		Duration	Plan A	Plan B	Plan C
Module 1	Azure Cosmos DB	2 W	✓	✓	✓
Module 2	Azure Databricks	2 W	X	✓	✓
Module 3	Azure Data Engineering, DP 200, DP 201	2 W	X	X	✓
Total Duration			**2 W**	**4 W**	**8 W**

Module 1: Azure Cosmos DB

Applicable for Plans A, B, C

Chapter 1: Introduction to Cosmos

Azure Cosmos Databases; Document Databases, Uses; Client & Server Side Operations; Globally Distributed Services; Multi-Model Database Env; Key benefits with Cosmos DB; Turnkey Global Distribution; Always On, Consistency Choices; Elastic Scalability of Throughput; Guaranteed Low Latency; Azure Account Registration; Azure Subscription, Resources;

Chapter 2: Azure Cosmos DB Configuration

Azure Cosmos DB Creations; Cosmos DB : API and Data Model; SQL, Mongo DB and Table Options; Graph, Cassandra APIs in Cosmos; Azure Cosmos DB .NET SDK; .NET Framework or .NET Core; Creating Azure Cosmos Account; Connecting to Azure Cosmos account; Configuring Project in Visual Studio; database and a container; Add items to Container; Querying Container, CRUD;

Chapter 3: Table API, Import

Creating Database Accounts; Azure Table Creation Options; Geo-Redundancy in Cosmos; Multi-Region Writes , Availability; Table Id and Throughput Options; Request units per second (RU/s); Data Explorer Options, Audits; Entities : Add Entity Options; Application Cloning in Cosmos; Applications in Visual Studio; Connection Strings in Code; Build, Deploy App in Cosmos;

Chapter 4. Import / Migrate Data

Data Migration Tool in Cosmos DB; Importing Data with Migration Tool; Command Line Installations; dt.exe, Enable Cosmos Table Log; AzCopy Command-Line Utility; Installation and Working Options; Table API Usage with Cosmos DB; Copy Data : One Table API to Another; Generally available Table API; Migrate Big Data to Azure Cosmos; Azure Data Factory for Data Imports; Real-time Challenges : Data Imports;

Chapter 5: Query Data

Data Explorer in Azure Cosmos; Azure Query Playground For SQL; SQL Queries and Syntax Checks; SQL Query Executions in Cosmos; Query Conditions and Predicates; Distribute Data in Cosmos; Add Global Database Regions; Replicate Data, Manual Failover; Store and Access Graph

Data; Working on NoSQL Data; Graph Database in Cosmos; Gremlin API with Azure Cosmos DB;

Chapter 6: Azure Cosmos Notebooks

Azure Cosmos DB Accounts; Enable notebooks in Cosmos; Jupiter Notebooks in Cosmos; Geo Redundancy, Apache Spark; Enabling Notebooks in Cosmos; Run Notebooks, Data Explorer; Query Tuning in Cosmos Db; Optimize Query Cost in Cosmos; Query Cost Throughput; Query Cost Storage Options; Read and Writes Costs; CRUD Analysis, Cost Regions; Query Metricsl;

Chapter 7: Power Shell for Cosmos

Azure Power Shell samples; Azure Cosmos DB - SQL (Core); Account, database and Container; Container with a large Partition key; Get RU/s and Update RU/s; Container with no index policy; Update Account in Cosmos DB; Update Cosmos Account's Regions; Failover Priority, Trigger Failover; Account Keys & Connection Strings; Cosmos Account with IP Firewall; Lock Resources;

Chapter 8: Monitoring, Partitions

Data Modeling, DB Partitions; Modeling Advantages: Cosmos; CQRS Requests for Analysis; User and Posts Container; Create, Retrieve User; Create, Edit Posts & Queries; Create, Edit Posts & Comments; Latency, RU Charge, Performance; Partitioned Containers Migration; System Partition Keys; Monitoring Azure Cosmos DB; Azure Logs; Azure Monitor Metrics in Portal;

Chapter 9: Azure Cosmos DB and ADB

Azure Data Bricks in Virtual Network; Creating Azure Cosmos DB Endpoint; Azure Cosmos DB Primary Connections; Data Explorer, Container Settings; Azure Data Bricks Cluster Setup; Auto Pilot, Workers and Drivers; Loading Data from Azure Cosmos; Azure Cosmos DB Security; Firewall Configurations, Virtual Nets; Private Endpoints, Encryptions; Cross Origin Resources in Cosmos; Certificate Based Authentication; Real-time Security Issues, Solutions;

Chapter 10: Visualizations with Cosmos DB

Azure Cosmos Database Containers; Power BI Reports from Cosmos DB; Using JSON Files with Power BI; Connection Strings and Connectors; Collection Keys and Throughput; Primary Key and Read Only Key; Dynamic Record Items in Power BI; Flattening and Transformations; Document Columns in Power BI; Custom Columns and Text Forms; Report Generation : Visualizations; Report Publish and Share Options; Dishoards, Gateways & Data Refresh;

Module 2: Azure Databricks

Applicable for Plans B, C

Chapter 11 : Introduction to Azure Databricks

Big Data, distributed computing; Introduction to Azure, Advantages; What is Spark?; What is Apache Spark; What is Scala; Need for Azure Data Analytics; Azure Data Engineer Technologies; Need for Azure Databricks; Azure Fundamentals; Azure Account Creation, Validations; Subscription and Resource Options; Resource Groups, Account Access; Billing Options and Resource Access; Azure Active Directory (Azure AD); Notifications, Azure Resources;

Chapter 12: Data Bricks Architecture

ETL Architecture with Data Bricks; Apache Spark analytics platform; Apache Spark Eco System; Spark Core API in Azure DB; Spark SQL Data Frames; Streaming Options, GraphX; Machine Leaning Components; Databricks Runtime; Enterprise security; Databricks ServerLess; Databricks Workflow; Databricks I/O; Integration with Azure; Integration with Data Warehouse; Integration with Power BI;

Chapter 13: Run Spark Job in Azure Data Bricks

Azure Databricks Service; Instance Configuration, Resource Group; Workspace Details, Pricing Tier; Standard, Premium and Trail Editions; Apache Spark and Azure AD Sync; Azure Databricks Networking; Workspace Failures, Retrials; Spark Cluster Creation in ADB; Cluster Mode and Pool Options; DataBricks Runtime Version; AutoPilot and Scaling Options; Worker Type and Capacity; Driver Type, Advanced Options; Notebook Creation, Azure Datasets; PySpark and SQL DB Access; Visualizations, Bar, Plot;

Chapter 14: Azure Databricks workspace in VN

VNet injection Options in Azure Data Bricks; Virtual Network Requirements, Options; Subnets and IP Address for Security; Virtual Address Space, CIDR Block; Public and Private Subnet Usage; Azure Databricks Service, Workspace; Network Security Groups, Arm Templates; Virtual Network Template; Databricks Runtime Architecture; Databricks Runtime Architecture; Better Performance with DBIO; Stronger Security with DBES; Significantly lower operational complexity; Rapid releases and early access; Databricks' SaaS offering; DATABRICKS COMPUTE ENGINES;

Chapter 15. Azure BLOB Storage to ADB

Configure Azure Storage Account; Creating Storage Container; Storage Options and Uploads; Scala notebook in Azure Databricks; Scala Notebooks Creations, Usage; Mounting Blob Storage into ADB; Generating SAS Connections; Data Reads, Transformations, Writes; Azure Databricks File System; Distributed File System; Mounting Storage Objects to DBFS; DBFS Root, Stoage Types; Metadata and Mount Point Paths; Accessing Databricks File System; DBFS API & Local Access; Python and Scala Code for DBFS;

Chapter 16: Azure SQL Server and ADB

Azure SQL Database from Azure Databricks; Deploying Azure SQL Server; Deploying Azure SQL Database; Query Editor and DB Access Options; CSV file upload to Azure Databricks Cluster; Creating Scala Notebooks for Imports; Scala Functions, Display Options; Transformations, Data Load Functions; Azure Cosmos DB and ADB; Azure Data Bricks in Virtual Network; Creating Azure Cosmos DB Endpoint; Azure Cosmos DB Primary Connections; Data Explorer, Container Settings; Azure Data Bricks Cluster Setup; AutoPilot, Workers, Drivers; Loading Data in Cosmos;

Chapter 17: Azure Synapse with ADB

Creating Azure Synapse Database; Loading Data from Azure Synapse; Azure Synapse connector; Spark Driver for Azure Synapse; Spark Executors for Azure Synapse; Azure Synapse to Azure storage; Global Hadoop configuration in ADB; Real-time Recommendations; ETL with Azure Data Bricks; ETL Workflow with Azure Data Bricks; Data Lake Storage Gen2 Storage; Service Principal Usage, ETL; Data Transformations and Lookups; Master Key and Firewall in Synapse; Tenant ID and Azure Active Directory; Authentication Keys, Connections, ETL;

Chapter 18: Azure Data Bricks Pools

Azure Databricks Pools, Clusters; AutoScaling Options, Display Pools; Creating Pools, Pool Configurations; Pool Usage, Instance Options; Pool Configurations, Clusters; Data bricks Runtime

Engines; Pool Size and Auto Terminations; Editing a Pool, Deleting Pools; Structured Streaming; Data Streaming with ADB; Sensors, IoT devices, social networks; Data Stream, Unbounded Tables; Complete Mode. Append Mode; Update Mode, Loading Sample Data; Stream Initialization, Job Streaming; Querying a Data Stream, Spark Jobs;

Chapter 19: Azure Delta Lake, Delta Engine

Azure Delta Lake and Real-time Usage; Creating Delta Table with Python, T-SL; Creating Table using JSON and Delta; Data Partitions into Azure Delta Tables; Data Modifications and Stream Writers; Batch Upserts, Display Table History; Time Travel and Version Control; Table Optimization Techniques; z-order by Columns, Performance; Table Snapshots and Real-time Use; Delta Lake on Azure Data bricks; Ingest Data and Table Data Loads; COPY INTO SQL Command; Auto Loader, Table Batch Concepts; Schema Validations, Vacuum; Concurrency Control, Delta Lake FAQs;

Chapter 20: ADB with ADF

Azure Data Factory for ETL Operations; Azure Data Bricks for ETL Operations; Creating Datasets in Data Factory; Creating Datasets in Azure DataBricks; Creating Data Frames in Azure Data Bricks; Creating Pipelines in Data Factory; Creating Pipelines in Azure DataBricks; Comparing ADB with ADF For ETL; Azure Databricks with Power BI; Using Azure Data Bricks for Data Visualizations; Azure Data Bricks Visualizations, Data Analytics; DataBricks Access in Power BI; Personal Access Token, Cluster Details; Spark Connections with Power BI; Power BI Navigator, Security; Integration, Power Query and DAX;

Module 3: Azure Data Factory [ADF]

Applicable for Plans C

Chapter 21: AZURE BI INTRODUCTION

Basic Vocabulary: ETL, DWH, Data Flow; Azure Data Factory (ADF) Pipelines; SSIS versus ADF : Basic Differences; ADF : Workflow & Data Source Types; Versions and Editions of SQL Server; SSMS Tool Installation, Connections; SSIS DB Creation and Connection; SSIS Catalog and Real-time Use; SSDT (Visual Studio) Installation; Azure Data Studio Tool Installation; SQL Server Connection @ SSDT and ADS Tools; Azure BI Components: ADF,AAS, ADLS; Azure Data Factory Architecture; ADF : Data Processing, Movement; Data Mash up and ETL Components; Data Warehouse, Analysis; End to End Implementation;

Chapter 22: AZURE SQL SERVER, DATABASE, STORAGE

Azure SQL Server Deployments; Azure Logical Servers, Usage; Data Center Locations, Collation; Azure Resource Groups and Usage; Azure SQL Database Deployment; Azure Pricing Tiers: Basic Plan; Standard and Premium Plans; Azure SQL DB Connections @ SSMS; Azure SQL DB Connections @ ADS; Azure Storage : Purpose, Tools; Azure Storage Explorer Tool; Azure Storage Container, Use; Azure Blob Storage and Types; LRS, GRS and RA - GRS Storage; Hot Access, Replication Options; File Storage (Uploads) to Azure; Azure Storage Security Options;

Chapter 23: AZURE DATABASE MIGRATIONS, SYNAPSE

Database and Data Migration Options; Azure Data Migrations with DMA Tool; On-Premise Database Assessments; Migration Scope, Object Scripting; Schema Migration and Data Migrations; Database Exports into BACPAC Files; Azure Storage Container For Imports; Database Imports using Azure Portal; Database Imports using SSMS Tool; Azure Data Warehouse (ADW); Key Benefits of Azure Synapse; Scaling Options, Compute Engine; ADW Connection from SSMS Tool; Massively Parallel Processing (MPP)

Sharding Benefits with MPP in DWH; Control Node and Compute Node; Data Distribution Components;

Chapter 24 : ADF BASICS, COPY DATA TOOL, DATA LAKE

Azure Data Factory Architecture; ADF : Data Processing Components; ADF : Data Movement Components; Data Pipelines and ADF Activities; Data Orchestration and Datasets; ADF Resource Creation in Azure; Pipeline Concept, Linked Services; SSIS Lift and Shift with ADF; Azure Data Factory (ADF) : Purpose; Creating ADF Resource in Azure, Use; Pipeline Creation; Azure Data Lake; Copy Data Tool in ADF Portal, Use; Linked Service Creation in ADF; Dataset Creation, Connection Reuse; Staging Dataset with Azure Storage; ADF Pipeline Deployments, Triggers; Polybase Benefits for Data Staging;

Chapter 25: ADF PORTAL, INTEGRATION RUNTIME

Copy Data From Azure SQL Database; Azure SQL Data Warehouse Sinks; Data Sets Creation, ADF Data Pipelines; Activities : Move & Transform for ETL; Actions and View Activity Runs; ADF Pipeline Monitoring and Refresh; ADF Triggers - Manual, Automatic; Azure Database to Azure Data Warehouse; Data Integration Units [DIU] Usage; Azure Data Factory Runtime Engine; On-Premise Data Extraction and Loads; Self Hosted IR in Copy Data Tool; IR Engine Installation, Key Concepts; IR Connections & Authentication Methods; IR Tool Configuration and Monitoring; Azure Storage Linked Service, Reuse;

Chapter 26: Incremental Loads with ADF - 1

Incremental Data Loads with ETL; Incremental Loads in ADF Portal; Source Data & Watermark Column Watermark Tables : Creation, Use; Time Stamp For Incremental Loads; Configuration Procedures for ETL Using Lookup Activities in ADF; Lookup Transformation with ADF; Watermark Column Extraction Activity Connections in ADF Portal; Expressions and Usage in ADF; Generating New Watermark Columns Lookup Output Row Level Audits; Stored Procedures and Dynamic Values; Parameterized Values for Data Extraction Using ADF Run IDs for Connections; Parameter Imports and Staging; Working with Lookup in ADF Output()& Item() Functions;

Chapter 27 : Incremental Loads with ADF - 2

Incremental Loads - Multiple Tables; Data Prep for Watermark Store; For Each Activity and Real-time use Lookup Activities Use in ADF; Linked Service Parameters, Use; Copy Data Activity for ETL in ADF Stored Procedure Activity for ETL; Data Pipelines Parameters in ADF; Object Name Parameters, Iterations Using Self Hosted IR Engine in ADF; Incremental Loads with Azure Storage; Dynamic Sources and Destinations Data Preview & Dynamic Content; Linked Services & Datasets Creations; ADF Pipeline Executions, Run IDs Dynamic File Formats and Connections; ADF Pipeline, Trigger Execution Runs;

Chapter 28: DATA FLOW IN ADF

Data Flow Task in ADF; Creating Data Flow Task from Controller and Activities; Using Transformations in Data Flow Task; Creating Pipelines and Source / Sink Connections; Debugging with Spark Cluster; Limitations; Data Preview; Tuning Settings for Data Flow; Data Flow Expressions for Lookup, Conditional Split, Sort, Aggregate and Joins; Auto Creation and Manual Creation of Sink Datasets; Using Joins in ADF : Data Flow; Debugging, Validations and Publish;

Chapter 29: DATA WRANGLING IN ADF

Data Wrangling with Azure Data Factory; Differences between Data Flow and Data Wrangling; Data Source and Data Sink Settings; Power Query Interface in ADF; Online Editors and Real-time Usage; Transformations with Data Wrangling Tasks; Linking Data Flow Activity with Wrangling Activity; Task Dependencies; Parameterized Values for Data Extraction Using ADF Run IDs for Connections; Parameter Imports and Staging; Working with Lookup in ADF Output()& Item() Functions;

Chapter 30 : ADF MONITORING, EVENTS, AZURE DATA LAKE

Azure Data Lake Storage Environment; Azure Data Lake Gen 1 and Gen 2; Using ADF to get data from Data lake ADF Integration with Azure Data Store; Copy Data Performance Monitoring Monitor, Activity Runs and Options; Data Partition Recommendations in ADL; Size Limitations and ADF Activities Azure Data Warehouse Management; Threat Detection, Threat Protection; Resource classes, ADW Queries Resource Class Operations & Precedence; Backups and Restores with ADWH; Export and Import Options with ADWH TDE Option and Maintenance Schedules; Security Management, Role Based Access; Azure DWH Alerts and Rule Creations;

Module 4: Azure Synapse

Applicable for Plans C

Chapter 31: Azure Synapse Intro

Need for Azure Synapse; Data engineers, Data scientists; IT professionals, Business analysts; Database administrators; Data security and privacy; Subscription and Azure Resources; Azure Synapse Configuration; Azure SQL Server Creation; Azure Server Firewalls; Synapse Firewalls, Remote Access; Azure Synapse Pool (DWH); Azure Synapse Analytics Engine; DWU Units and Scaling Options; Start and Pause Options;

Chapter 32: Azure Synapse Architecture

Scaling Options, Compute Engine; cDWU & DWU : Real-time Usage, Pricing; Massively Parallel Processing (MPP); Sharding with MPP in DWH; Control Node and Compute Node; Data Distribution Components in ADW; DMS : Data Movement Service, DMVs; Synapse Tables : Types; Integration Tables, Schemas; Table Persistence, Regular Tables; Temporary Tables, External Tables; Distributed Tables, Table Partitions; Hash Functions, Distributed Tables; Round-robin Distributed Tables; Replicated Tables in Azure DWH;

Chapter 33: Azure Analytics Workspace

Azure Synapse and Spark Technologies; Azure Synapse SQL Concepts; Azure Synapse Pipelines, Usage; Creating Workspace in Azure Synapse; Azure Data lake Storage Gen 2; Edit, Delete Options: Command Bars; Azure Synapse Studio : IDE; Storage Accounts Browsing; Query Files and Storage Blob; SQL Scripts, Spark Notebooks; Spark Pools in Synapse; SQL On-demand in Azure Synapse ; Logical Data Warehouse Concepts; Client Tools, PARQUET File Format;

Chapter 34 : Basic ETL Concepts with SSMS

SSMS Client Tool, Access Options; Azure Data Warehouse Access; Performance Levels in Synapse; Client IP Address, Fully Qualified Names; Creating Logins and Users for ETL; COPY Statement and Data Loads; Azure Data Lake Storage; ETL Operations with Azure Storage; Azure Storage Architecture, ADLS; Azure Storage Gen 1 and Gen 2; COPY Statement; Optimize Column Store Compression; Optimize Statistics. Column Store Indexes; Fact Tables & Dimension Tables;

Chapter 35 : Azure Synapse with Data Factory

Azure Data Factory with Synapse; Creating Factory Resource, Options; ADF Versions and Editor Options; Author and Monitor Options in ADF; Creating Pipelines, Copy Data Too l; Loading Data into Azure Synapse; Loading Data From Azure Synapse; Real-time Considerations; Incremental Loads with Files; Table Mapping and Polybase; Staging with ADF; Debugging, Validations; Publishing Options with

Pipelines; Triggering Options, Task Cadence;

Chapter 36: Bulk Loads with Synapse

Bulk Loads with COPY Data; High-throughput Data Ingestion; FIELDTERMINATOR in Copy Data; ROWTERMINATOR in Copy Data; LABEL Options, Data Loads; Load Monitoring in Synapse; Azure Synapse - SQL Tables; Copy Data with Multiple Tables; Authentication Methods; Storage account key with LF; Shared Access Signatures (SAS); CRLF as the row terminator; Managed Identity Authentication; Azure Active Directory Authentication;

Chapter 37 : Data Warehouse Units (DWUs)

Azure Data Warehouse Units; Synapse Resource Pools; Azure Service Level Objective; T-SQL Scripts for SLO Options; Performance Tiers - Gen1, Gen2 ; Data Warehouse Units (DWU); Capacity Limits, Permissions; DWU settings and Reports; Important DMVs, DMFs; Scaling Workflow in Azure Synapse; Performance Tuning Options ; Ordered Clustered Column Store Index; Non-Ordered Column Store Index; Query Performance, CCI, CTAS;

Chapter 38 : Azure Synapse Management

Backups and Restores in Synapse; User-defined restore points; Creating and Using Restore Points; Resting Existing Restore Point; Restoring Deleted Resource Pool; Restore from Geo-Backup SQL Pool; Secure Database in Azure Synapse; Connection security with Synapse; Authentication Modes and Options; Authorization Techniques in Synapse; Database Encryption : Data Security; Transparent Data Encryption (TDE); Certificates and AES-256 Algorithm; Important Security Scripts in T-SQL;

Chapter 39 : Azure Monitoring

Manageability with Synapse SQL pool; Monitor your Azure Synapse Analytics; SQL pool workload using DMVs; VIEW DATABASE STATE; Monitor connections, Memory; Monitor Query Execution; Monitor Active Queries, CTAS; Find top 10 queries longest running queries; Polybase Loads, Catalog Loads; Suspended States, TempDB; sys.dm_pdw_exec_requests; Investigate the Query Plan; Data Movement operations; SQL on the distributed databases; Monitor Waiting Queries;

Chapter 40 : Azure Synapse with ADB

Creating Azure Synapse Database; Loading Data from Azure Synapse; Azure Synapse connector; Spark Driver for Azure Synapse; Spark Executors for Azure Synapse; Azure Synapse to Azure storage; Global Hadoop configuration in ADB; Real-time Recommendations; Azure Synapse with Power BI; Report Generation with Synapse; Real-time Recommendations; Dataset Creation, Data Flows; Power Query and ETL; Visualizations, Data Modeling; Real-time Recommendations;

End to End Implementation, Resume, Mock Interview

Email : contact@sqlschool.com
Skype: SQL School Training Institute
Website: www.sqlschool.com

Call Us (India) : 24 x **7**
+91 9666 44 0801
+91 9666 64 0801

Call Us (USA / Canada) : 24 x **7**
+1 956. 825. 0401

Latest Schedules available at: https://sqlschool.com/Register.html
Address: #101, UMA Residency, Beside Metro Station Gate – D, SR Nagar, Hyderabad, India. Map
Website: **https://sqlschool.com**