



SPEED MACHINE LEARNING WORKFLOWS TO DRIVE FASTER INSIGHTS

Deploy models from pilot to production with an ML Ops solution delivered as a service



HPE GreenLake brings the cloud experience—self-serve, pay-per-use, scale up and down and managed for you—to apps and data everywhere, in edges, colocations, and data centers. It enables you to free up capital, boost operational and financial flexibility, and free up talent to accelerate what's next for you.

Machine learning (ML) models quickly deliver insights from large datasets and enable enterprises to make data-driven business decisions. To fully realize the value of ML—including new revenue streams and improved customer experiences—enterprises need to implement fully operational ML workflows.

Yet, machine learning operations (ML Ops) can be challenging. It requires standardization, while at the same time supporting the experimentation required in the ML workflow and ensuring collaboration between data scientists, developers, and IT operations professionals. Addressing the entire ML lifecycle truly eliminates siloed environments. And to stay current, an ML Ops solution needs to build and adopt the best innovations from the open source community without vendor lock-in.

Teams also require fast access to scalable, secure infrastructure no matter where data lives. By reducing the time to procure, set up, and deploy ML infrastructure, your company

can significantly improve productivity for data scientists and developers, thereby speeding the time to insight discovery.

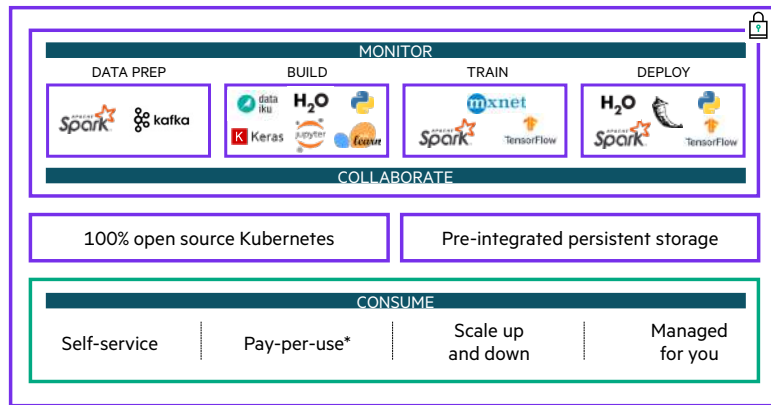
Organizations need an end-to-end solution that brings DevOps-like speed and agility throughout the ML lifecycle—from pilot to production. [HPE GreenLake](#) now makes it possible to realize faster time to value with operationalized ML across your business.

Reduce the typical six-to-nine month infrastructure procurement cycle to weeks with HPE GreenLake. It's delivered to your on-premises locations and monitored and managed to reduce risks, control costs, and free up IT resources. This enables data scientists to quickly operationalize their ML models.

Standardize ML workflows, improve productivity, and accelerate time to value.

“Artificial intelligence and machine learning are changing the future while disrupting the present. While it is deceptively easy to launch successful AI pilots, moving toward an organization-wide program to realize the full value potential is fiendishly hard. Using ML Ops to effectively manage and govern the AI lifecycle from experimentation to production is the business’s next competitive frontier.”¹

– Ritu Jyoti, Program Vice President, Artificial Intelligence Research, Global AI Research Lead, IDC.



* Minimum commitment may apply

FIGURE 1. Cloud services platform for ML Ops with HPE GreenLake

SPEED IT RESOURCES TO DATA SCIENTISTS

Operationalize the end-to-end ML lifecycle with a fully managed service from HPE GreenLake. It’s a turnkey, enterprise-grade solution with integrated security and a choice of data science open-source and ISV tools. It can support any open-source tool or framework allowing enterprise data science teams to keep up with the latest innovations from the open-source community. Data scientists can rapidly deploy, test, and run ML models and collaborate with developers to analyze results. And because it’s HPE GreenLake, you only pay for the infrastructure you actually use.

Accelerate business outcomes

- Gain insights faster with seamless deployment and improved collaboration among data scientists, developers, and IT operations

Get to work fast

- Spin up ML development, test, and production environments in minutes with a plug-and-play, turnkey solution that leverages containers
- Connect with existing data lakes or data warehouses for near-instant access to data resources wherever they live

Reduce risk

- Achieve out-of-the-box, enterprise-grade security and access controls with data governance, ML model auditability, and provenance tracking

Improve productivity

- Reduce latency issues typically associated with public cloud ML modeling. Ready-to-go infrastructure can be delivered quickly to speed modeling timetables

Enable scalability

- Avoid up-front costs, and pay only for the infrastructure and services consumed, with elastic capacity to support business demand. Also, help eliminate hidden costs such as cleaning, storing, aggregating, labeling, and creating data flows that surround building your own AI/ML platform

Simplify IT

- Offload the heavy lifting of managing the infrastructure with HPE GreenLake for installation, implementation, maintenance, and management

SPEED UP BUSINESS OUTCOMES

Jumpstart ML projects

Enable rapid deployment across the enterprise with intuitive graphical user interface, plus simple provisioning, and standardized toolsets

Speed insight generation

Unlock the value of data no matter where it lives—in data centers, co-location facilities, and at the edges—with a turnkey solution that enables teams to spin up environments in minutes, rather than days or weeks

Improve compliance

Unify visibility and control with services that plug-and-play with existing compliance tools, while avoiding the risk of moving or transferring data into public clouds

Get started with an ML Ops strategy today.

LEARN MORE AT
hpe.com/greenlake

¹ IDC’s Worldwide Consumption-Based IT Infrastructure Taxonomy, 2020

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



Get updates