



Untangling Data Mesh, Fabric, and Lakehouse

Modern Methods for Cloud Data Warehousing

- Origins, drivers, meaning, and detailed functionality of data lakehouse, fabric, and mesh
- Benefits – business and technical – and lessons learned for each approach
- An in-depth comparison of data fabric and mesh with data warehouse, lake, hub, and lakehouse
- Possibilities and challenges of new database and data management technologies in Cloud, on-premises, and hybrid environments
- The central role of context-setting information, knowledge systems, and metadata
- Using data virtualization and preparation as tools for integration of all types of content and data in Cloud, on-premises, and hybrid environments
- Practical planning and implementation steps from data warehouse / lake to data lakehouse, fabric, or mesh



Two day seminar by
Barry Devlin

AdeptEvents

LANGUAGE
English

VENUE
Utrecht / Hilversum / Virtual

TIME
9:30 – 17:00 hours

REGISTRATION
www.adeptevents.nl



Untangling Data Mesh, Fabric, and Lakehouse

Modern Methods for Cloud Data Warehousing

The transition to cloud data warehousing is well under way. Close examination confirms that the underlying principles are the same as traditional data warehousing. So why are there three competing architectural design patterns: Data Fabric, Data Mesh, and Data Lakehouse? And why not a simple Cloud Data Warehouse?

But what are these design patterns? Are they truly novel or simply marketing hype? In essence, all three are competing technological responses to the growing needs of digital transformation. If you are facing the urgent challenges of delivering high-value, consistent, and near real-time information across cloud and on-premises environments, understanding these approaches and their differences is critical. Learn if you should move beyond your existing data warehouse or lake and, if so, how.

Barry Devlin explains and positions data fabric, mesh, and lakehouse, as well as other concepts, old and new, using a logical digital information systems architecture framework. Exploring existing and emerging technologies as well as organisational issues, methodologies, and implementation approaches, Barry will help you decide if one of these new approaches is right for your business needs, existing technical environment, and current skills.

The seminar

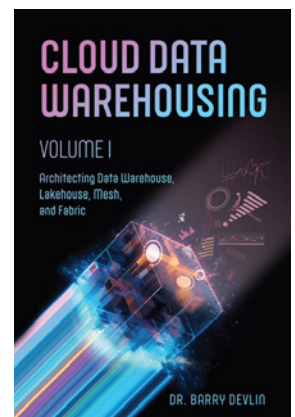
In the era of cloud data warehousing, three new architectural design patterns have recently emerged to challenge traditional data warehouse, logical data warehouse, and data lake thinking. Each has its pros and cons, strengths and weaknesses. *Data fabric* proposes active metadata and knowledge graphs to power a logical data warehouse approach. *Data mesh* suggests that a domain-oriented, self-service approach based on microservices thinking should be adopted, eliminating data copies almost entirely. *Data lakehouse*, as the name implies, attempts to combine the best of data warehouse and data lake, as well as promising transactional consistency within its scope. Today's business, with its conflicting needs for data timeliness vs. consistency, immediate vs. correct decisions, and

information-informed competition, places extensive new demands on cloud data warehousing systems. These demands have led to the emergence of the lakehouse, fabric, and mesh architectural patterns honed for the complex distributed and network-centric environments that are already common. However, the prevalence of legacy systems, historical data management issues, as well as existing and evolving complications in information meaning and usage mean that traditional approaches and knowledge cannot be readily abandoned.

The key questions, therefore, are if and how data lakehouse, fabric, and mesh address these new needs, how differently each does that, where they improve on existing approaches or create new problems, and how they can coexist with or replace established data warehouses and lakes.

To answer these questions, Dr. Barry Devlin compares and contrasts all these architectural design patterns, old and new, using as a foundation the conceptual and logical architectures first defined in "Business unIntelligence." Existing and emerging technologies for data storage, preparation, and virtualization; data catalogs and knowledge graphs; and other tools, both on-premises and Cloud, are described and analysed. Also explored are a wide variety of organisational issues, methodologies, and implementation approaches that are often as important in assessing solutions as the underlying technologies.

The course has been revised and updated during 2023 during the writing of "Cloud Data Warehousing, Volume I: Architecting Data Warehouse, Lakehouse, Mesh, and Fabric" published June 2023 and "Cloud Data Warehousing, Volume II: Implementing Data Warehouse, Lakehouse, Mesh, and Fabric". Barry is offering a special discount on Volume I to course attendees.





What you will learn

- Origins, drivers, meaning, and detailed functionality of data lakehouse, fabric, and mesh
- Benefits – business and technical – and lessons learned for each approach
- An introduction to the technical rationale, structure and components of the base conceptual and logical architecture and its business and technical uses
- An in-depth comparison of data fabric, mesh, and lakehouse with data warehouse, lake, and logical data warehouse patterns based on this architecture
- Possibilities and challenges of new database and data management technologies in cloud, on-premises, and hybrid environments
- The central role of context-setting information, knowledge systems, and metadata

- Adaptive processes as the basis for data preparation, information creation, and insight discovery
- Using data virtualization and preparation as tools for integration of all types of content and data in cloud, on-premises, and hybrid environments
- Practical planning and implementation steps from data warehouse / lake to data lakehouse, fabric, or mesh.

Who should attend?

- Enterprise, systems, solutions and data warehouse architects
- Systems, strategy and business intelligence managers
- Data warehouse, data lake and IT systems designers and developers
- Data and database administrators
- Tech-savvy business analysts.



Course description

1. The Path to the Present

- A brief history of decision-making support
- Data warehouse (hub & spoke and star schema) and marts: business, technology drivers, and challenges
- Logical data warehouse: business, technology drivers, and challenges
- The emergence and impact of big data, the Internet of Things and artificial intelligence
- Data lake: business, technology drivers, and challenges
- Introducing data lakehouse, data fabric, and data mesh

2. Architectural View I: Information as Foundation

- Conceptual and logical architecture views
- Thinking Spaces: Information, Process, and People
- Key information considerations – timeliness/consistency, structure/context, and reliance/usage
- From silos and layers to pillars – supporting multiple storage and processing technologies
- Information types: process-mediated data, human-sourced information, machine-generated data, and context-setting information
- Introducing the “classic” architectural design patterns

3. Information Storage Technologies and Architectural Design Patterns

- Relational technology and its evolution
- Hadoop and NoSQL Stores
- Cloud storage technologies

4. Deep Dive I: Data Lakehouse

- Vendor-defined architecture described
- Exploring data lakehouse – what does it really mean?
- Products, tools, and techniques
- Object stores as a foundation
- Example implementations
- Architectural design pattern
- Pros and cons

5. Context is Everything: Modernising Metadata

- From DIKW to the manifest meaning model
- Information, knowledge, meaning, decision, action

- Metadata as context-setting information – sources and stores, tools and techniques, including data glossary, data dictionary, and data catalog
- Modelling, ontologies, and knowledge graphs

6. Deep Dive II: Data Fabric

- Analyst-defined architecture described
- Exploring data fabric – what does it really mean?
- Products, tools, and techniques
- Active metadata, semantic knowledge graphs, and data virtualization
- Example implementations
- Architectural design pattern
- Pros and cons

7. Architectural View II: Process as Intermediary

- Conceptual and logical architecture views
- Merging of business and IT processes
- Defining adaptive, closed-loop processes across business and IT
- The new role of users in “application development” – opportunities and dangers
- Evolution of SOA to orchestration

8. Evolution of Information Preparation

- Data Preparation, ETL, Replication, Data Warehouse Automation, Wrangling, and Data Virtualisation
- Data pipelines and data ops
- Batch, real-time and Lambda architectures
- Streaming, messaging, immutable logs and Kappa architecture

9. Deep Dive III: Data Mesh

- Analyst-defined architecture described
- Exploring data mesh – what does it really mean?
- Domain-driven design and data as a product
- Focus on governance and development responsibilities
- Products, tools, and techniques
- Technology considerations, including Service Oriented Architecture and Microservices
- Architectural design pattern
- Pros and cons

Course description

10. Architectural View III: People and Organisation

- Conceptual architecture view
- Models for decision making and action taking
- Emerging AI and ML implications
- Support by data lakehouse, data fabric, and data mesh

11. Migration and Implementation

- Evolution – not revolution – a methodology for successful migration
- The Staged Implementation Roadmap (SIR)
- Organisational considerations; changes in IT culture and responsibilities
- Migration routes to data lakehouse, fabric, and mesh

12. Ethical and Economic Considerations

- Augmenting and/or Automating decision making and action taking with AI
- Technical, legal, and ethical issues with data collection, anonymisation and surveillance
- Bias, privacy erosion, facial recognition, affective computing and other dangers
- Surveillance capitalism and other economic uses of big data and AI

13. Conclusions and Wrap-up

- Comparing data lakehouse, fabric, and mesh with one another
- Outlook for further developments
- Readiness assessments for data lakehouse, data fabric, and data mesh
- Conclusions.



BARRY DEVLIN

Dr. Barry Devlin is among the foremost authorities on business insight and one of the founders of data warehousing, having published the first architectural paper in 1988. With over 40 years of IT experience, including 20 years with IBM as a Distinguished Engineer, he is a widely respected industry analyst, consultant, speaker and author of the seminal book, "Data Warehouse – from Architecture to Implementation" (1997) and the ground-breaking "Business unIntelligence – Insight and Innovation beyond Analytics and Big Data" (2013). His latest book "Cloud Data Warehousing, Volume I: Architecting Data Warehouse, Lakehouse, Mesh, and Fabric" was published mid 2023.

As founder and principal of 9sight Consulting, Barry provides strategic consulting and thought leadership to buyers and vendors of BI solutions. He continues to develop new architectural models and provides international, strategic thought-leadership. He is now based in Cornwall, and is a senior data architect with the UK's meteorological service (Met Office).

Information



DATE AND TIME

The workshop will take place once or twice a year with the exact date and time available on our website. The programme starts at 9:30 am and ends at 5:00 pm on both days. Registration commences at 8.30 am and we recommend that you arrive early.

If we need to run 3 virtual half day sessions, the programme starts at 9:00 am and ends at 1:00 pm. Please log in well in advance to check your video and audio settings.

VENUE

Adept Events works with several venues in and near Amersfoort and Utrecht. We strive to provide you with the location details as soon as possible. The exact venue will be on our website and in the confirmation e-mail that you will receive one week prior to the event. Always check our website prior to your departure to ensure you have the exact location and directions.

The virtual seminars and workshops will be offered through a live video stream from our video studio. The virtual classroom is equipped with multiple professional cameras and microphones and we have a two person crew available so that we can offer a similar experience to that of the traditional classroom and, more important, still offer adequate interaction with the speaker.

HOW TO REGISTER

Please register online at www.adeptevents.nl. For registering by print, please scan the completed registration form and send this or your Purchase Order to seminars@adeptevents.nl. We will confirm your registration and invoice your company by e-mail therefore please do not omit your e-mail address when registering.

REGISTRATION FEE

Taking part in this two-day workshop will only cost € 1305 when registering 30 days beforehand and € 1450 per person afterwards (excl. 21% Dutch VAT). This also covers documentation, lunch, tea/ coffee.

Note: This seminar may also be offered 'Online' or as 'Face-to-face and live streaming'. In that situation, the prices for attending online differ from the prices listed here. On the Registration Fee page of our website you will always find the current rates for all available formats of this seminar.

In completing your registration form you declare that you agree with our **Terms and Conditions**.

Team discounts

Discounts are available for group bookings of two or more delegates representing the same organization made at the same time. Ten percent off for the second and third delegate and fifteen percent off for all delegates when registering four or more delegates (all delegates must be listed on the same invoice).

This cannot be used in conjunction with other discounts.

All prices are VAT excluded.

PAYMENT

Full payment is due prior to the workshop. An invoice will be sent to you containing our full bank details including BIC and IBAN. Your payment should always include the invoice number as well as the name of your company and the delegate name.

Payment by credit card is also available. Please mention this in the Comment-field upon registration and find further instructions for credit card payment on our **customer service page**.

Cancellation Policy

Cancellations must be received in writing at least three weeks before the commencement of the workshop and will be subject to a € 75,- administration fee. It is regretted that cancellations received within three weeks of the workshop date will be liable for the full workshop fee. Substitutions can be made at any time and at no extra charge.

Cancellation Liability

In the unlikely event of cancellation of the workshop for any reason, Adept Events' liability is limited to the return of the registration fee only. Adept Events will not reimburse delegates for any travel or hotel cancellation fees or penalties. It may be necessary, for reasons beyond the control of Adept Events, to change the content, timings, speakers, date and venue of the workshop.

INFORMATION



+31(0)172 742680



<https://adeptevents.nl/mfl-en>



seminars@adeptevents.nl



@AdeptEventsNL / <https://twitter.com/AdeptEventsNL>



<https://linkedin.com/company/adept-events>



<https://facebook.com/AdeptEventsNL>



Visit our Business Intelligence and Data Warehousing website <https://biplatform.nl> and download the App



Visit our website on Software Engineering, <https://release.nl> and download the App

IN-HOUSE TRAINING

Would you like to run this course in-company for a group of persons? We can provide a quote for running an in-house course, if you offer the following details. Estimated number of delegates, location (town, country), number of days required (if different from the public course) and the preferred date/period (month). Please find more info on the **In-house page on our website**.