

Bring your DataOps to the Data Lakehouse

Benefits & Challenges



#Conf42 #DevOps #2025

Jorge Loaiciga-Rodriguez
Senior Platform Engineer / Consultant @ Perelyn



Introduction

Companies face recurring challenges in data projects

- Inconsistent data quality
- delayed insights
- poor alignment between stakeholder/developers
- etc.

Applying DataOps and embracing Data Lakehouse architectures isn't just another fix — it's a paradigm shift that systematically reduces these challenges while driving insight delivery and trust in data.





Learning Outcomes

1

Understanding

An understanding of what DataOps is, why it matters and how it approaches problems in modern data architectures.

2

Perspective

An appreciation for the huge opportunity that DataOps represents for your organisation or carrier.



About Myself

+6

Years of experience as data professional

12

Orgs supported¹ in their IT transformation journey

+24

Months in love with applied DataOps principles



Jorge
Loaiciga-Rodriguez
Senior Platform
Engineer / Consultant

jorge.loaiciga-rodriguez
@perelyn.com



¹ Consultation & Implementation



Agenda

1 Case-Study	2 Key Concepts	3 Adoption Challenges	4 Conclusions
 <p>FICSIT Incorporated</p> <ul style="list-style-type: none">▪ The issues addressed by DataOps▪ The grand vision behind DataOps	<ul style="list-style-type: none">▪ Cost▪ Performance▪ Under-standability	<ul style="list-style-type: none">▪ Technical▪ Organisational	<ul style="list-style-type: none">▪ Summary▪ What comes next?



Understanding the Problem

The story of FICSIT Inc.

FICSIT



“FICSIT Inc. is cost-effective & efficient. We do not waste”

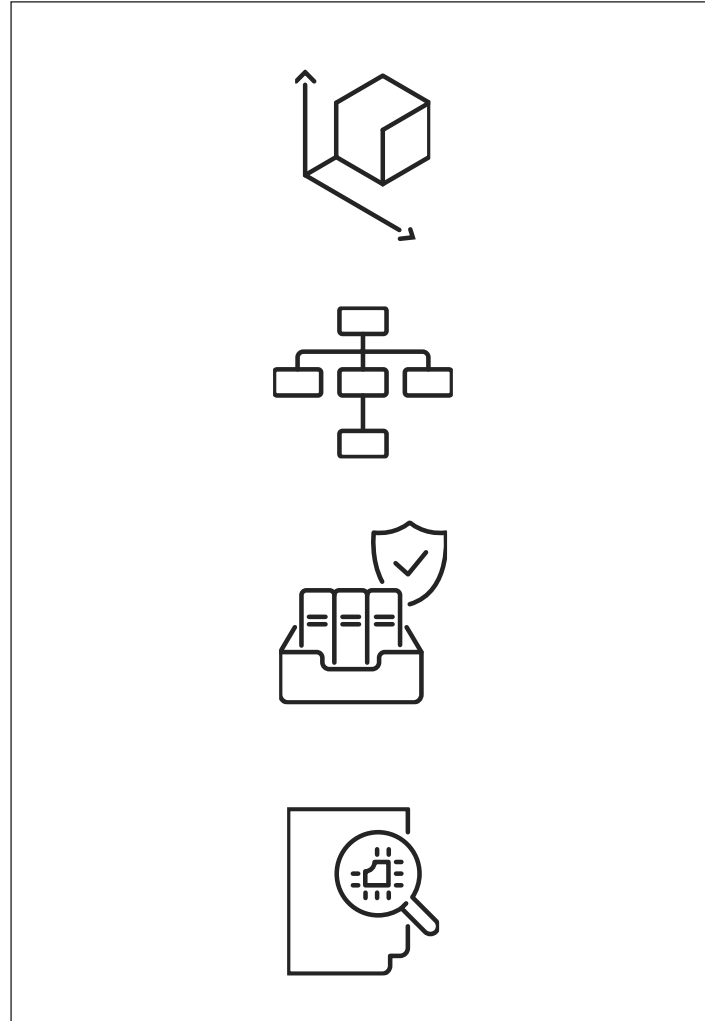
- Headquarters on Earth¹
- 500k+ employees
- R&D, Aerospace Engineering & Mining
- Data-driven and cutting-edge
 - 1st Chief Data Officer launched digital transformation program to drive growth
 - Central data department with ~500 engineers
 - Investment in Cloud-native compute/storage, and Martian BI tooling

¹ Allegedly

IT Department



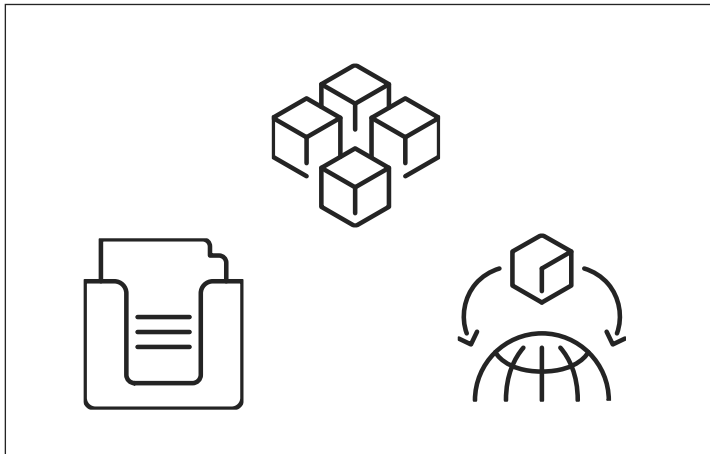
Data Engineering Department




Customer Support Department



Manufacturing Department

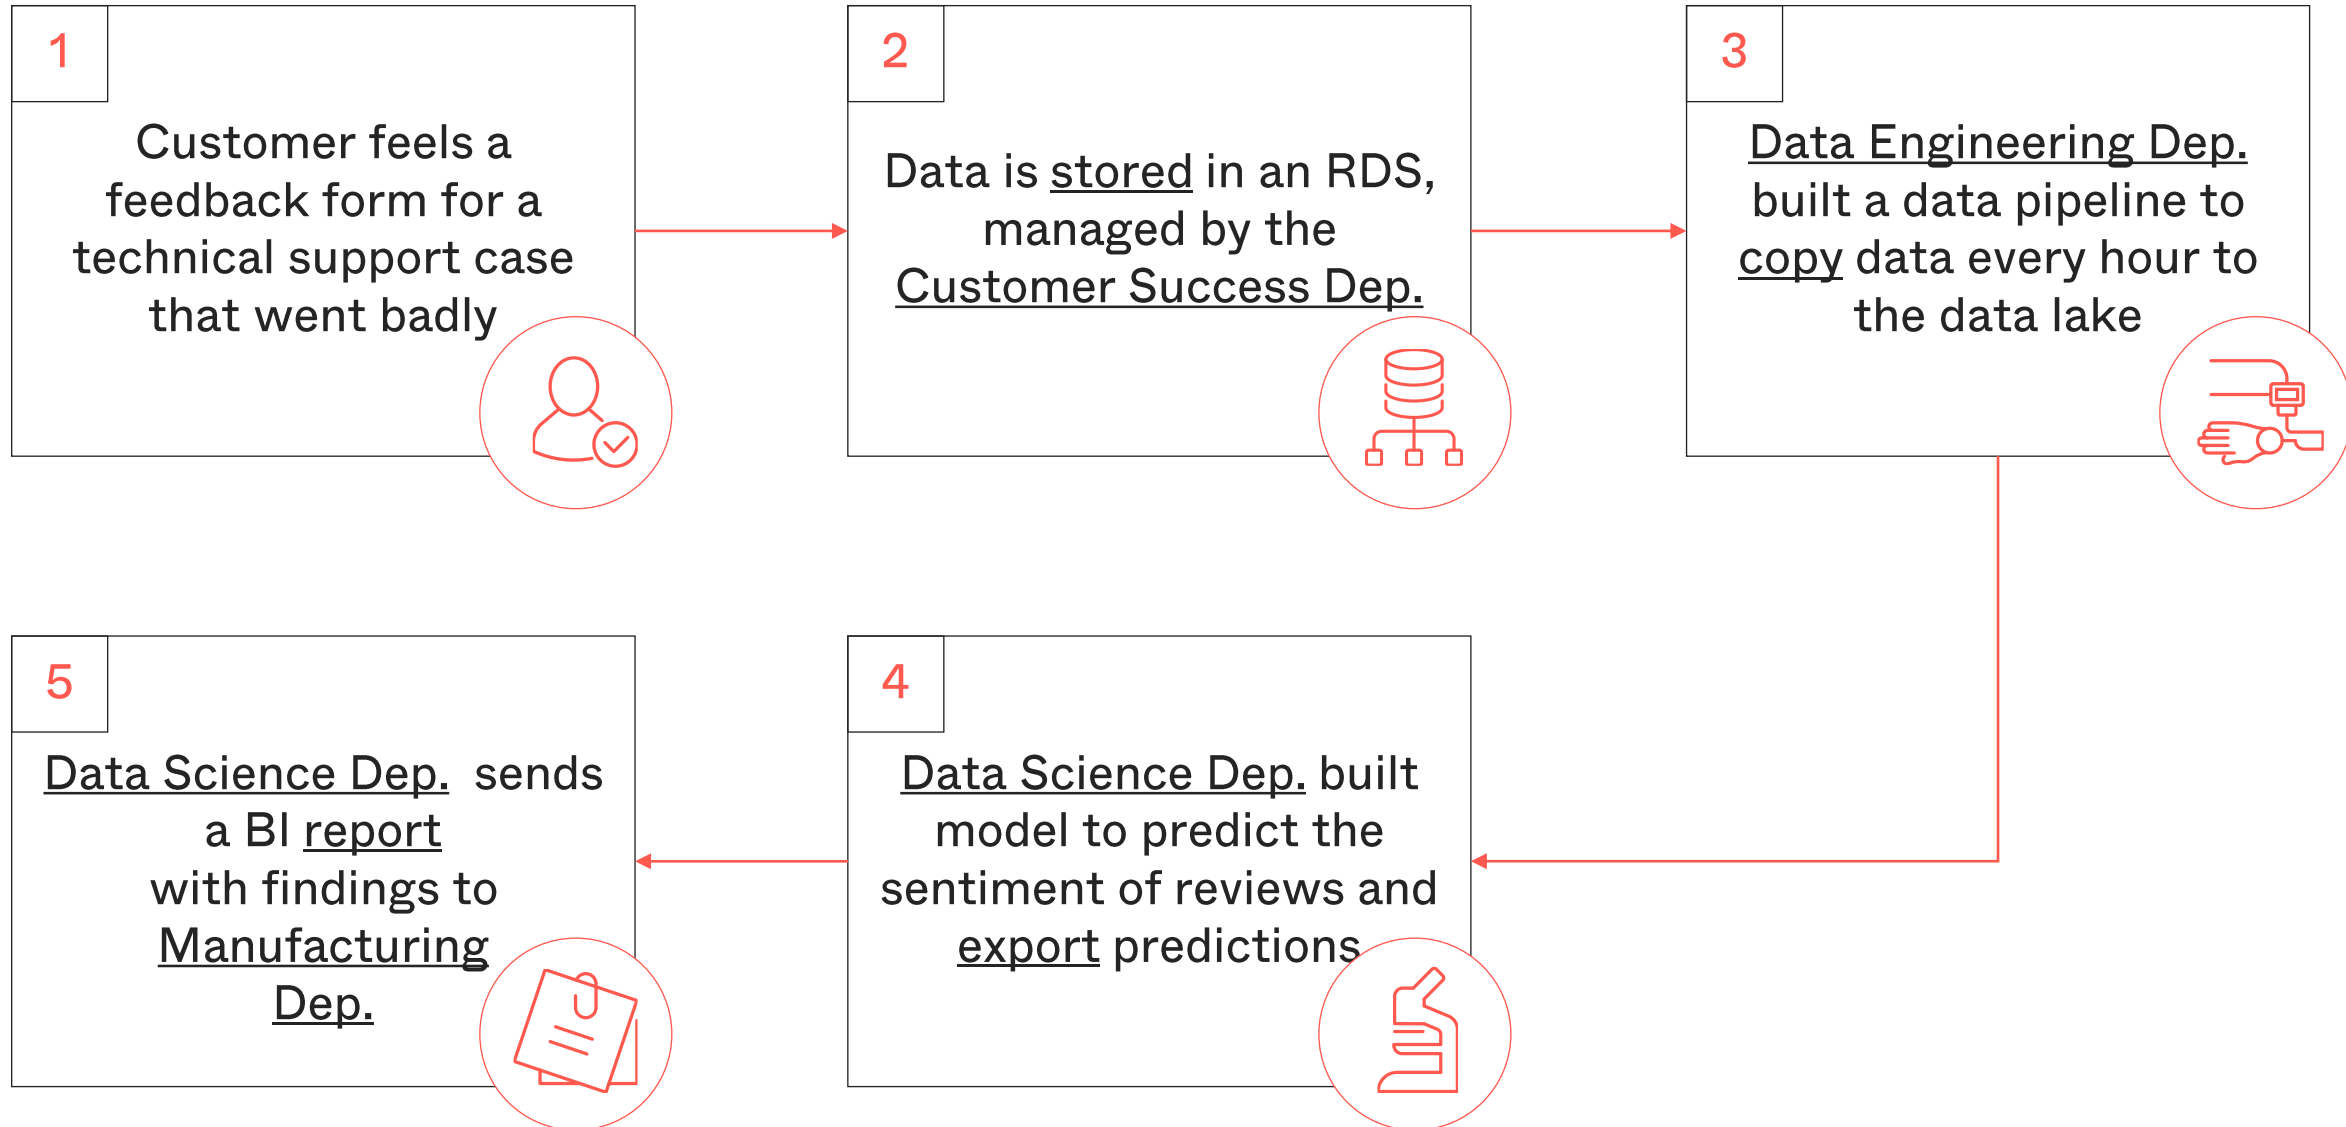


Data Science Department






Example Workflow: Customer Support Reviews



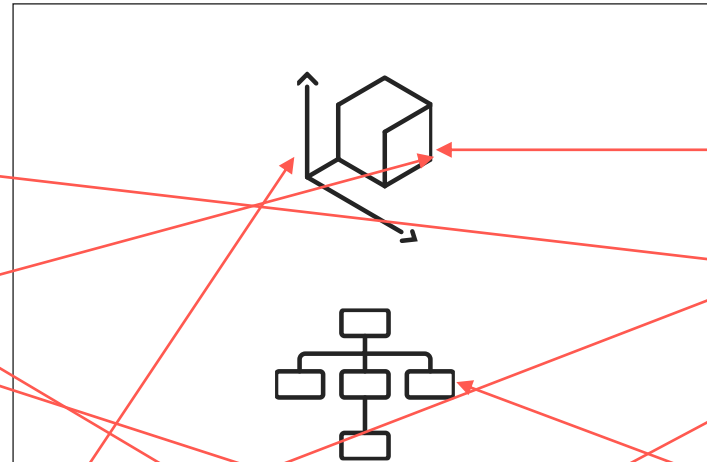


Scaling up the Workflow Count

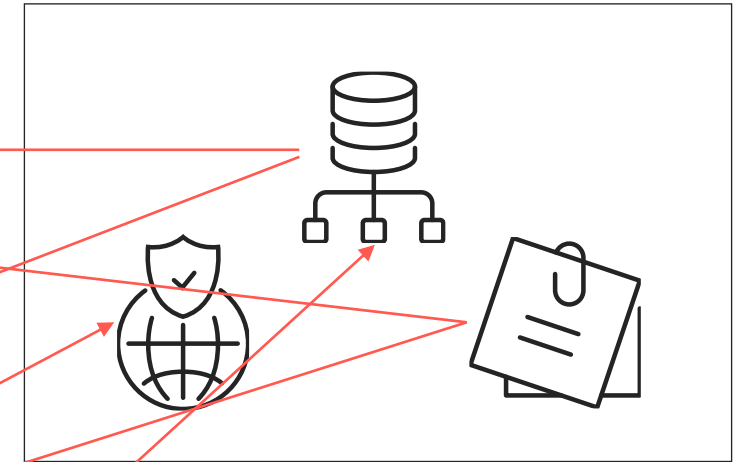
IT Department



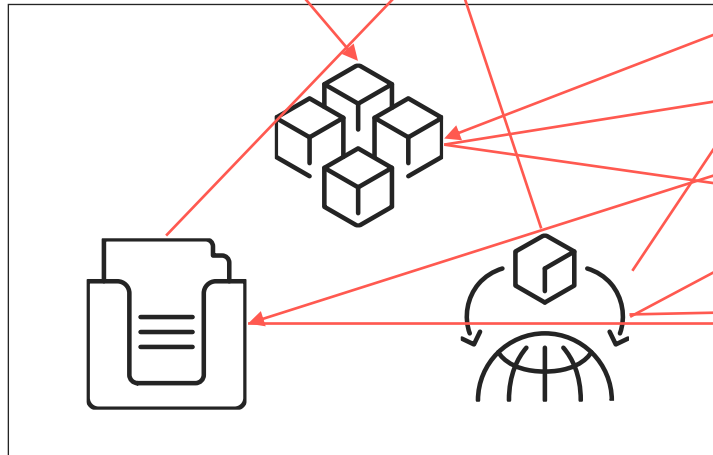
Data Engineering Department




Customer Support Department



Manufacturing Department



Data Science Department



“The IT architecture grew organically, there are data silos everywhere”

– Chief Executive Officer

“We maintain hundreds of pipelines, moving data between all departments, cyclically dependent even!”

– Head of Data Engineering

“Many obscure column names and transformations across data stores. 6 months in and onboarding is still going”

– Data Science Newest Recruit

“Too many systems,
all with different
permission schemes.
The monthly bill is also
unpredictable”

– Head of IT

“Can I trust data in that dashboard. Last time I checked it looked funny. It also takes forever to load”


– Managers (Dashboard Users)



Data Management Capability Assessment



1	Data Ingestion	2	Data Storage	3	Data Modeling	4	Data Reporting
	<ul style="list-style-type: none">No data silos1 data copy2 data formats1 access control schemeNo metadata management6 orchestration & monitoring panels		<ul style="list-style-type: none">4 data silos4 data copies3 data formats2 access control schemeNo metadata management4 orchestration & monitoring panels		<ul style="list-style-type: none">2 data silos5 data copies4 data formats3 access control scheme2 metadata management panels2 orchestration & monitoring panels		<ul style="list-style-type: none">1 data silo5 data copies5 data formats4 access control scheme3 metadata management panels1 orchestration & monitoring panels



“What if we introduced DataOps and molded into a Data Lakehouse architecture?”



Ideal Scenario

1	Data Ingestion	2	Data Storage	3	Data Modeling	4	Data Reporting
1 lake							
1 data copy ¹							
1 central panel for metadata management							
1 data format (open-source)							
1 central orchestration & monitoring panel							
1 central access control scheme							

¹ Not counting backups for disaster recovery


How are the following affected?

1. Costs
2. Performance
3. Understandability



DataOps on Data Lakehouse

Transactional



Source: chartio.com/blog



Storage & Syncing

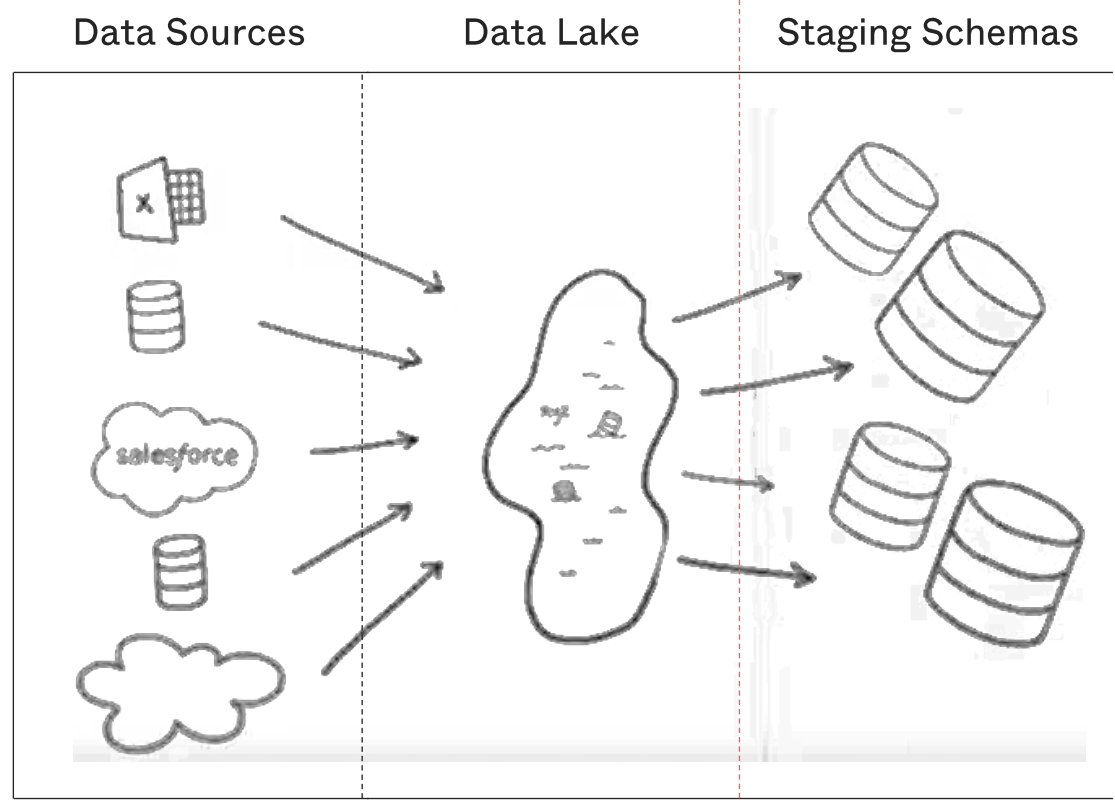
- System integration tends to produce far more unstructured/semi-structured data than structure data.
- In the pre-cloud era, storage was expensive vs. In modern data platforms, storage is cheap.
- Open-source data formats like Iceberg and Delta Lake are now widely supported. Free high-performance read and write (columnar, binary, ACID-support).



DataOps on Data Lakehouse

Transactional

Operational



Source: chartio.com/blog



Staging Schemas

- The strength of the system is introduced at this step
- Well-define stages for each of the common data refining operations
 1. Language translation
 2. Column renaming
 3. Column type formatting
 4. Storage standard format
 5. Storage partitioning & compression
 6. Model quality testing and freshness testing
 7. Business logic introduction
 8. Inter-dependency orchestration¹
- Easy to build, troubleshoot and **version**

¹ With tools such as dbt



DataOps on Data Lakehouse

Transactional

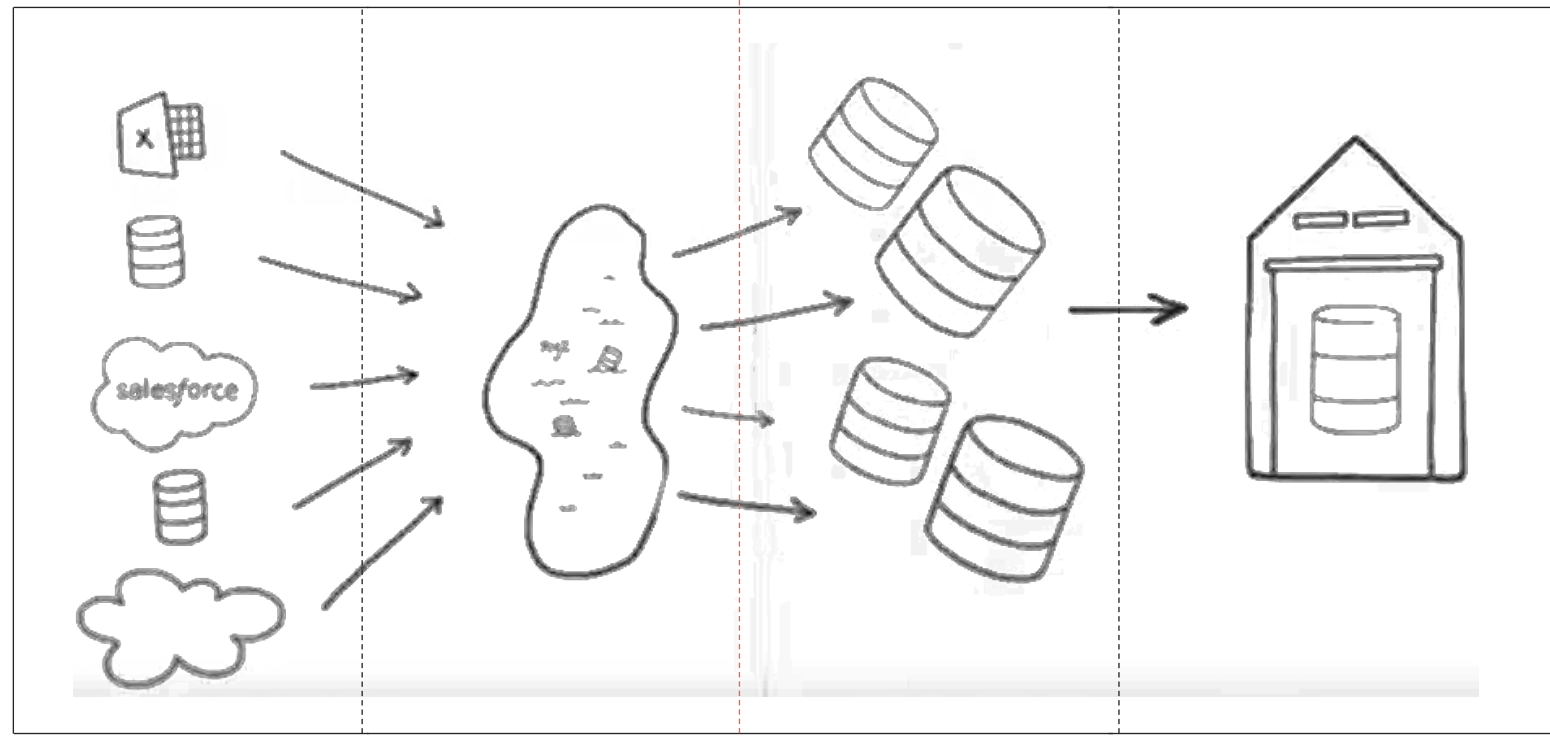
Operational

Data Sources

Data Lake

Staging Schemas

Data Warehouse





Operational Data Lakehouse

- If staging schemas are perfect for data exploration, the warehouse is perfect for data exploitation.
- Offers the computation power needed to meet even the time-critical use-cases.
- Can introduce fancy materialization strategies (incremental loading of records), or Machine Learning-based transformations.
- **Data** it reads and writes **stay within the one-lake.**



DataOps on Data Lakehouse

Transactional

Operational

Analytical


Data Sources

Data Lake

Staging Schemas

Data Warehouse

Wide Tables





Data Consumption from Wide Tables

- 25-50% more performant than running joins¹
- Extremely intuitive for analysts and business users
- BI tools tend to favor wide tables as input. Setup and maintainability of dashboards suffer otherwise.
- ✘ You might end up with duplicated columns with different **versions** of the data values

¹ fivetran.com/blog



DataOps on Data Lakehouse

Transactional

Operational

Analytical

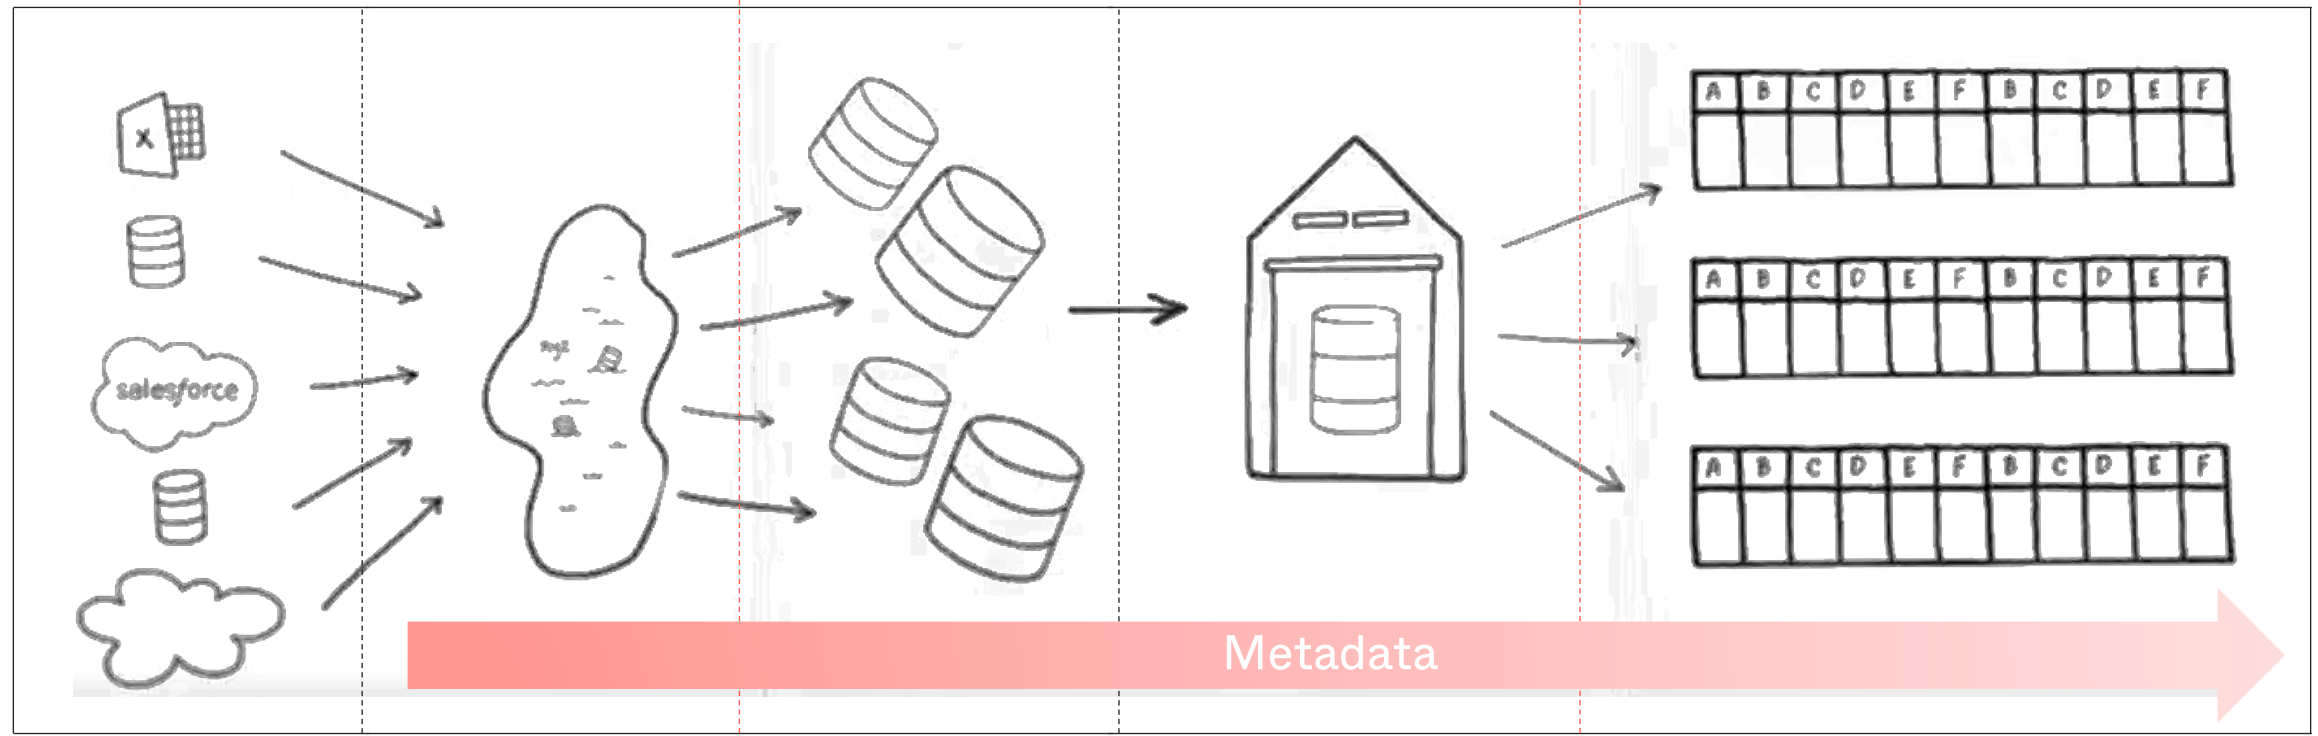
Data Sources

Data Lake

Staging Schemas

Data Warehouse

Wide Tables





Typical Challenges of DataOps adoption

Technical

- Existing long-running licences for enterprise tools.
- Too much trivial testing (i.e testing only for not null and unique values)
- Model version deprecation without proper communication to downstream users
- Fail to create a mockup of new use-cases in coordination with stakeholders

Organizational

- The conversions report you built does not match with the one from that other team
- Cool, but I am not a technical person..., why don't you talk to the engineers?
- We've have worked with tool X for years. Now you want me change into Y?



Change



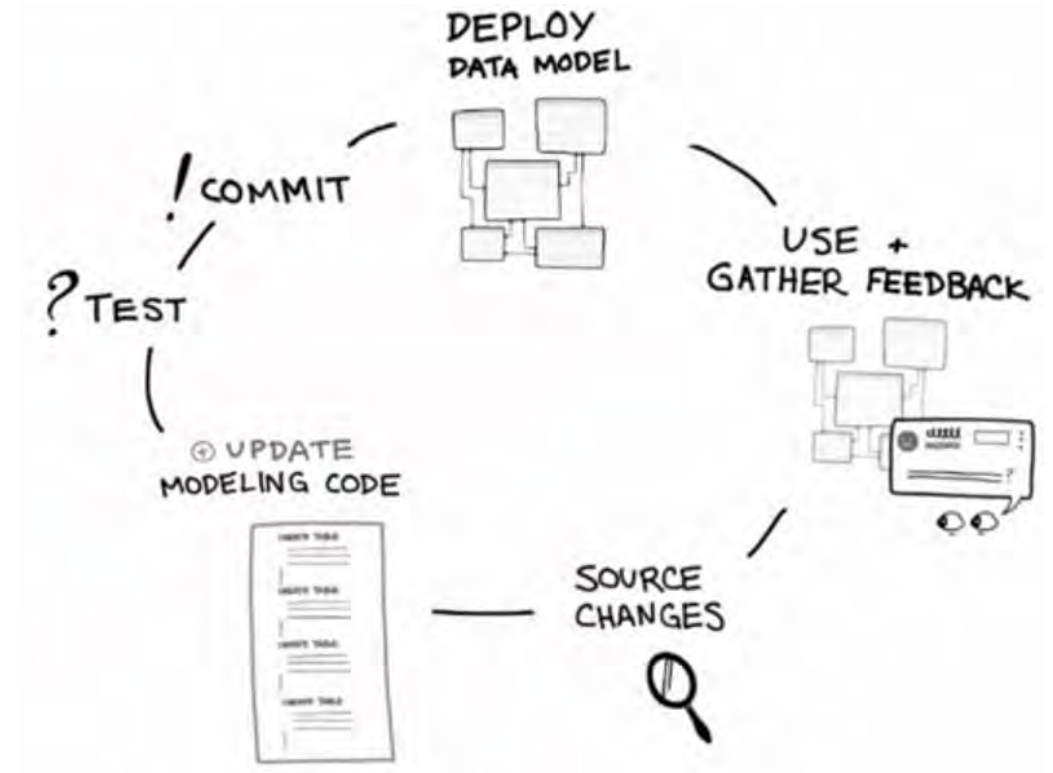
**“The only
thing
constant is
change”**

– Heraclitus of Ephesus



DataOps embraces change

- Business requirements change
- KPIs change
- Data changes
- Sources change
- Tools change
- People change
- Best-practices change





Putting it all together

Data Lakehouse & open-source formats are a response to system integration and cloud billing-models. **Cost-effective** thanks to cheap storage.

Wide table(s) are intuitive, performant, and even a requirement for data consumption tools. Its downside is outweighed by its benefits and **performance**.


Metadata is key to removing silos and tracking, monitoring and managing data as it moves toward the end products. **Understandability** from start to end!

DataOps is a paradigm shift. The way we treat data, metadata, and adapt to change itself.



How to get started?

1. Review the 18 principles of the DataOps manifest @ dataopsmanifesto.org
2. Read the best-practice and reasoning provided by the dbt developers @ docs.getdbt.com
3. Check out “Cloud Data Management — 4 stages for informed companies”, by Dave Fowler et.al.
4. Join the dbt Slack community @ getdbt.com/community



Bring your DataOps to the Data Lakehouse

Benefits & Challenges



#Conf42 #DevOps #2025

Jorge Loaiciga-Rodriguez
Senior Platform Engineer / Consultant @ Perelyn