



INTELLIGENT BUSINESS STRATEGIES PRESENTS



**Mike Ferguson**

**Practical Guidelines for  
Implementing a Data Mesh**

(Data Catalog, Data Fabric, Data  
Products, Data Marketplace)

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

## OVERVIEW

Most companies today are storing data and running applications in a hybrid multi-cloud environment. Analytical systems tend to be centralised and siloed like data warehouses and data marts for BI, Hadoop or cloud storage data lakes for data science and stand-alone streaming analytical systems for real-time analysis. These centralised systems rely on data engineers and data scientists working within each silo to ingest data from many different sources, clean and integrate it for use in a specific analytical system or machine learning models. There are many issues with this centralised, siloed approach including multiple tools to prepare and integrate data, reinvention of data integration pipelines in each silo and centralised data engineering with poor understanding of source data unable to keep pace with business demands for new data. Also, master data is not well managed.

To address these issues, a new approach has emerged attempting to accelerate creation of data for use in multiple analytical workloads. That approach is Data Mesh. Data Mesh is a decentralised business domain-oriented approach to data ownership and data engineering to create a mesh of reusable data products that can be created once and shared across multiple analytical systems and workloads. A Data Mesh can be implemented in a number of ways. These include using one or more cloud storage accounts on cloud storage, on an organised data lake, on a Lakehouse, on a data cloud, using Kafka or using data virtualisation. Data products can then be consumed in other pipelines for use in streaming analytics, Data Warehouses or Lakehouse Gold Tables, for use in business intelligence, data science and other analytical workloads.

This 2-day class looks at data mesh in detail and examines its strengths, and weaknesses. It also looks at the strengths and weaknesses of data mesh implementation options. Which architecture is best to implement this? How do you co-ordinate multiple domain-oriented teams and use common data infrastructure software like Data Fabric to create high-quality, compliant, reusable, data products in a Data Mesh. Also, how can you use a data marketplace to share data products? The objective is to shorten time to value while also ensuring that data is correctly governed and engineered in a decentralised environment. It also looks at the organisational implications of Data Mesh and how to create sharable data products for master data management AND for use in multi-dimensional analysis on a data warehouse, data science, graph analysis and real-time streaming analytics to drive business value? Technologies discussed includes data catalogs, data fabric for collaborative development of data integration pipelines to create data products, DataOps to speed up the process, data orchestration automation, data marketplaces and data governance platforms.

## AUDIENCE

This seminar is intended for business data analysts, data architects, chief data officers, master data management professionals, data scientists, IT ETL developers, and data governance professionals. It assumes you understand basic data management principles and data architecture plus a reasonable understanding of data cleansing, data integration, data catalogs, data lakes and data governance.

## LEARNING OBJECTIVES

Attendees will learn about:

- Strengths and weaknesses of centralised data architectures used in analytics
- The problems caused in existing analytical systems by a hybrid, multi-cloud data landscape
- What is a data mesh and how does it differ from a data lake and a data lakehouse?
- What benefits does data mesh offer and what are the implementation options?
- What are the principles, requirements, and challenges of implementing these approaches?
- How to organise to create data products in a decentralised environment so you avoid chaos

- The critical importance of a data catalog in understanding what data is available
- How business glossaries can help ensure data products are understood and semantically linked
- An operating model for effective federated data governance
- What software is required to build, operate and govern a data mesh of data products for use in a data lake, a data lakehouse or data warehouse?
- What is data fabric software, how does it integrate with data catalogs and connect to data in your data estate
- An Implementation methodology to produce ready-made, trusted, reusable data products
- Collaborative domain-oriented development of modular and distributed DataOps pipelines to create data products
- How a data catalog and automation software can be used to generate DataOps pipelines
- Managing data quality, privacy, access security, versioning, and the lifecycle of data products
- Publishing semantically linked data products in a data marketplace for others to consume and use
- Consuming data products in an MDM system
- Consuming and assembling data products in multiple analytical systems like data warehouses, lakehouses and graph databases to shorten time to value

## MODULE 1: WHAT IS DATA MESH, A DATA LAKE AND A LAKEHOUSE? WHY USE THEM?

This session looks at the challenges facing companies trying to become data driven and at the strengths and weaknesses of current centralised data architectures used in analytics. It then introduces Data Lakes, Data Lakehouse and Data Mesh as potential ways to address current problems. It explores the pros and cons of each of these, explains how you can enable the creation of trusted, reusable data products in a Data Mesh using different architecture options such as cloud storage accounts, data lakes, cloud data platforms, etc., for use in multiple analytical workloads. It also asks if combining multiple architecture approaches is advantageous or not.

- Data complexity in a hybrid, multi-cloud environment
- The growth in new data sources
- Centralised data architectures in use in existing analytical systems – data warehouse, data lakes
- Pros and cons of Data Lakes
- The merging of data warehouses and data lakes
- The move from just data science to multi-purpose data lakes
- What is a Data Lakehouse?
- How does a Data Lakehouse work?
- Pros and cons of a Data Lakehouse
- What is a Data Mesh?
- Data Mesh principles
- How does decentralised Data Mesh work?
- What is a data product?
- What types of data product can you build?
- Decentralised development of data products
- Pros and cons of Data Mesh
- What are the challenges with this decentralised approach?
- Is data management software ready for Data Mesh?
- How will Data Mesh impact your current IT organisation and data culture?
- Is federated data governance possible?
- What are the architectural options for implementing Data Mesh?

- Implementing Data Mesh on Data Cloud Vs Kafka Vs Cloud Storage Vs Lakehouse Vs Data Virtualisation
- Implementation requirements to create data products
  - Federated operating model
  - Common business vocabulary
  - Data producers and data consumers
  - Architecture independence
  - A unified data platform for building any pipeline to process any data
  - DataOps – component-based CI/CD pipeline development
  - Distributed pipeline execution
  - Reusable, semantically linked data products
  - Governance of a distributed data landscape
- Key technologies: Data Fabric, Data Catalogs, data classifiers, Data Marketplace, Data Automation tools
- Vendor's offerings in the market – Alation, AWS, BigID, Cambridge Semantics, Collibra, Dremio, Global IDs, Google, IBM, Informatica, Microsoft, Oracle, Qlik, Qlik (Talend), SAP, SAS, Software AG, Starburst Data, StreamSets, IDERA WhereScape, Truist Zalani

## MODULE 2: METHODOLOGIES FOR CREATING DATA PRODUCTS

This session looks at how to produce business ready, reusable data products for use by data consumers in multiple analytical use cases who need it to drive business value. It also looks at how master data products can also be produced for use in master data management.

- Creating a program office
- Decentralised development of data products in a Data Mesh using a cloud storage based Data Lake, a Lakehouse, a data cloud or data virtualisation
- The special and critical case of master data
- A best practice step-by-step methodology for building reusable data products

- How does structured, semi-structured and unstructured data impact the methodology?
- Applying DataOps development practices to data product development?

### MODULE 3: USING A BUSINESS GLOSSARY TO DEFINE DATA PRODUCTS

This session looks at how you can create common data names and definitions for your data products in a business glossary so data consumers can understand the meaning of the data produced and available in a Data Mesh. It also looks at how business glossaries have become part of a data catalog

- Why is a common vocabulary relevant?
- Data catalogs and the business glossary
- The Data Catalog market, e.g., Alation, Amazon Glue, Cambridge Semantics ANZO Data Catalog, Colibra Catalog, data.world, Denodo Data Catalog, Google Data Catalog, Hitachi Vantara Lumada, IBM Watson Knowledge Catalog, Informatica Axon and EDC, Microsoft Purview Data Catalog, Oracle, Qlik Catalog, SAP, Truist Zalon Data Platform
- Roles, responsibilities, and processes needed to manage a business glossary
- Jumpstarting a business glossary with a data concept model
- Defining data products using glossary terms
- Using a catalog and glossary to ensure data products are semantically linked?

### MODULE 4: STANDARDISING DEVELOPMENT AND OPERATIONS IN A DATA MESH, DATA LAKE OR LAKEHOUSE

This session looks at how to standardise the setup in each business domain to optimise development of data products in a Data Mesh,

- The importance of a program office
- Implementing Data Mesh on a single cloud Versus a hybrid multi-cloud environment
- Implementing Data Mesh on a Data Lake or Lakehouse
- Standardising the domain implementation process – ingest, process, persist, serve

- Creating zones in a domain cloud storage account, a Data Lake or Lakehouse to produce and persist data products in a Data Mesh
- Using Kafka as an option to persist data products in a Data Mesh
- Selecting Data fabric software as a platform for domain-oriented teams to build data products
- Steps-by-step data product development
  - Data source registration
  - Automated data discovery, data quality profiling, sensitive data detection, governance classification, lineage extraction and cataloguing
  - Data ingestion
  - Global and domain policy creation for federated governance of classified data
  - Data product pipeline development
  - Standardising on best practice and taking complexity away from citizen data engineers
  - Data product publishing for consumption

### MODULE 5: BUILDING DATAOPS PIPELINES TO CREATE MULTI-PURPOSE DATA PRODUCTS

This session looks at designing and developing modular DataOps pipelines to produce trusted data products using Data Fabric software

- Collaborative pipeline development & orchestration to produce data products
- Designing component based DataOps pipelines to produce data products
- Using CI/CD to accelerate development, testing and deployment
- Designing in sensitive data protection in pipelines
- Processing streaming data in a pipeline
- Handling schema drift in a pipeline
- Processing unstructured data in a pipeline using ML
- Generating data pipelines using Data Automation tools
- Using data observability to monitor and improve pipelines
- Making data products available for consumption in a Data Mesh using a data marketplace

- Publishing data products on Kafka as topics – Is this a good option?
- The Enterprise Data Marketplace – enabling information consumers to shop for data
- Serving up trusted data products for use in multiple analytical systems and in MDM
- Consuming data products in other pipelines for use in data warehouses, lakehouses, data science sandboxes, graph analysis, Kafka and MDM

## MODULE 6: IMPLEMENTING FEDERATED DATA GOVERNANCE TO PRODUCE AND USE COMPLIANT DATA PRODUCTS

With data highly distributed across so many data stores and applications, on-premises, in multiple clouds and the edge, many companies are struggling to govern data throughout its lifecycle. This is critically important in a Data Mesh where federated computational data governance is a fundamental principal, data product development is decentralised, and data products are shared and consumed across the organisation. It is also paramount across the whole hybrid multi-cloud data landscape. This session looks at how this can be achieved.

- What is involved in federated data governance?
- How do you implement this across a hybrid, multi-cloud distributed data landscape?
- Understanding compliance obligations
- Types of data governance policies
- Understanding Global Vs local policies when creating a Data Mesh
- Defining sensitive data types
- Using the data catalog for automated data profiling, quality scoring and sensitive data type classification
- Defining and attaching policies to classified data in a data catalog
- Creating sharable master data products and reference data products for MDM and RDM
- Ensuring data quality in data product development
- Protecting sensitive data in data product development for data privacy compliance
- Governing data product version management
- Governing consumer access to data products containing sensitive data
- Prevent accidental oversharing of sensitive data products using DLP
- Governing data retention of data products in-line with compliance and legal holds
- Monitoring and data stewarding to ensure policy enforcement
- Data catalog and data fabric technologies to help govern data across a distributed data landscape
  - Types of data governance offerings
  - Alation, Ataccama, Collibra, Confluent Schema Registry and Catalog, Dataguise
  - Google Cloud IAM, Data Catalog, BigQuery, Dataplex and DLP
  - IBM Cloud Pak for Data, Watson Knowledge Catalog, Optim & Guardium
  - Hitachi Vantara
  - Immuta, Imperva
  - Informatica EDC and Axon
  - Microsoft Purview
  - Okera, OneTrust Data Governance Suite
  - Oracle Enterprise Data Management Cloud
  - Privitar
  - SAP Data Intelligence
  - Software AG StreamSets DataOps Platform
  - Qlik (Talend), TopQuadrant

## 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)

