



**Doing now what patients need next**





### Customer

**Roche Diagnostics**, part of multinational healthcare company Roche (F. Hoffmann-La Roche AG).

### Solution

**DataOps.live** enables a key capability for the self-service data and analytics infrastructure as part of a data mesh solution, providing orchestration and automation, integrating Snowflake and other tools in a **#TrueDataOps** approach.

### Requirement

Improved data management, collaboration and self-service analytics to empower data product teams and deliver on the company's stated purpose—*Doing now what patients need next*.

### Outcome

A data-driven business, with product teams adapting faster and delivering more while maintaining governance and security; Teams are onboarded rapidly; average MVP time reduced from six months to 6–8 weeks; 120 monthly releases compared to one release every three months before.

The DataOps.live platform is helping data this global pharmaceutical giant to orchestrate and benefit from next-generation analytics on a self-service data and analytics infrastructure consisting of Snowflake and other tools using a data mesh approach.

“I’m inspired every day by our purpose: Doing now what patients need next,” says Omar Khawaja, Global Head of Business Intelligence, Roche Diagnostics. “As a leader in pharmaceuticals and diagnostics, with ambitions to lead in clinical insights, data plays a key role in our business. We want to be a truly data-driven company.”

He says DataOps.live “provides a full end-to-end orchestration experience” to enable domain agility: data product teams across data domains can build and adapt data products faster, and more effectively, but without compromising on data governance and security and applying tried and tested software engineering best practices to data and analytics. “DataOps.live has moved to the heart of everything we do.”

### Background

Founded in 1896, Roche is the world’s largest biotech company, employing 90,000 people in 100 countries. Diagnostics is a core division, focused on delivering life-changing value. With Roche providing a number of solutions to help fight Covid-19, data has a central role to play in enabling the company to, for example, evaluate its products, supply chains and workforce in light of changing global demands and patient needs.

“It’s about speed,” says Paul Rankin. Head of Data Platforms. “The quicker you can get information to the people making the decisions, the better we can fulfil our promise of giving patients what they need next.”

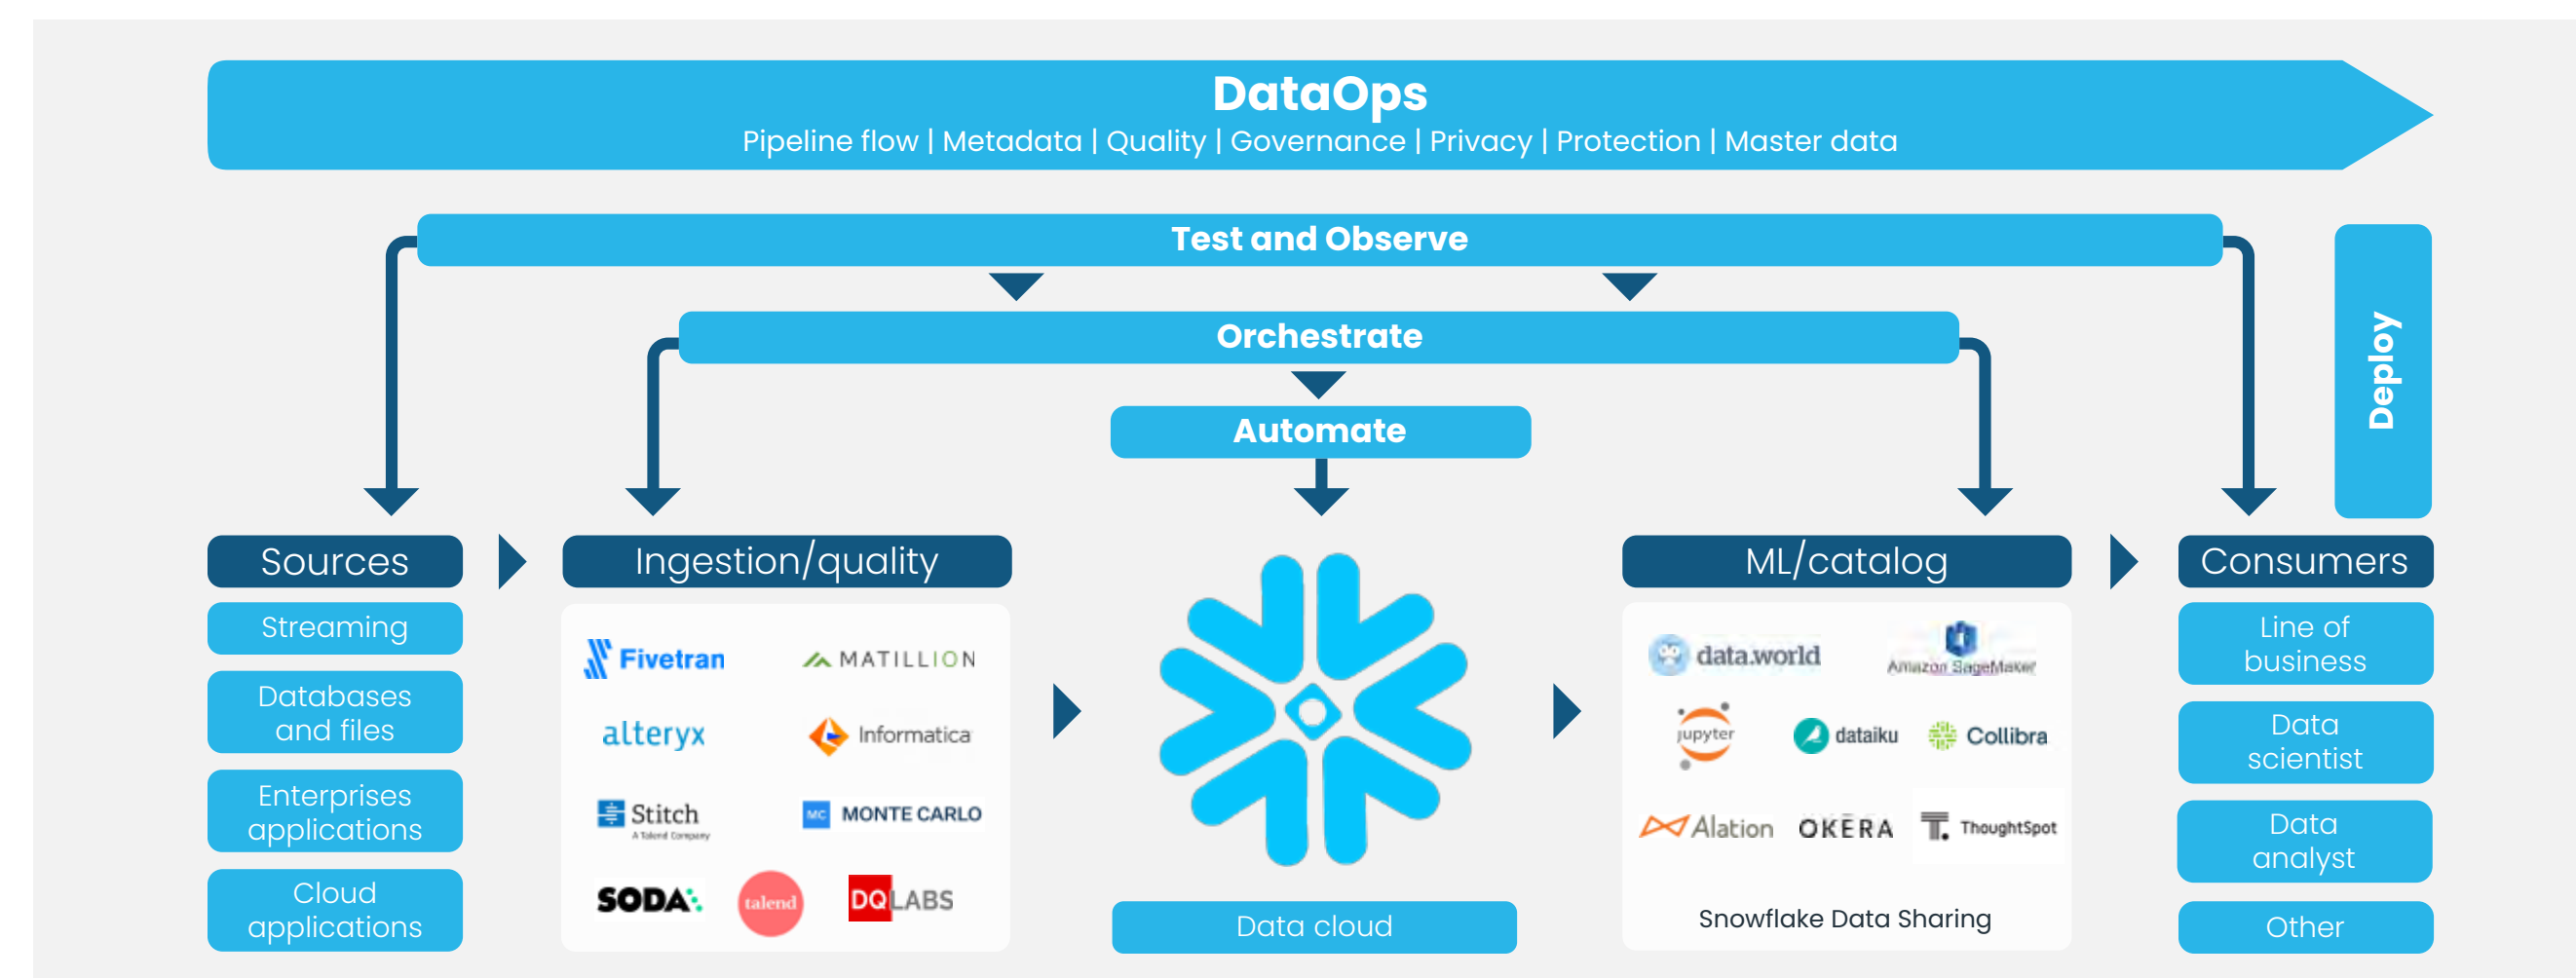
### Requirements

Omar Khawaja says, “You can’t carry on doing the same things and expect different results. We wanted to move the needle further on the dial and become a more agile data-driven business, which led to a pioneering data mesh and true DataOps approach.”

With huge volumes of data available internally and externally, Omar continues, “We ha’d undertaken initiatives to unlock the value of the data, like data lakes and data warehouses, but each had its limitations. With an ecosystem using Snowflake as our core data platform capability and tools such as Talend, Tableau, Alteryx and Collibra, there was no effective way to orchestrate activity, and automate the creation and lifecycle management of Snowflake environments.

“I chose to pursue a forward-looking data stack and, on the architecture side, a novel data mesh approach. This would address the critical issues around our people and our unique decentralized culture, to empower and give ownership to teams, creating and using our tech stack in a very different way while still having the federated governance on top that we need.”

### 1. DataOps: Automate | Orchestrate | Observe | Deploy








Roche Diagnostics has moved DataOps.live to the heart of everything we do—across data ingestion, data modeling, generating and visualizing insights, access management and security, data quality and observability, and data cataloging.

Omar Khawaja,  
Global Head of BI, Roche Diagnostics



**The solution:  
data mesh and DataOps.live**

The answer lay in decentralized domain-driven data products, where DataOps capabilities are implemented using DataOps.live, with some unique features enabled by Snowflake and other tools.

This secure platform for orchestration, lifecycle and release management supports the set-up and enforcement of policies meaning teams can be autonomous and innovative while also adhering to governance and security requirements in data access, residency, encryption and more.

“We wanted to implement data mesh and not a version of data mesh, which is why we worked with the originator of the concept Zhamak Dehghani and her team,” Omar adds. “We wanted ‘the gold standard.’”

DataOps.live was brought in to provide the unified orchestration and release management infrastructure.


“Implementing data mesh requires building numerous pipelines to extract, transform, prepare and share data within each domain and across many domains,” says Justin Mullen, co-founder of DataOps.live. “We provide the pipeline building and orchestration so all of this can happen.”

With embedded secrets management enablement of all data tools in extraction, modelling and transformation, the solution is integrated into core enterprise source code repositories, mirroring changes back and forth in real time.

Roche Diagnostics’ Data Management & Architecture teams act as an enabler and support function. They include specialists in data management, data operations and governance as well as technologies such as Snowflake, Talend and all other tools, with engineers from Roche, Accenture and DataOps.live working together.

This approach, recommended by DataOps.live, enables Roche to onboard new data products in a very short period of time.

**2. Building a DataOps Centre of Excellence**





**DataOps.live enables us to pull all this together in terms of orchestration, deployment, release management, CI/CD—and to do it at scale. It’s a complete game changer.**

**Paul Rankin**  
Head of Data Platforms, Roche Diagnostics

**Outcomes:  
faster creation and release  
of improved data products**

“We started from zero and began putting significant deliverables in the hands of users after only nine months,” says Omar Khawaja, “which speaks volumes about ‘speed to value.’” He says this is due to strong collaborations and partnerships inside and outside the business. Specifically, the DataOps.live platform has enabled:

- Robust data governance for compliance, with centralized policy definitions and enforcement.
- Greater business agility with teams working autonomously but within an overall framework: improved speed in meeting internal and external demands.
- Heightened data awareness and best practice sharing across teams.
- Time savings and improved efficiency through end-to-end automation and metadata management, for example more accurate and usable Data Catalogs.

“DataOps.live is exploiting every functionality Snowflake provides, bringing us the true DataOps practices we need,” Omar says. “It enables our teams to create the data products we require, using all governance best practices like code check-in and check-outs, and allows multiple data engineers to work concurrently without creating a bottleneck or interfering with each other.”

Paul Rankin says, “Federating this approach across hundreds of developers in 20+ data product teams was the challenge. We knew we’d need a powerful tool so developers could follow a DevOps methodology; the data is continuously changing and the pipeline is continuously questioned. This could never have been done with a centralized approach.

“DataOps.live enables us to pull all this together in terms of orchestration, deployment, release management and CI/CD—and to do it at scale. We’re talking about ROI in terms of saving thousands of hours and dollars in processing and developer time.”

With the platform build commencing in early 2021, by summer 2022 more than 1,300 users, developers and consumers had been onboarded, including 40+ data product teams, with two teams being onboarded each month. “We’re delighted with adoption. The DataOps.live team was amazing on the enablement side.

“The knowledge transfer was done, they helped to compile a blueprint and playbook, supporting us with the devolved architecture to align with the data mesh concept. I’ve been impressed by their ability to be agile and adapt to our way of deployment.”

Critically, average MVP time has reduced from six months to only 6–8 weeks. The number of releases has risen to 120+ per month, compared to just one release every three months prior to DataOps.live and data mesh. The new platform has enabled the integration of more than 15 additional capabilities and partners into the Roche Diagnostics data ecosystem.

Omar Khawaja adds, “We will continue to simplify processes through automation as we encounter new requirements of different data types, new steps to follow, new features of other tools, and how DataOps.live can continue orchestrating and tying all that together.”

