

TECHNOLOGY TRANSFER PRESENTS

MIKE FERGUSON

**Designing, Operating and Managing
a Multi-purpose Data Lake**

ONLINE LIVE STREAMING

NOVEMBER 22-23, 2021



info@technologytransfer.it
www.technologytransfer.it

ABOUT THIS SEMINAR

Most organisations today are dealing with multiple silos of information. These include cloud and on-premises based transaction systems, multiple Data Warehouses, Data Marts, Master Data Management (MDM) systems, NoSQL databases, Cloud storage, Hadoop and Content Management systems. In addition the number of data sources is increasing. It is not surprising that many companies are struggling with knowing what data is available, whether they can trust it and how they go about integrating it. Many have ended up managing information in silos with different tools being used to prepare and integrate data. In addition, both IT and Business users are now integrating data and the danger is total chaos.

The question is, do we let this continue or is there another way to govern and unify data across an increasingly complex data landscape that can shorten time to value?

This 2-day seminar looks at how a new DataOps approach proposes a new Data Architecture where data is ingested and organised in a Data Lake and IT and Business Users collaborate to unify and build ready-made trusted data assets that are published in a Data Marketplace and available to others to consume and use to drive value. The objective is to reduce time to value, avoid reinvention while governing and unifying data quickly in a multi-Cloud, multiple Data Store, hybrid Computing environment.

AUDIENCE

This seminar is intended for business data analysts doing self-service data integration, data architects, chief data officers, master data management professionals, content management professionals, database administrators, big data professionals, data integration developers, and compliance managers who are responsible for data management. This includes metadata management, data integration, data quality, master data management and enterprise content management.

The seminar is not only for 'Fortune 500 scale companies' but for any organisation that has to deal with Big Data, small data, multiple data stores and multiple data sources. It assumes that you have an understanding of basic data management principles as well as a high level of understanding of the concepts of data migration, data replication, metadata, data warehousing, data modelling, data cleansing, etc.

LEARNING OBJECTIVES

Attendees will learn:

- How to define a strategy for producing trusted Data as-a-Service in a distributed environment of multiple Data Stores and data sources
- How to organise data in a centralised or distributed data environment to overcome complexity and chaos
- How to design, build, manage and operate a Logical or centralised Data Lake within their organisation
- The critical importance of an Information Catalog in understanding what data is available as a service
- How data standardisation and Business Glossaries can help make sure data is understood
- An operating model for effective distributed information governance
- What technologies and implementation methodologies they need to get their data under control
- Collaborative curation of trusted, ready-made data products and publishing them in a data marketplace for people to shop for data
- How to apply methodologies to get Master and Reference Data, Big Data, Data Warehouse, Data and unstructured data under control irrespective of whether it be on-premises or in the Cloud
- Fuelling rapid 'last mile' analytical development to reduce time to value

OUTLINE

1. Establishing a Data Strategy for rapid unification of trusted Data Assets

This session introduces the Data Lake together with the need for a data strategy and looks at the reasons why companies need it. It looks at what should be in your data strategy, the operating model needed to implement, the types of data you have to manage and the scope of implementation. It also looks at the policies and processes needed to bring your data under control.

- The ever increasing distributed data landscape
- The siloed approach to managing and governing data
- IT data integration, self-service data preparation or both? – data governance or data chaos?
- Key requirements for Data Management
 - Structured data – Master, reference and transaction data
 - Semi-structured data – JSON, BSON, XML
 - Unstructured data - Text, video
 - Re-usable services to manage data
- Dealing with new data sources - Cloud data, sensor data, social media data, smart products (the Internet of Things)
- Understanding scope of your Data Lake
 - OLTP systems
 - Data Warehouses
 - Big Data systems
 - MDM and RDM systems
 - Data virtualisation
 - Streaming data
 - Enterprise Content Management
- Building a business case for Data Management
- Defining an enterprise data strategy
- A new collaborative approach to governing, managing and curating data
- Introducing the Data Lake and Data Refinery
- Data Lake configurations – what are the options?
 - Centralised, distributed or Logical
- DataOps – a component-based approach to curating trusted data products
- The rising importance of an Information Catalog
- Key technology components in a Data Lake and information supply chain – including data fabric software
- Using Cloud storage or Hadoop as a data staging area and why it is not enough
- Implementation run-time options – the need to ex-

- cute in multiple environments
- Integrating a Data Lake into your Enterprise Analytical Architecture

2. Information Production Methodologies

Having understood strategy, this session looks at multiple methodologies and the technologies needed to help apply it to your structured and multi-structured data to bring it under control. It also looks at how platforms like Hadoop and common data services provide the foundation to manage information across the enterprise

- Information production and information consumption
- A best practice step-by-step methodology structured Data Governance
- Why the methodology has to change for semi-structured and unstructured data
- Methodologies for structured vs. multi-structured data

3. Data Standardisation & the Business Glossary and the Information Catalog

This session looks at the need for data standardisation of structured data and of new insights from processing unstructured data. The key to making this happen is to create common data names and definitions for your data to establish a shared business vocabulary (SBV). The SBV should be defined and stored in a Business Glossary and is important for information consumers to understand published data in a Data Lake. It also looks at the emergence of more powerful Information Catalog software and how business glossaries have become part of what a Catalog offers.

- Semantic data standardisation using a shared Business Vocabulary within an Information Catalog
- SBV vs. taxonomy vs. ontology
- Why is an SBV relevant in a Data Lake and a Logical Data Warehouse?
- Approaches to creating an SBV
- Business Glossary products storing SBV business data names Alteryx Connect Glossary, ASG, Collibra, Global IDs, Informatica Axon, IBM Information Governance Catalog, SAS Business Data Network

and more.

- Planning for a Business Glossary Organising data definitions in a Business Glossary
- Key roles and responsibilities - getting the operating model right to create and manage an SBV
- Formalising governance of business data names, e.g. the dispute resolution process
- Business involvement in SBV creation
- Beyond structured data - from Business Glossary to Information Catalog
- What is an Information Catalog?
- Why are Information Catalogs becoming critical to Data Management?
- Information Catalog technologies, e.g. Alation, Alteryx Connect, Amazon Glue, Apache Atlas, Collibra Catalog, Cambridge Semantics ANZO Data Catalog, Denodo Data Catalog, Google Data Catalog, IBM Information Governance Catalog & Watson Knowledge Catalog, Informatica EDC & Live Data Map, Microsoft Azure Data Catalog, Qlik Catalog, Hitachi Vantara Lumada, Zoloni Data Platform
- Information Catalog capabilities

4. Organising and Operating the Data Lake

This session looks at how to organise data to still be able to manage it in a complex data landscape. It looks at zoning, versioning, the need for collaboration between business and IT and the use of an Information Catalog in managing the data.

- Organising data in a centralised or distributed Data Lake
- Creating zone to manage data
- New requirements for managing data in centralised and distributed Data Lakes
- Creating collaborative Data Lake projects
- Hadoop as a staging area for enterprise data cleans-ing and integration
- Core processes in Data Lake operations
- The data ingestion process
- Tools and techniques for data ingestion
- Implementing systematic disparate data and data relationship discovery using Information Catalog software
- Using domains and Machine Learning to automate and speed up data discovery and tagging
- IBM Watson Data Catalog, Informatica CLAIRE, Silwood, Hitachi Vantara Lumada, Data Catalog

- Automated profiling and tagging and cataloguing of data
- Automated data mapping
- The data classification and policy definition processes
- Manual and automated data classification to enable governance
- Using tag based policies to govern data

5. The Data Refinery Process

This session looks at the process of refining and curating data in an information supply chain to produce trusted data products.

- What is a Data Refinery?
- Key Requirements for refining data
- The need for multiple execution engines to run in multiple environments
- Options for refining data - ETL versus self-service data preparation
- Key approaches to scalable ETL data integration using Apache Spark
- Self-service data preparation tools for Spark and Hadoop e.g. Alteryx Designer, Azure Data Factory Wrangling Flows, Enterprise Data Preparation, IBM InfoSphere Advanced Data Preparation, Azure Data Factory Wrangling Flows, DataRobot (Paxta), SAS Data Preparation, Tableau Prep Builder, Tamr, Talend, Trifacta
- Automated data profiling using analytics in data preparation tools
- Executing Data Refinery jobs in a distributed Data Lake using Apache Beam to run anywhere
- Approaches to integrating IT ETL and self-service data preparation
- Apache Atlas Open Metadata & Governance
- Joined up analytical processing from ETL to analytical pipelines
- Publishing data and data integration jobs to the Information Catalog
- Mapping produced data of value into your DW and Business Vocabulary
- Data provisioning - publishing trusted, ready-made data products into an Enterprise Data Marketplace
- The Enterprise Data Marketplace – enabling information consumers to shop for data
- Provisioning trusted data using data virtualisation, a Logical Data Warehouse and on-demand informa-

tion services

- Consistent Data Management across Cloud and on-premise systems

6. Unifying Big Data, Master Data & Data Warehouses data to drive Business value

This session looks at how the data refining processes can be applied to governing, unifying and provisioning data across a Big Data, MDM and traditional Data Warehouses to drive new Business value. How do you deal with very large data volumes and different varieties of data? How do you load and process data in Hadoop? How should low-latency data be handled?

- A walk through of end-to-end Data Lake operation to create a Single Customer View
- Types of Big Data & Small Data needed for single customer view and the challenge of bringing it together
- Connecting to Big Data sources, e.g. Web logs, clickstream, sensor data, unstructured and semi-structured content
- Ingesting and analysing clickstream data
- The challenge of capturing external customer data from social networks
- Dealing with unstructured data quality in a Big Data environment
- Using graph analysis to identify new relationships
- The need to combine Big Data, Master Data and data in your Data Warehouse
- Matching Big Data with Customer Master Data at scale
- Governing data in a Data Science environment

7. Information Audit & Protection – Governing Data across a distributed Data Landscape

Over recent years we have seen many major brands suffer embarrassing publicity due to data security breaches that have damaged their brand and reduced customer confidence. With data now highly distributed and so many technologies in place that offer audit and security, many organisations end up with a piecemeal approach to information audit and protection. Policies are everywhere with no single view of the policies associated with securing data across the enterprise. The number of administrators involved is often difficult to

determine and regulatory compliance is now demanding that data is protected and that organisations can prove this to their auditors. So how are organisations dealing with this problem? Are the same data privacy policies enforced everywhere? How is data access security co-ordinated across portals, processes, applications and data? Is anyone auditing privileged user activity?

This session defines this problem, looks at the requirements needed for Enterprise Data Audit and Protection and then looks at what technologies are available to help you integrate this into your data strategy

- What is Data Audit and Security and what is involved in managing it?
- Status check - Where are we in data audit, access security and protection today?
- What are the requirements for enterprise data audit, access security and protection?
- What needs to be considered when dealing with the data audit and security challenge?
- Automatic data discovery and the information catalog – a huge help in identifying sensitive data
- What about privileged users?
- Using a Data Management platform and Information Catalog to govern data across multiple data stores
- Securing and protecting data using tag based policies in an Information Catalog
- What technologies are available to protect data and govern it? – Apache Knox, Cloudera Ranger, Dataguise, Hortonworks Ranger, IBM (Watson Data Platform, Knowledge Catalog, Optim & Guardium), Immuta, Informatica Secure@Source, Imperva, Micro Focus, Okera, Privitar
- Can these technologies help in GDPR?
- How do they integrate with Data Governance programs?
- How to get started in securing, auditing and protecting your data

INFORMATION

<p>PARTICIPATION FEE</p> <p>€ 1100</p> <p>The fee includes all seminar documentation.</p> <p>SEMINAR TIMETABLE</p> <p>9.30 am - 1.00 pm 2.00 pm - 5.00 pm</p>	<p>HOW TO REGISTER</p> <p>You must send the registration form with the receipt of the payment to: TECHNOLOGY TRANSFER S.r.l. Piazza Cavour, 3 - 00193 Rome (Italy) Fax +39-06-6871102</p> <p>PAYMENT</p> <p>Wire transfer to: Technology Transfer S.r.l. Banca: Cariparma Agenzia 1 di Roma IBAN Code: IT 03 W 06230 03202 000057031348 BIC/SWIFT: CRPPIT2P546</p>	<p>GENERAL CONDITIONS</p> <p>DISCOUNT</p> <p>The participants who will register 30 days before the seminar are entitled to a 5% discount.</p> <p>If a company registers 5 participants to the same seminar, it will pay only for 4.</p> <p>Those who benefit of this discount are not entitled to other discounts for the same seminar.</p> <p>CANCELLATION POLICY</p> <p>A full refund is given for any cancellation received more than 15 days before the seminar starts. Cancellations less than 15 days prior the event are liable for 50% of the fee. Cancellations less than one week prior to the event date will be liable for the full fee.</p> <p>CANCELLATION LIABILITY</p> <p>In the case of cancellation of an event for any reason, Technology Transfer's liability is limited to the return of the registration fee only.</p>
---	---	--

MIKE FERGUSON
DESIGNING OPERATING AND MANAGING A MULTI-PURPOSE DATA LAKE

November 22-23, 2021

Registration fee: 1100€

first name

surname

job title

organisation

address

postcode

city

country

telephone

fax

e-mail



Stamp and signature

Send your registration form with the receipt of the payment to:
Technology Transfer S.r.l.
Piazza Cavour, 3 - 00193 Rome (Italy)
Tel. +39-06-6832227 - Fax +39-06-6871102
info@technologytransfer.it
www.technologytransfer.it

If anyone registered is unable to attend, or in case of cancellation of the seminar, the general conditions mentioned before are applicable.

SPEAKER

Mike Ferguson is Managing Director of Intelligent Business Strategies Limited. As an analyst and consultant he specialises in Business Intelligence and Enterprise Business Integration. With over 39 years of IT experience, he has consulted for dozens of companies on Business Intelligence Strategy, technology selection, enterprise architecture, and data management. 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 and European Managing Director of Database Associates. He teaches popular master classes in Operational Business Intelligence, New Technologies in DW and BI for the Agile Enterprise, Big Data Multi-Platform Analytics, Master Data Management and Enterprise Data Governance.