

Where Everyday AI Meets the Data Cloud

Business and technical users can quickly create and deploy analytics and AI solutions powered by Snowflake, all while maintaining enterprise visibility and trust. With Dataiku, all users have access to data directly in Snowflake — without additional configuration.

Anyone (not just engineers) can prep data, train machine learning models, build large language model (LLM)-powered apps, and deploy workflows into production with Dataiku's visual, SQL, and Python components. Plus, an easy-to-understand visual flow helps teams build readable, reusable, and production-ready pipelines.



Build Production Assets Quickly

With Dataiku, organizations can build and deploy projects quicker than ever before. For example, by centralizing data in Snowflake and then building a parameterized post-marketing campaign analysis workflow in Dataiku, Air Canada data scientists were able to accelerate from one analysis in two weeks to twelve analyses in 3.5 hours.

96% Resource utilization improvement at Air Canada¹

Empower Teams

Everyone in the business can inject analytics and AI into their day-to-day with Dataiku and Snowflake. Dataiku's visual tools, backed by Snowflake virtual warehouses and AI services, helped Thrive Skilled Pediatric Care (Thrive SPC) move from a web of spreadsheets to a central, automated reporting system without having to hire a new team of data scientists.

3% margin improvement at Thrive SPC²

Reduce Risk

Build trust as you expand access to data, analytics, and AI in a secure, collaborative environment. At Novartis, implementing Dataiku brought transparency that has bolstered trust in data and the insights derived from it, reducing the risk of erroneous decisions based on inaccurate or misleading information.

10+ teams easily & securely share data at Novartis³

What Dataiku + Snowflake Clients Are Saying

"[Dataiku] was for us, the best overall ROI because of the skills that were required to use it. We didn't need a DevOps engineer, data scientist, and similar skillsets."

— Mike Hoyt, SVP of Technology and Business Process at Thrive SPC



"Dataiku enabled the orchestration of the processes, allowing us to automate, create the flow, and schedule our processes as opposed to having to manually re-code and treat these projects as ad-hoc analysis, which has reduced analysts' efforts to a few hours one day a week. The seamless integration with Snowflake (to easily access the data) and PowerBI (to easily showcase the outputs) also helped streamline the whole process."



Enterprise-Ready AI at Your Fingertips

Foster Collaboration

- Business and analytics teams can take on advanced analytics projects together in a single collaborative, visual environment.
- Build trust in AI-driven systems by bringing business users to the table.
- Reduce friction and handoffs with IT teams by building, deploying, and monitoring projects all in one platform.

Cover All Stages of the Analytic Lifecycle

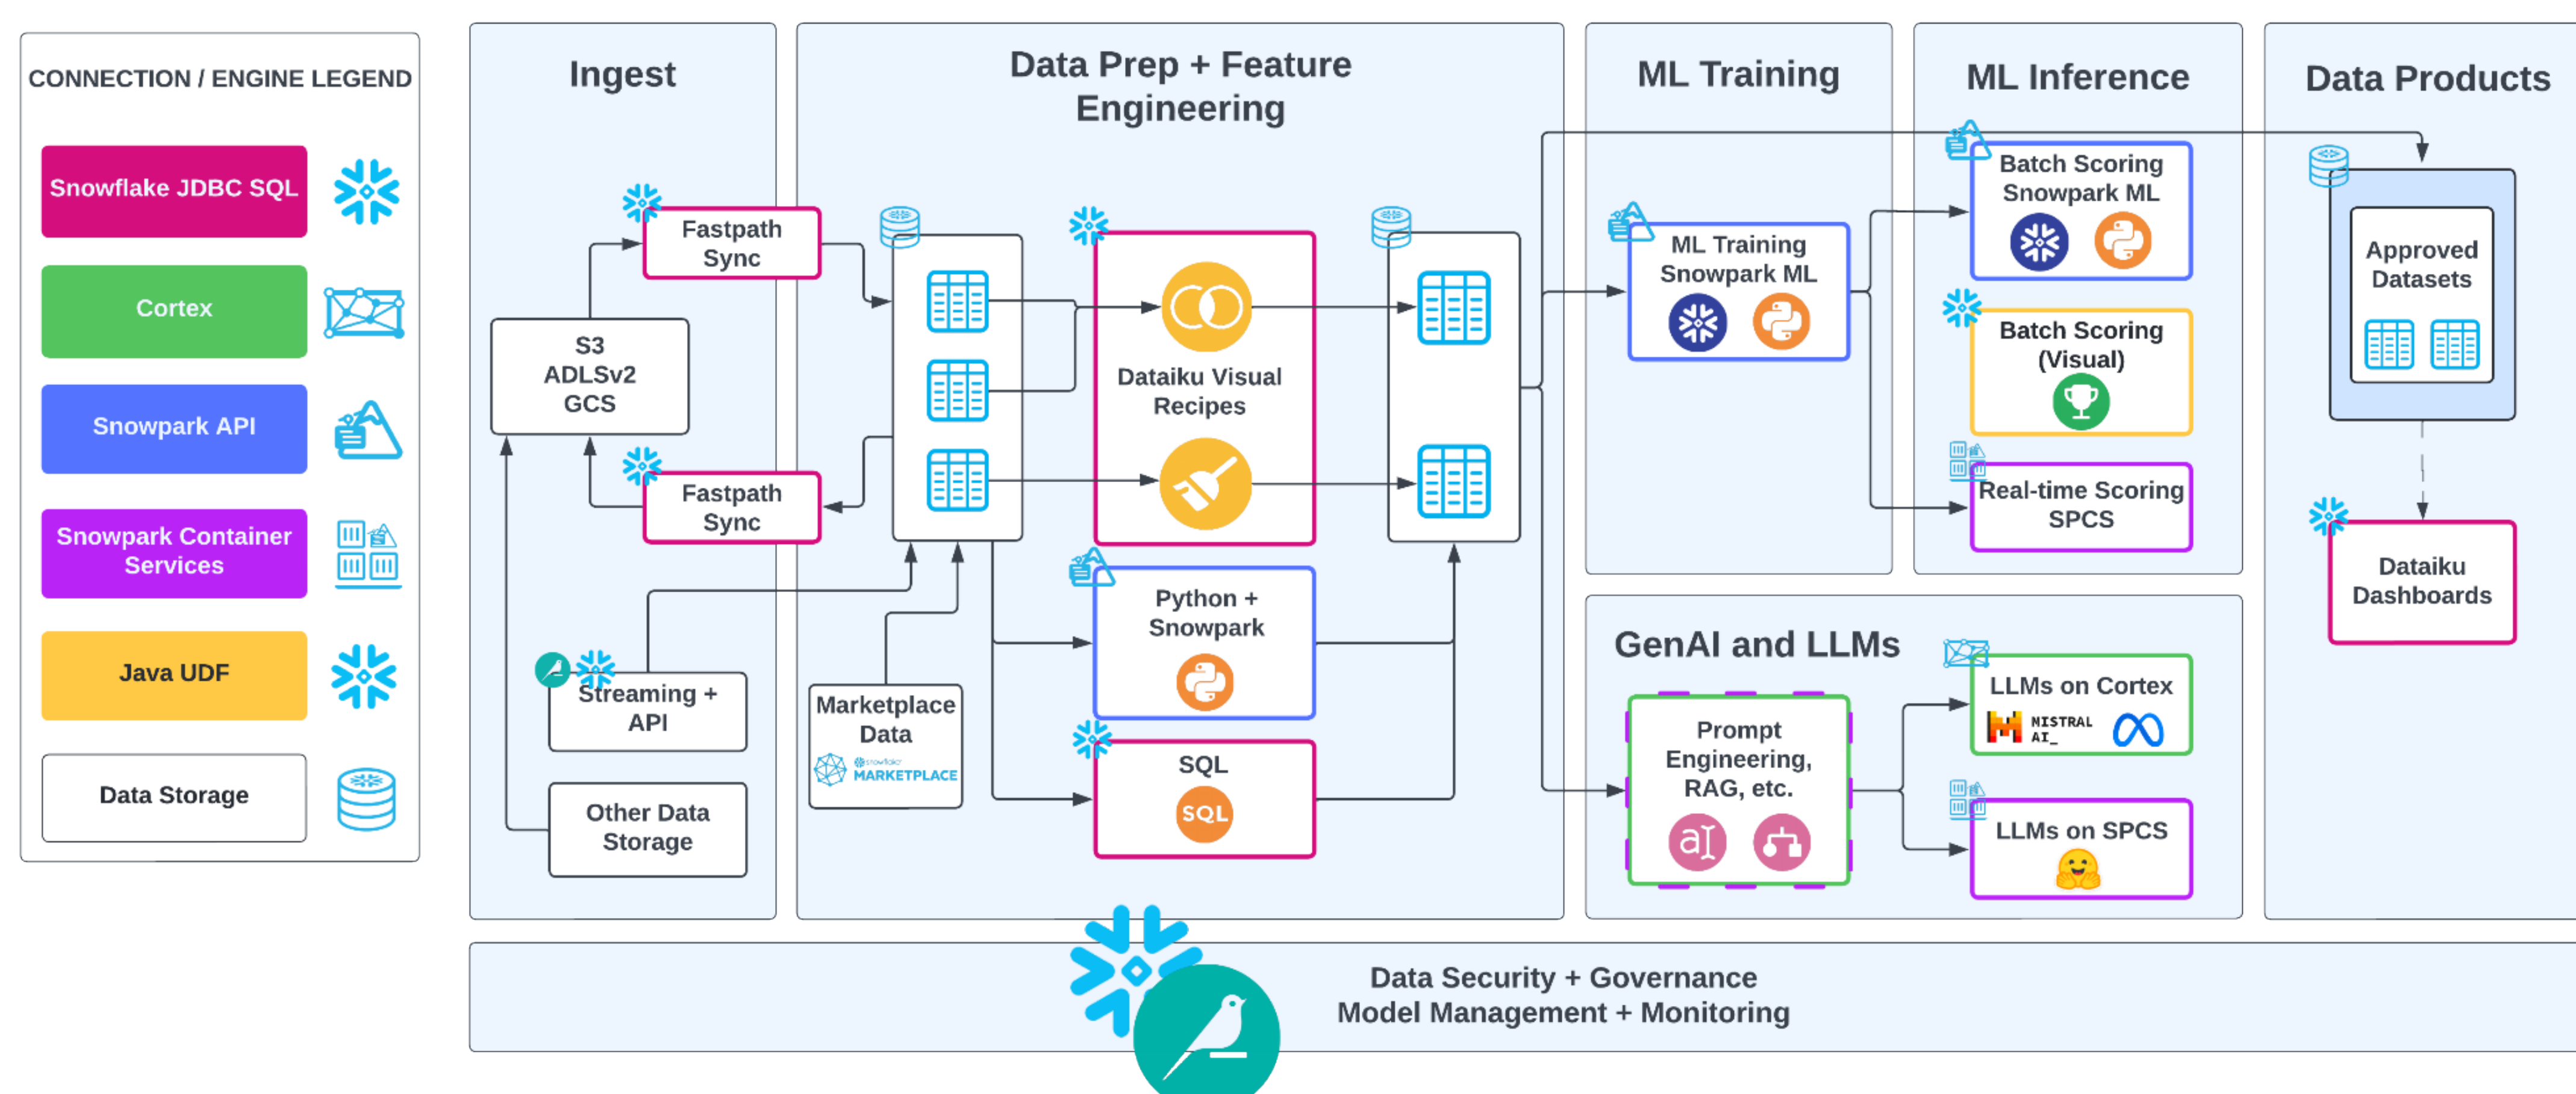
- Ingest and clean data with Dataiku's visual recipes, pushing SQL compute to Snowflake's powerful virtual warehouses.
- Train traditional machine learning (ML) models in parallel with Snowpark ML, analyze performance and bias with helpful charts, and track experiments.
- Automate and deploy entire pipelines for batch workloads and API services on Snowpark Container Services for real-time use cases with just a few clicks.

Run Jobs Faster With Push-Down Computation

- Develop sophisticated analytic pipelines in Dataiku using visual and SQL recipes, pushing compute to Snowflake with zero data egress.
- Domain experts can leverage the computational power of Snowflake without having to write SQL.
- Engineers can write Snowpark Python code in notebooks and code recipes to query and process data, engineer features, and build more advanced applications.

Scale the Use of Generative AI

- The Dataiku LLM Mesh allows IT to fully control access to Generative AI tools, helping business teams build safe, secure, enterprise-ready applications.
- Connect to top performing, generalized LLMs in Snowflake Cortex or niche, fine-tuned LLMs in Snowpark Container Services.
- Engineer prompts, build retrieval augmented generation (RAG) pipelines, and more with Dataiku's visual and Python-based LLM tools.



1. <https://community.dataiku.com/t5/Dataiku-Frontrunner-Awards/Air-Canada-Post-Campaign-Analysis-Automation/ta-p/36592>
 2. <https://www.dataiku.com/stories/detail/thrive-spc/>
 3. <https://community.dataiku.com/t5/Dataiku-Frontrunner-Awards/Novartis-Automated-Commercial-Analytics-For-Brand-Uptake/ta-p/36756>