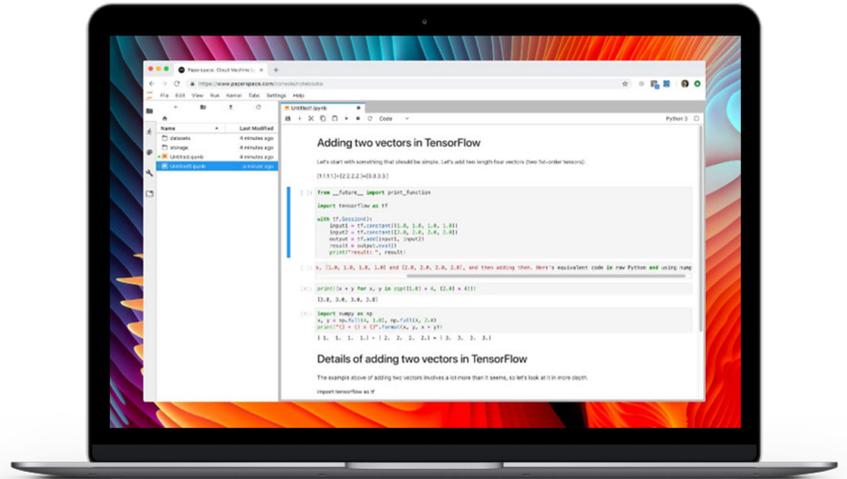


BUILD, TRAIN, AND DEPLOY DEEP LEARNING MODELS AT SCALE.

From exploration to production, Gradient enables teams to quickly develop, track, and collaborate on deep learning models.



We put data scientists first.

Gradient removes the time sink of infrastructure management and setup by providing a ready-to-use platform and essential tools. For organizations, it reduces project costs by maximizing the efficiency of their data science teams and hardware resources.

Scalability

Gradient enables running many concurrent experiments and conducting large-scale distributed deep learning.

Zero DevOps required

Gradient fully automates infrastructure orchestration and reallocates resources seamlessly to maximize efficiency.

Full Flexibility

Gradient integrates into any existing ML stack and supports TensorFlow, PyTorch, as well as custom Docker containers.

Infrastructure-Agnostic

Gradient can run on any public cloud, private cloud, or on-premise. Easily switch between solutions without changing any workflows.

AI Ops Control Plane

User-friendly UI & CLI for team members across Jupyter notebooks, experiments, models, deployments, and GradientCI®.

Lower Infrastructure Cost

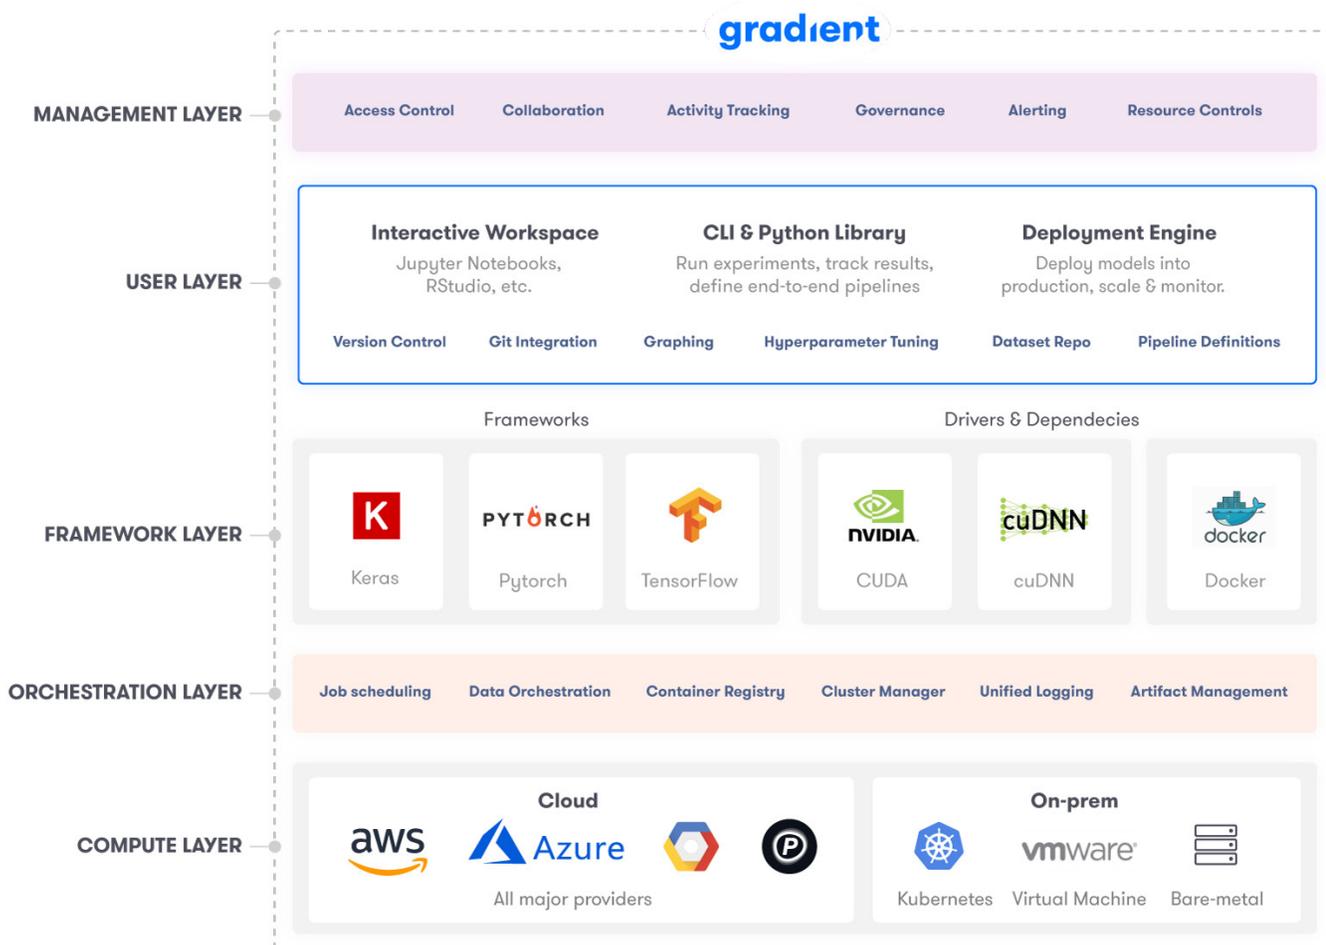
Gradient optimizes resource utilization and enables lower computing costs through its innovative cluster management technology.

GRADIENT MAKES DEEP LEARNING WORKFLOWS EFFICIENT

Gradient is the serverless operating system for AI. It supports any underlying infrastructure and abstracts away the complex orchestration process by providing concise and intuitive interfaces.

For data scientists, Gradient provides the freedom to use the tools they are familiar with. In a working environment that lacks time-consuming DevOps tasks and resource orchestration, teams can focus on what matters: training algorithms and creating value.

For organizations, Gradient significantly reduces the infrastructure costs of deep learning research. Gradient's focus on optimizing the utilization of available resources and flexible deployment scenarios can save up to 80% of compute costs compared to non-optimized systems.



Gradient is a product of Paperspace, a cloud computing company. Funded by investors such as Intel Capital, Y Combinator and Battery Ventures, the company has offices in Brooklyn, Copenhagen, Amsterdam, and Gdansk.

INTERESTED IN LEARNING MORE?

Visit www.gradient.cloud You can contact our sales team at **(718) 705-3826** **+31621884054** for EMEA, or email us at hello@paperspace.com