



Azure Synapse for End-to-End Business Intelligence Training Course

Ref: #BUI8788



Course Introduction / Overview:

This course provides a comprehensive exploration of Microsoft Azure Synapse Analytics, the limitless analytics service that brings together enterprise data warehousing and Big Data analytics. It is designed to guide participants from foundational concepts to advanced implementation techniques, enabling them to build scalable, end-to-end business intelligence solutions. In the modern data landscape, organizations struggle to unify disparate data sources and analytical workloads. This course directly addresses this challenge by demonstrating how Synapse integrates data ingestion, preparation, management, and serving for immediate BI and machine learning needs. We will delve into the principles of modern data architecture, echoing the concepts discussed by thought leaders like Ralph Kimball in his seminal work "The Data Warehouse Toolkit," and apply them within the cloud-native framework of Azure. BIG BEN Training Center has meticulously structured this program to ensure participants not only learn the tools but also understand the strategic vision behind creating a unified analytics platform, transforming raw data into actionable insights that drive business value. This journey covers everything from setting up a Synapse workspace to building complex data pipelines, querying vast datasets with SQL and Spark, and visualizing results in powerful dashboards.

Target Audience / This training course is suitable for:



- Data Engineers.
- Business Intelligence (BI) Developers.
- Data Analysts and Scientists.
- Database Administrators.
- IT Professionals seeking to upskill in cloud data services.
- Solutions Architects designing data platforms.
- Technology Managers overseeing data initiatives.

Target Sectors and Industries:

- Finance and Banking Sector.
- Healthcare and Life Sciences.
- Retail and E-commerce.
- Manufacturing and Supply Chain.
- Telecommunications and Media.
- Energy and Utilities.
- Government and Public Sector Agencies.

Target Organizations Departments:

- Information Technology (IT).
- Business Intelligence and Analytics.
- Data Engineering and Data Management.
- Finance and Accounting.
- Marketing and Customer Insights.
- Operations and Logistics.
- Research and Development.

Course Offerings:



By the end of this course, the participants will have able to:

- Architect and deploy a complete Azure Synapse Analytics workspace.
- Integrate diverse data sources using Synapse Pipelines and data flows.
- Master data warehousing concepts using dedicated SQL pools for high-performance analytics.
- Perform exploratory data analysis on a data lake using serverless SQL pools.
- Leverage Apache Spark pools for large-scale data engineering and machine learning tasks.
- Secure the Synapse environment using robust access control and data protection features.
- Optimize query performance and manage costs effectively across different compute engines.
- Develop and connect interactive Power BI reports to Synapse for enterprise-level reporting.
- Implement a unified data governance and management strategy within the Synapse ecosystem.

Course Methodology:



The training methodology at BIG BEN Training Center is designed to be immersive, practical, and highly interactive, ensuring participants gain tangible skills they can apply immediately. We move beyond theoretical lectures by focusing on a hands-on, lab-intensive approach where participants build, configure, and manage a real-world Azure Synapse Analytics environment from scratch. Each module is supported by practical exercises that reinforce key concepts, from ingesting data and building pipelines to writing complex queries and creating visualizations. Our expert instructors facilitate engaging discussions, encouraging participants to share their unique challenges and collaborate on solutions. The course incorporates real-world case studies to illustrate how different industries leverage Synapse to solve complex business problems. This problem-solution format helps participants develop critical thinking and architectural design skills. Continuous feedback is a cornerstone of our approach, with instructors providing personalized guidance during lab sessions and group activities. This blend of expert instruction, practical application, and collaborative learning ensures a comprehensive and effective educational experience.

Course Agenda (Course Units):

Unit One: Introduction to Azure Synapse Analytics



- The Modern Data Warehouse and Analytics Landscape.
- Overview of the Azure Synapse Analytics Architecture.
- Exploring the Core Components: SQL Pools, Spark Pools, and Pipelines.
- Navigating the Azure Synapse Studio Workspace.
- Provisioning and Configuring Your First Synapse Workspace.
- Understanding Synapse Storage with Azure Data Lake Storage Gen2.
- Comparing Synapse with Other Azure Data Services.

Unit Two: Data Ingestion and Integration with Synapse Pipelines

- Introduction to Data Integration and ETL/ELT Processes.
- Creating Linked Services and Datasets to Connect to Data Sources.
- Building Your First Pipeline with the Copy Data Activity.
- Orchestrating Complex Workflows with Control Flow Activities.
- Transforming Data at Scale with Mapping Data Flows.
- Ingesting Data from On-Premises and Cloud Sources.
- Scheduling, Triggering, and Monitoring Pipeline Executions.

Unit Three: Enterprise Data Warehousing with Dedicated SQL Pools

- Architecting a Modern Data Warehouse with Dedicated SQL Pools.
- Understanding Massively Parallel Processing (MPP) Architecture.
- Designing Tables and Choosing Distribution and Indexing Strategies.
- Loading Data Efficiently using PolyBase and the COPY Statement.
- Writing and Optimizing T-SQL Queries for Performance.
- Managing Concurrency, Workloads, and Resources.
- Implementing Security Features like Row-Level Security and Data Masking.

Unit Four: Big Data and Exploratory Analytics with Serverless and Spark Pools



- Querying the Data Lake Directly with Serverless SQL Pools.
- Using the OPENROWSET Function to Analyze Various File Formats.
- Introduction to Apache Spark and its Role in Big Data.
- Creating and Managing Apache Spark Pools in Synapse.
- Data Engineering and Analysis using Spark Notebooks with Python and Scala.
- Integrating Spark Pools with SQL Pools for Unified Analytics.
- Leveraging Spark for Machine Learning Model Training and Scoring.

Unit Five: Visualization, Governance, and Management

- Connecting Power BI to Azure Synapse for Interactive Reporting.
- Optimizing Power BI Performance with Direct Query and Composite Models.
- Implementing a Comprehensive Security Model for the Entire Workspace.
- Monitoring Synapse Workspace Activity, Costs, and Performance.
- Integrating with Azure Purview for Data Discovery and Governance.
- Best Practices for CI/CD and Source Control with Synapse.
- Final Project: Building an End-to-End BI Solution.

FAQ:

Qualifications required for registering to this course?

There are no requirements.

How long is each daily session, and what is the total number of training hours for the course?

This training course spans five days, with daily sessions ranging between 4 to 5 hours, including breaks and interactive activities, bringing the total duration to 20 - 25 training hours.

Something to think about:



As the lines between data lakes and data warehouses blur into the 'lakehouse' architecture, how does this paradigm shift impact traditional data governance strategies and the role of the data engineer?

What unique qualities does this course offer compared to other courses?

This course distinguishes itself by offering a holistic, end-to-end perspective on building modern analytical solutions, rather than focusing on isolated features. Its core strength lies in its integrated curriculum, which seamlessly blends the three critical pillars of data analytics: data integration, enterprise data warehousing, and big data processing. We move beyond the "what" and delve deeply into the "why" and "how," teaching participants the strategic thinking required to choose the right tool within Synapse—be it a dedicated SQL pool, a serverless endpoint, or a Spark cluster—for the specific business problem at hand. The curriculum is built around practical, real-world scenarios that mirror the complex challenges organizations face today, ensuring that the skills learned are directly transferable to the workplace. Unlike courses that may provide a purely technical overview, this program emphasizes architectural best practices, performance optimization, and robust security implementation. It is designed not just to train technicians, but to cultivate well-rounded data professionals who can design, build, and manage scalable, secure, and cost-effective analytics platforms that deliver true business value.