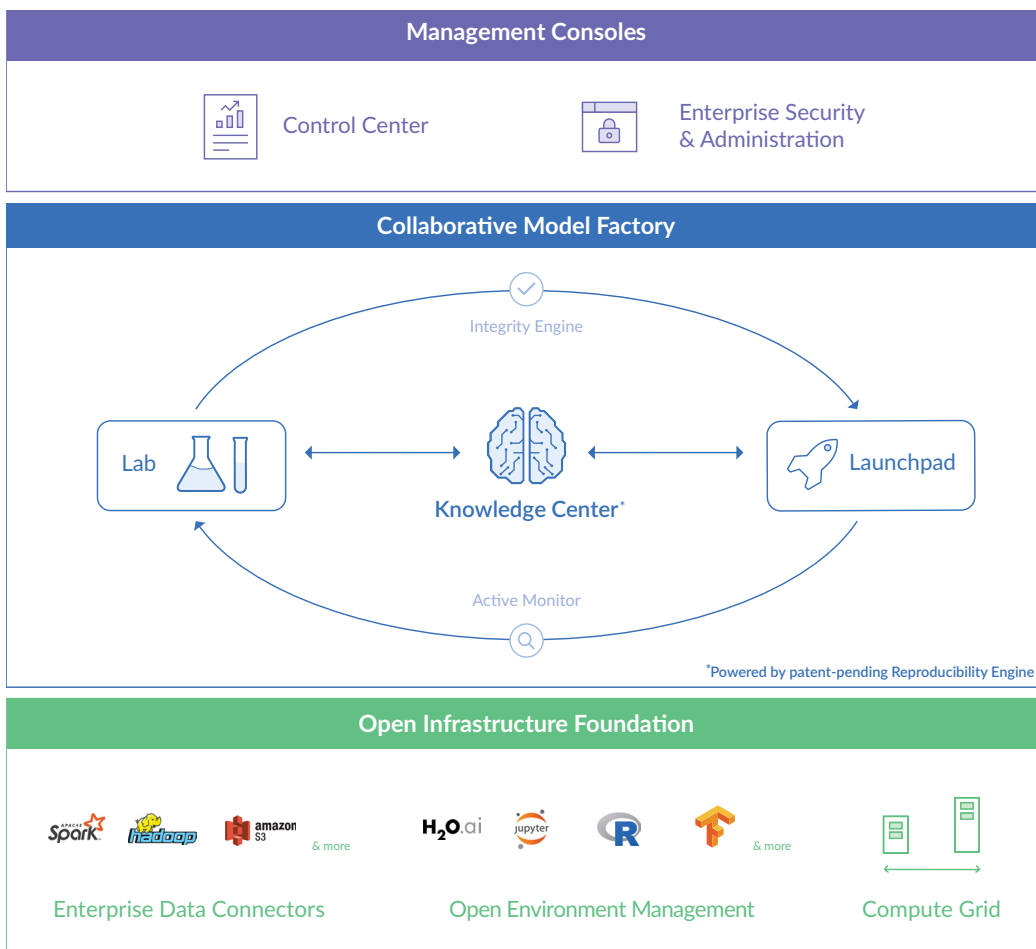


Drive Breakthrough Research, Deliver High-Impact Models

The Domino data science platform empowers data scientists to develop and deliver models with open access to the tools they love. [Domino Lab](#) supports both interactive and batch experimentation with all popular IDEs and notebooks (Jupyter, RStudio, SAS, Zeppelin, etc.).

With [Domino Launchpad](#) data scientists can easily deliver and understand the impact of their model products. [Domino Open Infrastructure Foundation](#) orchestrates all the relevant hardware and software, while the [Reproducibility Engine](#) automatically captures and organizes work to avoid reinventing the wheel.



Domino addresses the three biggest problems data scientists face

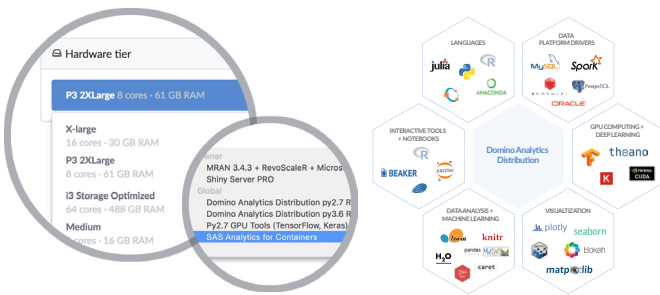
- DevOps headaches** Data scientists are forced to use outdated, underpowered tools and hardware or waste time juggling infrastructure themselves.
- ModelOps bottlenecks** Promising models have little impact on the business, frustrating data scientists and stakeholders.
- Redundant work** Data science teams are not enabled with knowledge management tools or processes, leading to duplicative efforts and missed opportunities for innovation.

TRUSTED BY



Domino Features and Benefits for Data Scientists

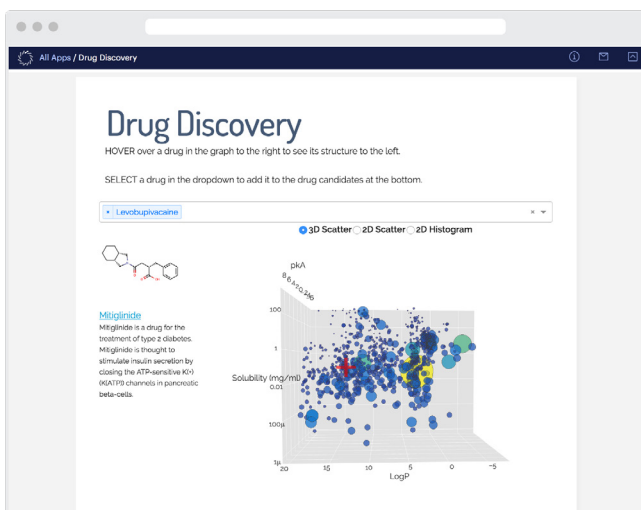
Eliminate DevOps barriers to model development



Compute infrastructure automation Select whichever hardware is appropriate for the task at hand, whether quick EDA on sampled data or many GPUs for intensive deep learning training workloads. Vertically and horizontally scale with [Compute Grid](#).

Open tooling platform Domino comes pre-configured with popular languages, tools, and packages such as R, Python, SAS, Jupyter, RStudio, Tensorflow, and H2O. You can run code in a native development environment such as RStudio, SAS Studio, or JupyterLab. Build your own custom Docker environment with any packages or drivers. [Environments are versioned](#) and can be shared among your colleagues while still safely sandboxed from others.

Launchpad reduces ModelOps bottlenecks

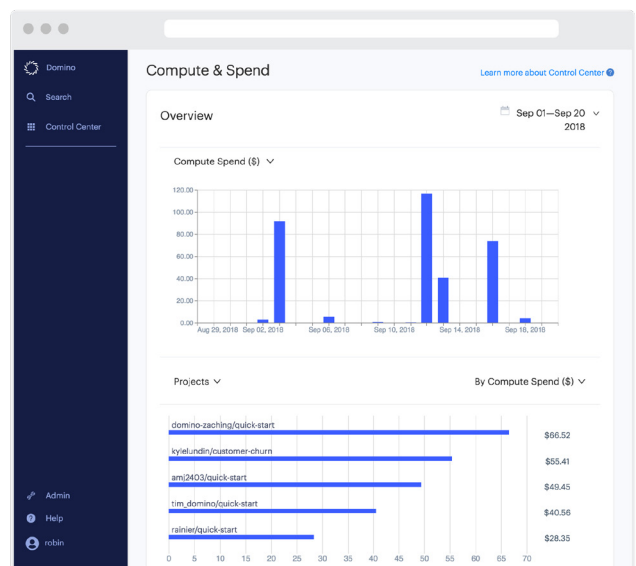


Multiple delivery modes Data scientists can run scheduled reports, use self-service web forms (Launchers), develop interactive apps built with Shiny or Flask, or call batch and real-time APIs.

Understand engagement and impact Teams can see how their model products and APIs are being consumed in the UI or via API.

Enterprise-grade delivery Delivered models are automatically versioned, secured via separate permissioning, and highly available.

Accelerate iteration velocity



Model product usage tracking Domino automatically captures usage data for model products, providing transparency into business impact and encouraging collaboration between stakeholders and data science teams.

Automatic versioning of model products Each version of a report, app, or API is versioned, minimizing risk to delivery and allowing easy roll-back.

Model lineage capture Domino's [Reproducibility Engine](#) automatically captures and organizes the dependencies and metadata in the model development process (data, code, environment, parameters, etc.), ensuring future collaborators can discover and understand how a model was developed and iterate on it if necessary.