



INTELLIGENT BUSINESS STRATEGIES PRESENTS



# Mike Ferguson

## **Modern Data Architecture**

- Building a Foundation for the  
Data Driven Enterprise

[www.intelligentbusiness.biz](http://www.intelligentbusiness.biz)  
[training@intelligentbusiness.biz](mailto:training@intelligentbusiness.biz)

## OVERVIEW

In the last five years, digital transformation has reshaped data management presenting new challenges and opportunities including:

- Data complexity caused by many more data sources with data now stored in SaaS applications, on multiple clouds, on-premises and streaming in from the edge
- Business units are buying data catalogs to help understand, govern and provision data
- Multiple new siloed analytical systems like streaming analytics, lakehouses and graph databases have appeared beyond the data warehouse offering alternative new data architectures
- Data modelling seems to have disappeared
- Data engineering is now happening everywhere and new technologies like Data Fabric and Modern Data Stacks have emerged offering way more than ETL

Simultaneously CEO's now see data and AI as strategic and needed in every part of the business. They are demanding a way forward to speed of development.

So how do you make sense of all this? Is there a future for the data warehouse? Is data modelling dead? With so many competing data architectures, which one is best? How do you meet all requirements and prevent chaos?

This new 2-day class looks at this problem. By attending it, you will learn how you assess your existing environment, look at the considerations, define future requirements and design a new modern data architecture that modernises your data warehouse, and makes it possible to merge it with multiple analytical workloads like data science, streaming analytics, and graph analysis. You will also learn how the architecture allows you to use a data catalog, and data fabric and data observability to build resilient DataOps pipelines to create a data mesh of reusable data products published in a data marketplace that help shorten time to value by enabling new insights and AI to be delivered more rapidly.

## AUDIENCE

CDOs, CIO's, IT Managers, CTOs, Business Analysts, data scientists, BI Managers, data warehousing professionals, enterprise architects, data architects, solution architects, Business Intelligence Specialists, IT strategists, Database administrators, IT consultants



## MODULE 1: WHAT IS DATA ARCHITECTURE?

This starting session looks at what data architect is, why companies need one and what current issues are with respect to data architecture today.

- What is data architecture?
- Why do you need a data architecture?
- What are the differences between data architecture, solution architecture and enterprise architecture?
- What are the main capabilities of a data architecture?
- Reference data architectures and their prod-and cons
  - What is the difference between Batch, Lambda, Kappa, Zero-copy integration?
- Popular data architectures and their pros and cons e.g. data warehouse, data lake, streaming analytics, transaction systems
- What data storage, data processing and data analytics technologies are used in data architecture
- What are the pros and cons of cloud computing in a data architecture?
- How does it all fit together across all environments?

## MODULE 2: ASSESSING YOUR EXISTING DATA ARCHITECTURE

This session looks at how to assess your existing data architecture to understand how well it is serving your business needs from both an operational and analytical perspective. This includes documenting the current problems and identifying stand-alone projects that currently bypass your architecture. The module will cover assessing your existing:

- Data sources you connect to
- Data ingestion techniques
- Data stores in use and their purpose
- Data stores on the cloud Vs on-premises
- Operational applications, processes and data flows
- Analytical data capabilities e.g. data warehouses, data lakes, feature stores
- Data integration data flows for analytics
- Tools used e.g. data catalog, data quality, data governance, ETL, self-service data preparation, data science
- Identifying and documenting issues with your data architecture e.g.,
  - Data complexity

- Siloed analytical systems
- Too many copies of data
- Data that is too big to move
- Brittle data engineering pipelines
- Unmanaged self-service data engineering
- Managing data in a distributed and multi-cloud hybrid computing environment
- Missing data & analytical capabilities
- Integration complexity
- Rapid growth in APIs
- How well does existing data architecture meet your business needs?
  - E.g. Inability to respond in a timely way
  - Are people bypassing your data architecture and creating their own?

## MODULE 3: NEW DATA ARCHITECTURE AND TECHNOLOGIES OPTIONS

Having understood data architecture capabilities, technologies and assessed your existing architecture, this module looks at multiple new data architectures. It also looks at how the emergence of open table formats, advances in SQL, universal data lake APIs, federated query, data catalogs, and data automation are changing data architecture to provide a data foundation for an AI-driven enterprise

- Data architecture options for operational processing
  - Data Integration hub
  - Message bus
- Data architecture options for analytical processing and their pros and cons
  - Zoned Data lake
  - Cloud Data warehouse
  - Lakehouse
  - Cloud Data Warehouse and data lake integration
  - Logical data warehouse
  - Data fabric and Data Mesh
  - Customer data platforms
- Data architecture options for both operational and analytical processing
  - Dataware zero copy integration for data collaboration
- Centralised versus Federated architectures
- Next generation analytical data architecture

- The impact of Open table formats
- The OneTable initiative
- Merging of data warehouse, lake house and streaming on multi-purpose open tables
- Separation of storage from compute
  - Multiple query and data engineering engines on shared data
- Universal execution engine for query and data engineering workloads e.g. Velox
- Universal administration of multiple engines
- Advances in SQL to support graph analytics
- Federated query across open tables on multiple clouds and on-premises storage
- Technology options and their pros and cons
  - Single vendor data and AI stack
    - Data fabric including:
      - Data catalog
      - Collaborative data engineering
      - Data Automation – using a data catalog and generative AI to auto generate pipelines, virtual views, APIs and architecture
      - DataOps and Data Observability
      - Data virtualization for on-demand data integration
    - Data platform
    - AI platform
    - End-to-end data and AI governance
  - Modern data stack
  - Best of breed tools
- Ensuring a common approach to integrating transaction processing applications
- Introducing continuous change and change data capture in transaction systems to lower data latency
- Ensuring tools are integrated to share metadata and enable reuse of business terms, data models, data transformations, data quality rules and data governance policies
- Data warehouse modernisation
  - Modernising ETL processing
  - Data Warehouse automation
  - Migration to the cloud
  - Virtual data marts
- Where does data modelling fit in?
- Data engineering
  - Supporting citizen data engineers, IT professional data engineers and code in data engineering pipelines
  - Establishing a universal approach to data ingestion
  - Incorporating DataOps, CI/CD, data orchestration and data observability into data engineering pipelines
  - Dealing with sensitive data to govern data privacy
  - Establishing a data marketplace to share data

#### **MODULE 4: CONSIDERATIONS FOR DESIGNING, MANAGING AND OPERATING A MODERN DATA ARCHITECTURE**

This module focusses on what you need to think about with respect to designing a data architecture and how you operate and manage it:

#### **MODULE 5: DEFINING YOUR FUTURE DATA ARCHITECTURE REQUIREMENTS**

- Ability to process any type of data
- Support hybrid multi-cloud as standard
- Documented sharable metadata rules on cleansing, transformation, matching and integration with full metadata lineage
- Governed data ingestion, staging and transformation
- Data automation to accelerate development
- Reusable high-quality data products available as a service
- Shared data products to reduce data redundancy
- Data contracts to govern data sharing
- Consuming data products via APIs
- Integrated data governance

- Integrated analytics e.g. BI and ML, ML and streaming, graph analysis and BI, graph and ML, streaming and graph analysis
- Scalability to handle
  - data volumes, data velocity and an increasing number of concurrent users using tools with natural language
  - Generative AI user interfaces
- History of changes to transaction data and/or data products used in transaction systems

### **MODULE 6: DESIGNING AN INNOVATIVE NEW DATA ARCHITECTURE**

In this session we focus on how to design a modern data architecture and merge workloads to create a multi-purpose platform supporting multiple analytical workloads.

- Data architecture design principles
- Defining the key data capabilities needed in your data architecture
- Alignment of data capabilities needed to meet business strategy objectives and priorities
- Designing a data concept model for your business
- Designing your data architecture and data flows to support:
  - Operational processing
  - Analytical processing
  - Knowledge management
- Merging your data warehouse, data lake, lakehouse, streaming and graph to create a new multi-purpose hybrid analytical data platform
  - Defining the core data concepts needed
  - Steps to migrating your data warehouse to the cloud
  - Adoption of open tables formats

- Using data fabric catalog metadata and generative AI to generate pipelines to build data products in open tables
- Supporting streaming data in open tables
- Consuming data products in your cloud data warehouse
- Building virtual data marts from data products
- Building virtual feature stores for data science from data products
- Supporting graph analysis in data warehouses

### **MODULE 7: GETTING STARTED**

This final session looks at what you have to do to get started with a new data architecture initiative. In particular, it looks at:

- Creating an action plan to bring it to life
- Change Vs rebuild?
- What order do you do this in?
- How do you minimise impact on the business while you re-architect?
- How do you deal with a backlog of change when you are also trying to re-architect?
- Pros and cons of build Vs automating development
- What new skills are needed?
- Delivering new business value whilst re-architecting
- How do you involve business professionals in the re-architecting effort?

## PRESENTER



**Mike Ferguson** is Managing Director of Intelligent Business Strategies Limited. As an independent IT industry analyst and consultant, he specialises in BI / analytics and data management. With over 40 years of IT experience, Mike has consulted for dozens of companies on BI/Analytics, data strategy, technology selection, data architecture, and data management. Mike is also conference chairman of Big Data LDN, the fastest growing data and analytics conference in Europe and a member of the EDM Council CDMC Executive Advisory Board. He has spoken at events all over the world and written numerous articles. Formerly he was a principal and co-founder of Codd and Date Europe Limited – the inventors of the

Relational Model, a Chief Architect at Teradata on the Teradata DBMS. He teaches popular master classes in:

- Big Data Architecture & Technology
- Building a Competitive Data Strategy for A Data-Driven Enterprise
- Data Catalogs – Governing and Provisioning Data in a Data Driven Enterprise
- Data Virtualisation & the Logical Data Warehouse
- Data Warehouse Modernisation
- Embedded Analytics, Intelligent Apps & AI Automation
- How to Govern Data Across a Distributed Data Landscape
- Migrating your Data Warehouse to the Cloud
- Modern Data Architecture
- Practical Guidelines for Implementing a Data Mesh (Data Catalog, Data Fabric, Data Products, Data Marketplace)
- Real-Time Analytics

## ONSITE TRAINING

All training can be given as onsite education, tailored to your company's requirements. For further details please contact us at [training@intelligentbusiness.biz](mailto:training@intelligentbusiness.biz)

