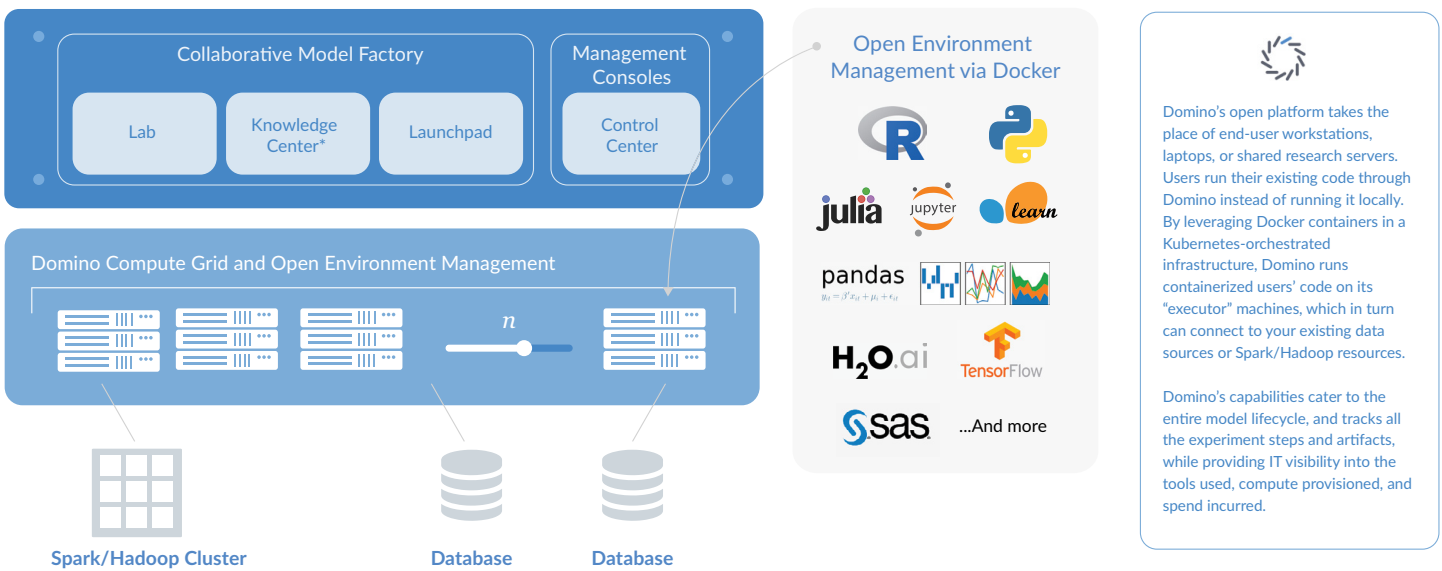DOMINO

Deliver and Govern High-Impact Models with Confidence

Domino provides a centralized platform for data scientists to build, validate, deliver, and monitor models on a scalable infrastructure managed by Domino Compute Grid. Domino Lab is a single place for data scientists to provision any compute environments for their experiments in one click, regardless of the tools used. Domino Launchpad enables IT and data science teams to publish models in a safe, scalable manner, while Domino Knowledge Center automates best practices and streamlines model tracking, versioning, and sharing. These tools empower data scientists to be more agile and productive, while Domino Control Center enforces ModelOps best practices to ensure consistency and transparency, contain costs, and reduce security risks.

Domino Data Science Platform

Collaborative Model Factory

Lab | Knowledge Center* | Launchpad

Management Consoles

Control Center

Domino Compute Grid and Open Environment Management

n

Spark/Hadoop Cluster | Database | Database

Open Environment Management via Docker

R | python | julia | jupyter | learn | pandas | $y_{it} = \beta' x_{it} + \mu_i + \epsilon_{it}$ | H2O.ai | TensorFlow | SAS | ...And more

Domino's open platform takes the place of end-user workstations, laptops, or shared research servers. Users run their existing code through Domino instead of running it locally. By leveraging Docker containers in a Kubernetes-orchestrated infrastructure, Domino runs containerized users' code on its "executor" machines, which in turn can connect to your existing data sources or Spark/Hadoop resources.

Domino's capabilities cater to the entire model lifecycle, and tracks all the experiment steps and artifacts, while providing IT visibility into the tools used, compute provisioned, and spend incurred.

* Powered by patent-pending Reproducibility Engine

Any organization that is investing in data science must empower its researchers to experiment rapidly and operationalize models that adhere to IT policies. The key pain points Domino addresses for the IT leaders include:

- **Friction in ModelOps** IT and data science teams need to collaborate to get models to production rapidly without long rewrite cycles. Furthermore, rapid iterations on models in production should be built into the deployment technology and process.

- **Shadow IT** Hiring PhDs and arming them with laptops and shared servers won't drive sustained competitive advantage. IT needs to provide centralized access to right tools and elastic compute to data scientists within their governance and purview.

- **Security, governance and cost controls** IT needs to ensure data science practices are carried out in a secure infrastructure with auditability and compliance. Costs of elastic compute environments need to be accounted for properly.

TRUSTED BY

A system of record for data science

With Domino, data science teams can acquire the tools they need to develop, validate, deliver and monitor models with speed, reproducibility and governance. Features include:

Open platform to speed up data science work Domino's open platform accelerates innovation by enabling data scientists to take advantage of the latest tools and packages. Domino ships with optimized distributions of tools such as R, Python, Jupyter, RStudio, SAS, Tensorflow, and H2O. IT can also proactively make the best-in-class tools available on customized Domino environments or allow data scientists to build their own in compliance with IT policies.


Knowledge Center for reproducibility and knowledge management Domino keeps code, data, results, discussions, and environments linked in one place, so anyone can share insights and find artifacts from past projects and apply them to future work. A centralized hub increases collaboration, transparency, and reproducibility, all of which help to speed research as they mitigate compliance and audit risk.

Open Infrastructure Foundation for faster cloud and on-premise deployment Domino's containerized architecture allows data science workloads to be deployed easily anywhere. Data scientists can provision compute resources with a single click, either in the cloud or across on-premises hardware. Data science workloads are ideal for the cloud because they are computationally intensive, bursty, and often require specialized hardware like GPUs. Domino enables turnkey data science workflows tightly integrated with AWS.

Launchpad to remove ModelOps friction Domino lets data scientists publish models as REST APIs, hosted interactive web apps (e.g., Shiny), or scheduled jobs for generating reports or running ETL tasks. This feature allows users to re-implement their work on production systems faster, so the business gains value more quickly. Access controls and gatekeeping features let you enforce governance processes to control deployment.
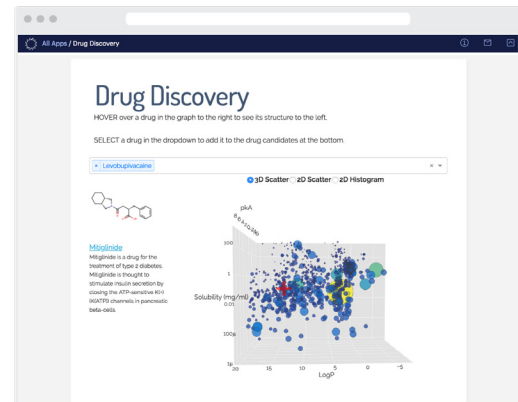
Control Center to strengthen security and governance... Analytical assets are kept in one central place—on managed and monitored infrastructure—instead of being spread across users' machines. Domino largely reduces vulnerabilities introduced by shadow IT. Granular access controls keep work secured, with activity logs and reports available to administrators for audit and compliance.
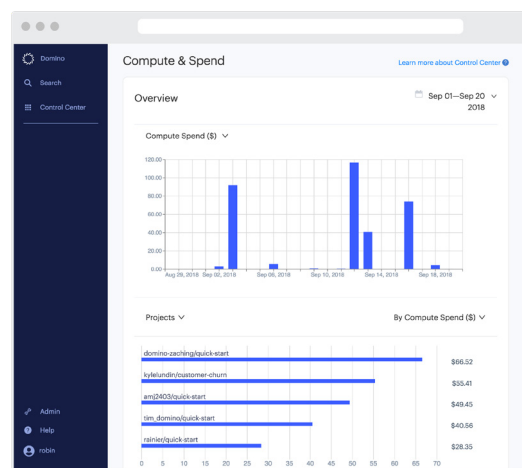
... And provide insights into costs With Domino, administrators have full insight into the compute costs of data science work. Compute and spend is accounted for and attributed to users and projects so CIOs and business line leaders can perform more accurate budgeting and planning.

---

dominodatalab.com

DOMINO