

Domino Data Lab Enterprise MLOps Platform

Data Center-Ready | Validated for NVIDIA AI Enterprise and NVIDIA -Certified Systems™

Domino makes data science at scale a reality

For companies with large teams of code-first data scientists, Domino is an Enterprise MLOps platform that accelerates the process of developing and productionizing data science work. Domino, in close collaboration with NVIDIA®, supports open, collaborative, reproducible model development, training, and management free of DevOps constraints - powered by efficient, end-to-end compute.

Data Center-Ready Enterprise MLOps

Develop, deploy, and manage GPU-accelerated data science workloads on existing enterprise infrastructure. Domino's validation for [NVIDIA AI Enterprise](#) pairs the Enterprise MLOps benefits of workload orchestration, self-serve infrastructure, and collaboration with cost-effective scale from virtualization on mainstream accelerated servers.

Under the hood, Domino automates the DevOps activities required to optimize utilization of powerful NVIDIA accelerated computing, eliminating the low-value configuration and debugging tasks performed by valuable researchers. Domino provides flexible, governable access to critical GPU resources, and it blends these workloads seamlessly with traditional infrastructure across a single system of record.

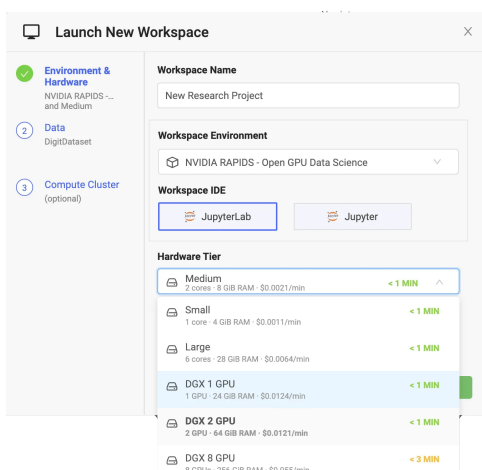


Figure 1: Choose the right hardware and software for any task. Use Domino Workspaces to run proprietary and open source tools side by side, governing scalable GPU Hardware Tier access with Domino's Compute Grid. Here, a JupyterLab IDE is configured to launch with access to NVIDIA Compute Resources

Democratize GPU access, unleash your AI infrastructure

Data Scientists and AI Researchers can focus on research instead of dev ops

Launch Domino Workspaces on-demand, with docker images configured with the latest data science tools, frameworks, and NVIDIA GPUS - with automatic storing and versioning of code, data, and results.

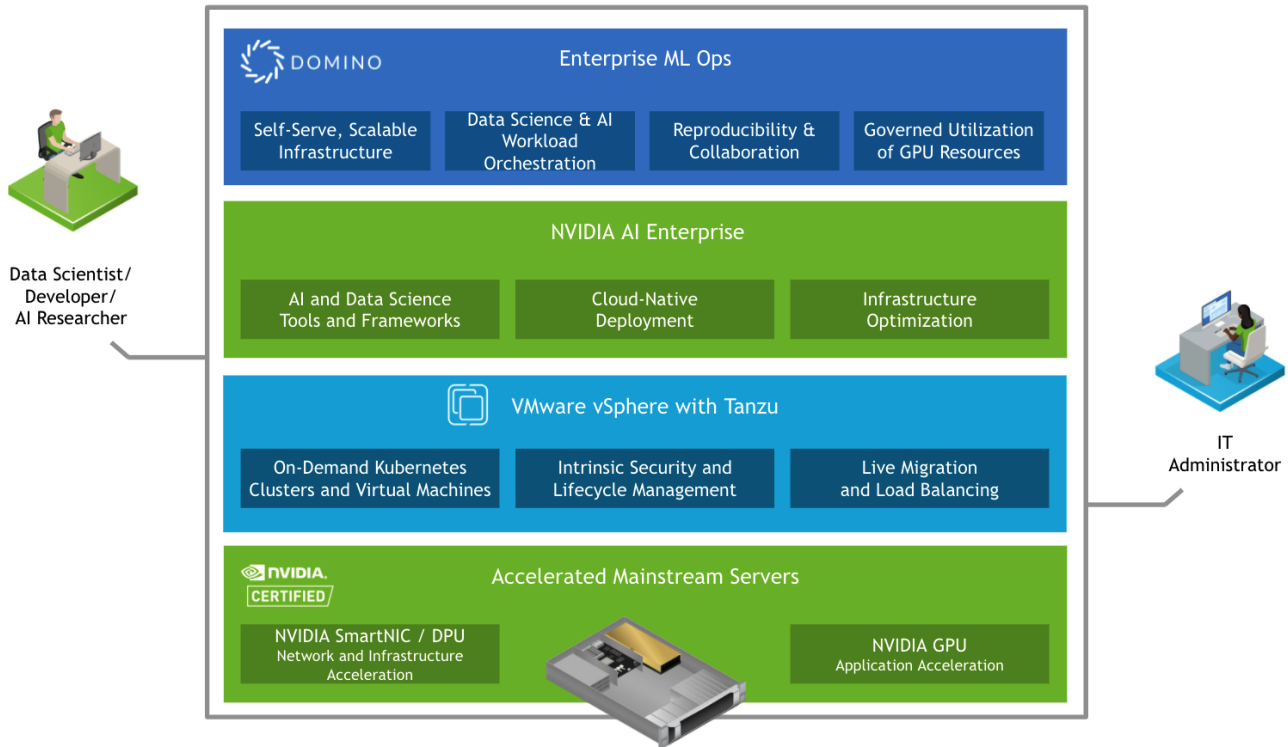
- ✓ Build and deploy models faster with improved collaboration, reuse, and reproducibility while continuously improving knowledge and best practices.
- ✓ Automatically launch Domino Workspaces on-demand, with docker images configured with the latest data science tools and frameworks - optimized with pre-allocated NVIDIA GPUs.
- ✓ Get more models into production faster, with cost-effective scale for enterprise-wide deployments in existing enterprise infrastructure.

IT Teams can orchestrate data science workloads centrally in existing infrastructure

Provide data scientists self-service access to their preferred IDEs, languages, and packages so they can focus on data science innovation. Reduce IT costs, management, and support burden with tools and NVIDIA infrastructure consolidated and orchestrated in a central location across projects and teams.

- ✓ Fold data science workflows into existing infrastructure with VMware vSphere® with Tanzu running on industry-leading NVIDIA-Certified™ systems for enterprise-grade security, manageability, and support.
- ✓ Achieve bare-metal performance for AI and data science workloads, optimized to cost-effectively scale up and out for enterprise wide deployments.
- ✓ Deploy data science workloads allocated with GPU resources, with governance and reporting to ensure proper utilization.

Cost-effective scale for enterprise data science and AI deployments



Unleash data science to accelerate model velocity

Deliver business value from models faster. Rapidly deploy, manage, and scale model-driven data science workloads in the modern hybrid cloud with ready-to-use containers - automatically configured and compatible with Domino - deployed on prem, in the cloud, or the edge.

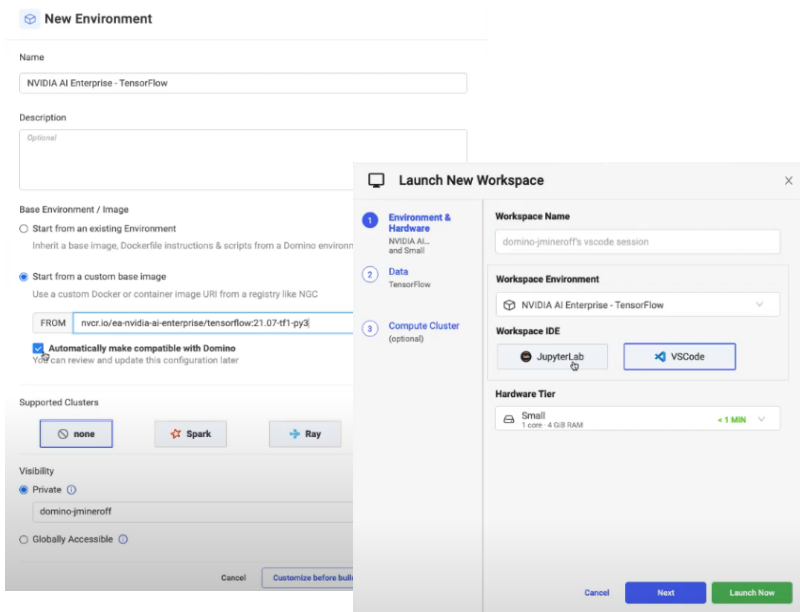


Figure 2: A Domino Environment is created from an [NVIDIA NGC TensorFlow container](#) image, automatically adding the appropriate Dockerfile instructions to ensure compatibility, configuring JupyterLab and VSCode for availability. Then, a workspace with pre-configured hardware tiers can be launched.

Optimized, certified, and supported

Optimized, certified, and supported on VMware vSphere with VMware Tanzu with [NVIDIA-Certified Systems](#), NVIDIA AI Enterprise suite is purpose-built to meet the demands of enterprise AI and data science, delivering the fastest start in AI development, effortless productivity, and revolutionary performance.

[Domino's validation for NVIDIA AI Enterprise](#) accelerates data science initiatives—at velocity and scale—by complementing NVIDIA AI Enterprise with the best-in-class Enterprise MLOps platform.

Learn More

dominodatalab.com/partners/nvidia

Domino Website: dominodatalab.com

Blog: blog.dominodatalab.com

Demo: dominodatalab.com/demo

NVIDIA Website:

<https://www.nvidia.com/en-us/data-center/products/ai-enterprise-suite/>

Blog:

blogs.nvidia.com/blog/category/enterprise