DigitalOcean delivers advanced AI capabilities with high-density infrastructure deployment

Customer story

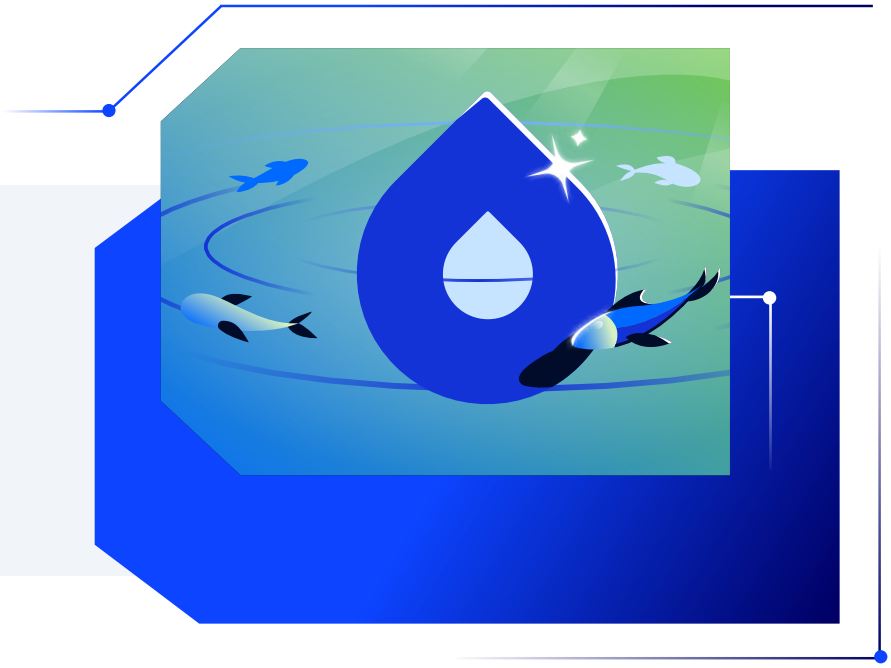

DigitalOcean

Modern data center design empowers customer innovation with scalable GPU infrastructure.

Flexential solution portfolio

Colocation

Interconnection

Flexential partnership highlights

- Flexential Atlanta-Douglasville data center enables expanded GPU computing through dual data halls accelerated by NVIDIA's Hopper GPUs and AMD Instinct GPUs, significantly increasing DigitalOcean's capacity to provide customers with scalable computing power on demand
- Advanced liquid cooling system supports 1,500-plus watts per square foot, maximizing computing power while minimizing operational costs
- FlexAnywhere® platform provides nationwide network connectivity and scalability, enabling DigitalOcean to expand GPU computing services and support AI development across multiple markets

The history

Headquartered in Broomfield, Colorado, DigitalOcean is the simplest scalable cloud platform that democratizes cloud and AI for growing tech companies around the world.

According to the Flexential State of AI Infrastructure Report, 44% of organizations face infrastructure constraints as their primary barrier to AI expansion. DigitalOcean tackles this challenge head-on by simplifying cloud computing and AI to allow builders to spend more time creating software that changes the world.

More than 600,000 customers trust DigitalOcean to deliver the cloud, AI, and ML infrastructure they need to build and scale their organizations.

Starting with their flagship Droplet virtual machines, DigitalOcean has evolved into a comprehensive cloud platform provider serving innovative businesses across the growth spectrum. From early-stage startups to rapidly scaling companies, DigitalOcean's suite of solutions, including virtual machines and managed Kubernetes, delivers enterprise-grade capabilities with the simplicity that small and medium-sized businesses demand.

> Flexential's Atlanta-Douglasville facility addresses our need to bolster GPU infrastructure quickly as our AI/ML business continues to grow. This expansion extends our footprint and supports the specific power and cooling requirements of high-density GPU environments, positioning us to meet our customers' expanding AI computing needs."

Adam Knapp,
Vice President of Engineering, DigitalOcean

The challenge

Growing AI demand requires high-density, high-performance infrastructure

With surging demand for AI/ML capabilities, DigitalOcean needed to rapidly scale its GPU infrastructure while maintaining operational efficiency and meeting rigorous power, cooling, and networking requirements for high-density computing environments.

"We saw a high demand from startups for dedicated high-performance GPU clusters to run custom AI model training," said Adam Knapp, Vice President of Engineering for DigitalOcean. "We responded with NVIDIA Hopper-accelerated bare metal solutions, scaling from two to 32 nodes, which drove our AI infrastructure revenue up 200%."

DigitalOcean's rapid expansion required a data center partner who could deliver on three critical factors: aggressive timelines, high-density requirements, and scalability for future growth.

"Building infrastructure for AI workloads is fundamentally different from traditional clouds," said Knapp. "The networking requirements for AI training demand high bandwidth and low latency. That's why we've invested heavily in technologies like NVIDIA NVLINK and InfiniBand with NVIDIA NCCL, which enables us to optimize our fabrics for seamless multi-node communications. But it's not just about internal traffic; we're also scaling our capacity to handle massive data movement in and out of the environment. And of course, power and cooling are equally critical in this equation."

Offering customers the highest level of physical and digital security was also paramount for DigitalOcean. Requirements included a comprehensive security framework that would protect critical assets 24/7. This meant implementing multiple layers of physical access control, from dual-factor authentication to biometric scanning, supported by continuous security monitoring and surveillance systems. These stringent security requirements were essential to safeguard customer data and protect high-value AI computing resources.

Details: DigitalOcean wanted more than a standard data center space. The company required a strategic partner who could bring AI and other high-performance initiatives online with the reliability of uniquely built space, power, and cooling capabilities.

The solution

Flexential delivers the foundation for speed, performance, and scalability

The Flexential Atlanta-Douglasville data center was uniquely positioned to deliver on DigitalOcean's requirements: an efficient high-density compute-intensive environment, featuring built-in liquid cooling capabilities and 100% service-level agreements on power, cooling, network, and bandwidth.

The deployment significantly increased DigitalOcean's capacity to provide customers scalable computing power on demand, ideal for AI/ML training and inference, deep learning, high-performance computing (HPC), data analytics, and graphics/video rendering tasks.

"Flexential's Atlanta-Douglasville facility addresses our need to bolster GPU infrastructure quickly as our AI/ML business continues to grow. This expansion extends our footprint and supports the specific power and cooling requirements of high-density GPU environments, positioning us to meet our customers' expanding AI computing needs."

Through the FlexAnywhere® platform, the facility provides DigitalOcean with the connectivity and scalability needed to expand its GPU computing services. This infrastructure foundation enables developers and businesses to build next-generation AI applications while maintaining the ability to scale seamlessly across markets.

"Our customers need to run production-grade AI applications, which demand lower latency, tighter orchestration integration, and efficient scaling," said Knapp. "We are modernizing our entire stack, not just the hardware. Our Gen AI platform is a fully managed service that enables customers to build and scale AI agents. Ultimately, we're not just selling GPUs, we are helping customers move from prototyping to production in a simplistic way with an AI stack that can scale."

In addition, Flexential eliminates security concerns by offering industry-leading capabilities that include round-the-clock security, dual-factor authentication access, biometric scanners, and monitored security cameras.

> We wanted to go bigger and be more deliberate, which led to our strategic investment in the Flexential Atlanta-Douglasville data center. Flexential could deliver the density, improved power envelope, and networking performance needed for the next wave of AI workloads. But our story doesn't end with training, we're seeing a dramatic shift towards inferencing, where AI models move from development into real-world applications requiring more efficient scaling."
>
> **Adam Knapp,**
> Vice President of Engineering

Key features: DigitalOcean's rapid expansion required a strategic partner who could deliver on aggressive timelines, high-density requirements, and scalability for future growth. The Flexential Atlanta-Douglasville facility emerged as the ideal solution, offering both immediate capabilities and long-term partnership positioned for growth.

Implementation: To meet customers' needs, DigitalOcean required rapid deployment with strict adherence to timelines and technical specifications.

Drawing on extensive experience delivering high-density environments, Flexential technical teams applied tested deployment methodologies and precise project execution to ensure the space was commissioned on schedule.

The implementation centered on enabling additional DigitalOcean GPU Droplets across two connected data halls, accelerated by NVIDIA's Hopper GPUs and AMD Instinct GPUs.

Each phase of the deployment was carefully managed by specialized teams, ensuring seamless integration of power, cooling, and networking infrastructure while maintaining strict timeline adherence.

"Speed is part of our DNA, and Flexential's infrastructure enables us to move fast while maintaining operational excellence," said Knapp. "Their purpose-built facilities and rapid deployment capabilities matched our pace of innovation, allowing us to quickly scale AI services without compromising on performance."

The results

Meeting AI demands without compromising performance

Performance: The high-density deployment at the Flexential Atlanta-Douglasville data center significantly increased DigitalOcean's capacity to provide customers with scalable computing power on demand.

Efficiency: Flexential high-capacity data centers are uniquely designed to handle the requirements of compute-intensive environments, while ensuring reliable performance and scalability.

Customer experience: DigitalOcean's AI infrastructure enables customers to quickly deploy GPU resources and seamlessly integrate AI capabilities from experimentation to production.

Transform AI possibilities into reality with infrastructure built for innovation

Discover Flexential High-Density Infrastructure Capabilities →