



ENTERPRISE MLOPS ACCELERATES AI MODELS IN ENERGY

Domino Data Lab Unleashes Data Science at Scale

This new phase of our collaboration with NVIDIA further enables enterprises to solve the world's most challenging problems by putting models at the heart of their businesses. Together, we are providing every company with the end-to-end platform to rapidly and cost-effectively deploy models enterprise-wide."

— **Thomas Robinson**, Vice President of Strategic Partnerships, Domino Data Lab

The energy industry presents unique challenges for developing and deploying effective machine learning models. From fossil fuel to renewable energy networks, heterogeneous edge infrastructure, complex enterprise architectures, and organizational complexity mean it's difficult to drive a seamless, low-friction machine learning process.

By harnessing NVIDIA's accelerated computing platforms, Domino Data Lab is helping energy companies eliminate roadblocks to deploying cutting-edge AI and machine learning to manage energy infrastructure.

Domino Democratizes GPU Access

Founded in 2013, Domino Data Lab powers model-driven businesses with its leading Enterprise MLOps Platform.

By democratizing access to NVIDIA's powerful NVIDIA DGX™ systems and AI frameworks, Domino is empowering teams to accelerate the development and deployment of data science workloads while increasing collaboration and governance.

Data science teams can deploy the Domino Kubernetes-native platform on-premises, in the cloud, or in the modern hybrid cloud to support optimized compute governance and greater model development scalability.

Governance

With Domino, data scientists have self-serve access to custom GPU-enabled resources from the Domino workbench. The platform provides a high standard of governance and access control to satisfy the strict requirements of enterprise IT organizations.

Permissions can be set to ensure employees without proper entitlements are not overconsuming valuable resources, while power users have full access to maximize the use of enterprise hardware.

Domino's support for NVIDIA Multi-Instance GPU (MIG) technology on the NVIDIA A100 Tensor Core GPU also enables admins to allow up to 56 concurrent notebooks or hosted models, each with an independent GPU instance.



Domino's Enterprise MLOps Platform accelerates research, speeds model deployment, and increases collaboration for code-first data science teams at scale.

PRODUCT BENEFITS

- > Self-serve GPU access
- > Automated DevOps activities
- > Shared platform with access to tools and infrastructure
- > Standardized workflows to track, monitor, and manage models
- > Data scientists free to focus on research

RESULTS

- > Optimized GPU utilization
- > Reduced model training time
- > Faster time to publish and deploy work

NVIDIA HARDWARE USED

- > NVIDIA DGX systems
- > NVIDIA-Certified servers

NVIDIA SOFTWARE USED

- > NVIDIA AI Enterprise
- > NVIDIA Fleet Command






Figure 1: Domino users can launch on-demand workspaces with the latest NVIDIA GPUs, optimized with open-source and commercial data science tools.

Scalability

With a single system of record, IT teams can eliminate infrastructure configuration and debugging tasks. NVIDIA DGX resources can be configured within the Domino compute grid, rather than depending on IT for one-off tools and deployments, reducing the time data scientists need to spend on DevOps.

In addition, Domino enables the automatic creation, scaling, and management of multi-node clusters, releasing them when training is done—eliminating the need for dedicated resources and maximizing GPU utilization.

Machine Learning Ops for More Efficient Maintenance

Effectively managing and monitoring renewable energy systems requires an extensive network of IoT sensors. One solar installation, for example, could have 600,000 data points to monitor.

With many states laying out aggressive timelines for building a sustainable energy future, the demand for machine learning in energy infrastructure has accelerated.

Electric utility companies can use the Domino platform as an orchestration tool to increase efficiency across the entire data science lifecycle—and focus on modeling efforts rather than periphery technology.

By launching its data science operations on Domino, AES, a Virginia-based energy company, was able to develop, deploy, and manage 50 machine learning models within a condensed timeline.

“We’re leveraging digital technology to expedite the shift to renewables and the deployment of new energy solutions,” said Sean Otto, Director of Analytics at the AES Corporation. “Our goal is to be the best at delivering viable business AI and machine learning models to support the future of energy, and Domino’s Enterprise MLOps Platform is helping us achieve that goal.”

With access to the compute infrastructure they needed, AES data scientists reduced model training time significantly, in some cases from 12 hours to only three minutes. This allowed them to launch 50 models in under two years.

A Snow Detection Model on Edge Infrastructure Maximizes Equipment Uptime

AI models help the AES team more efficiently manage a distributed network of solar panels and wind farms, as well as business and logistics operations.

One of the models AES launched on Domino was a snow detection model. Without technicians onsite, teams need to be dispatched into the field to service equipment and make repairs. If the technician arrives and the solar equipment is covered in snow, they may not be able to make the repair. This is a significant sunk cost for the operator.

Using the Domino platform, AES pushes snow detection models to NVIDIA-Certified Systems™ at the edge via **NVIDIA Fleet Command™**, a managed platform for container orchestration at the edge. With edge AI snow detection, operators receive alerts when there is snow present and gain insights into the extent of coverage and the melt rate. This enables them to make more informed decisions before they send technicians into the field to make repairs.

Powered by NVIDIA

Domino's enterprise MLOps platform is an NVIDIA DGX-ready software solution, tested and certified for use on DGX systems to deliver revolutionary performance.

Centralized support with version control for data science workspaces ensures stable and consistent access to cutting-edge deep learning compute resources and frameworks such as Keras, TensorFlow, Torch, and NVIDIA® TensorRT™.

When data science teams deploy their on-demand notebooks, they select tailored DGX resources and software for their tasks with administrator-controlled permissions. The Domino open platform streamlines workflows and scales, taking full advantage of the power of NVIDIA DGX systems and NVIDIA NGC™ optimized containers out of the box.

Using NVIDIA AI Enterprise, a comprehensive suite of optimized AI tools and frameworks, Domino empowers its customers to cost-effectively scale data science work by accelerating research, model development, and model deployment on mainstream accelerated servers.

Finally, the Domino platform uses Fleet Command to push trained models back out to edge devices so that equipment can operate at maximum uptime with minimum upkeep costs.




Figure 2. Domino accelerates data science initiatives by pairing NVIDIA-Certified Systems with its best-in-class Enterprise MLOps Platform.

Experience Domino on NVIDIA LaunchPad

Experience the difference of Domino accelerated by NVIDIA AI Enterprise on NVIDIA LaunchPad. NVIDIA LaunchPad provides free, immediate, short-term access to the necessary hardware, software, and data, so customers can experience end-to-end solution workflows in the areas of AI, data science, 3D design collaboration and simulation, and more.

The **Domino on NVIDIA LaunchPad** hands-on lab allows customers to use the Domino Enterprise MLOps Platform to develop, deploy, and monitor models using NVIDIA RAPIDS™. Domino on NVIDIA LaunchPad includes:

- > Temporary **dedicated** Domino deployment for up to five users
- > **Ready-to-use infrastructure** available in 24 hours
- > **A hands-on tour** of Domino's Enterprise MLOps Platform
- > Storage for **up to 1TB** of data
- > Ability to test workloads on **NVIDIA-Certified hardware**




Image courtesy of Domino

Figure 3. Domino provides data scientists with self-serve access to their preferred integrated development environments (IDEs), languages, and packages, so they can focus on data science innovation.

By partnering with NVIDIA, Domino Data Lab has created a single platform to support the end-to-end data science model lifecycle, empowering teams to bring machine learning into production with greater velocity and more efficient use of resources across industries.

Ready to Get Started?

To learn more about Domino on NVIDIA LaunchPad, visit:
www.nvidia.com/domino-launchpad