SOLUTION BRIEF

Domino Data Lab | Enterprise MLOps Platform

End-to-end data science lifecycle: Build, deploy, & monitor models in Snowflake

From data-driven to model-driven

Domino enables organizations to become model-driven and succeed with data science at scale. In today's world, failing to become model-driven puts organizations at risk of being disrupted or left behind by competitors.

Becoming model-driven requires data science maturity as a discipline—across people, processes, and technology. This starts with a team of sophisticated code-first data scientists who follow best practices with open and flexible technology that meets their needs and IT requirements. The full lifecycle from development to continuous improvement of production models must be seamless, secure, governed, and integrated.

More datasets, more production models, zero DevOps

Combine the flexibility of model building in Domino with the scalability and power of Snowflake's platform for in-database computation. Run models where data resides.

- ➔ Instant, seamless access to Snowflake data
- ➔ Use Domino to train models in Snowflake using Snowpark
- ➔ Deploy models directly from Domino to Snowflake for in-database scoring
- ➔ Monitor all production models for prediction accuracy

Get more models into production, iterate faster to compound knowledge, and increase business impact with the power of Domino and Snowflake.

Traditional challenges to data science scale

Inflexible Infrastructure

Data scientists require powerful compute, valuable and sensitive data, and the latest open-source tools

- **Result:** Data scientists do DevOps with different tools and bespoke hardware
- **Impact:** Less time for data science, shadow IT with operational and security risks, and slower innovation

Wasted Work

Data scientists work independently with different processes and tools, and struggle to find and build off prior work

- **Result:** Low visibility into past work and in-flight projects
- **Impact:** Redundant and slow work, poor governance, and low auditability and reproducibility

Productionization Pitfalls

Data scientists struggle to consistently track experiments, put them into production, and monitor model quality.

- **Result:** Low standardization for model deployment, monitoring, and retraining
- **Impact:** IT complexity, project delays, and bad business decisions from inaccurate models

The world's most successful enterprises rely on Domino

evidation · SCOR The Art & Science of Risk · gsk · BNP PARIBAS · Moody's Analytics · Red Hat · Arthur D. Little · ConocoPhillips

"Domino is an underpinning to our scaling from 10 to 150 projects."

NAJAT KHAN, Chief Data Science Officer, Global Head, Strategy & Operations for R&D

Janssen | Pharmaceutical Companies of Johnson & Johnson

"Domino automates decision making with models 800X faster than humans."

STIG PEDERSEN, Head of Machine Learning

Topdanmark

"The Domino platform brings order to the chaos."

MATT SEAMAN, Enterprise Operations Chief Data & Analytics Officer

Lockheed Martin

Dominodatalab.com

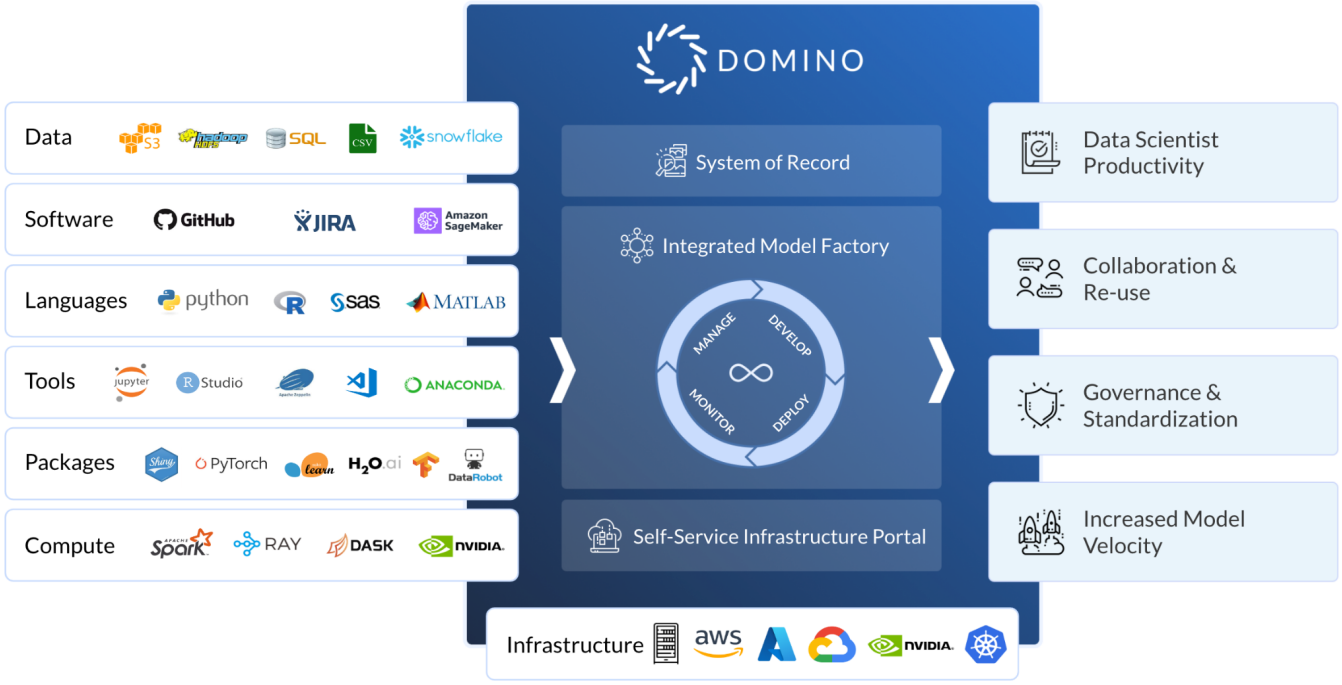

Domino & Snowflake: Modern ML stack

Simplify enterprise infrastructure with a common platform across IT and Data Science. Move more models to production, iterate faster, and collaborate more effectively to embed model-driven applications throughout the enterprise.

Accelerate Model Development

In a few clicks, data science teams have self-service data science workspaces with governed, secure access to data in Snowflake, pre-configured with curated tools, packages, frameworks, and compute for model development and training at scale - no DevOps required.

Flexible Model Deployment

Improve prediction response time for critical applications by deploying models and executing Python scoring code inside Snowflake Data Cloud, using the scalability and power of Snowflake for in-database computation. Simplify enterprise infrastructure with a common data and deployment platform.

Real-Time Model Monitoring

Simplify model management with automated prediction data capture pipelines and monitoring for models deployed to Snowflake Data Cloud. Ensure model accuracy with continuously updated data drift and model quality calculations to make better business decisions.

"Domino and Snowflake enable our data science team to connect to person-generated health data (PGHD) seamlessly and prototype rapidly, all within a secure environment. They are integral in our work to measure health in everyday life, and ultimately turn data into strategies to promote better health outcomes."

- Luca Foschini, Chief Data Scientist at Evidation

Loved by data scientists, trusted by IT

The Domino platform accelerates productivity and collaboration for teams of code-first data scientists with the most open and flexible access to tools, languages, and infrastructure. Domino also tames the "wild west" of data science by enabling IT with centralized resource management, security, and governance.

Built for Teams

- Centralized system of record for data science projects, experiments, and artifacts
- Seamless collaboration across teams while using different tools
- Full visibility to work, results, and insights of past and current projects

Open & Flexible:

- Broad support for the latest open-source and commercial tools, languages, and packages
- Pre-configured environments, including distributed compute and GPUs in a few clicks
- Single platform to consolidate tools, improve governance, and reduce support costs and burden

Integrated Workflows:

- Support for full MLOps lifecycle from data exploration to model monitoring and retraining on a single platform
- Robust security and governance with repeatability and reproducibility across the lifecycle
- Enforcement of data science best practices

Learn More

dominodatalab.com/snowflake
blog.dominodatalab.com
snowflake@dominodatalab.com

DOMINO snowflake