



Benchmarking AI infrastructure readiness

What IT leaders need to know
about **building future-ready,
resilient AI infrastructure**

The AI infrastructure imperative

As artificial intelligence moves from hype to real business results, the pressure is on IT leaders to deliver infrastructure that's ready for what comes next. According to the State of AI Infrastructure Report, 90% of companies are deploying or planning to deploy generative AI, but nearly half face significant roadblocks with their current IT environments.

51% of organizations expect a return on AI investment within a year, making rapid infrastructure readiness a top priority.

– [State of AI Infrastructure Report](#)

Companies are no longer just experimenting with AI; they're integrating it directly into products, workflows, and customer experiences. This sense of urgency isn't limited to a single industry; it's being felt everywhere, from fast-moving startups to established research universities.

With high financial expectations and relentless competitive pressure, IT teams are being asked to move faster than ever. Infrastructure decisions made now will set the pace for innovation and determine who leads as AI reshapes every part of the business.

"If you can cut 10 hours off a sales workflow or create a differentiated user experience with AI, you're either going to be first—or left very far behind.

– Adam Knapp, DigitalOcean

How ready is your AI infrastructure?

AI adoption is accelerating faster than most infrastructure strategies were built to handle. According to the State of AI Infrastructure Report:

44% of IT leaders cite infrastructure constraints as the biggest hurdle to scaling generative AI


86% of leaders are concerned about finding or developing enough AI talent

Overcoming infrastructure hurdles

As AI projects move from pilot to production, many organizations are discovering that their existing infrastructure can't keep up with the demands of advanced workloads. Power, cooling, network performance, and supply chain constraints are now central challenges for IT teams everywhere.

Getting infrastructure "AI-ready" takes more than just adding GPUs. Reliable power, scalable networks, data locality, and forward-thinking partnerships are all essential. Strong supply chain partnerships are just as important, as timing can make or break an AI deployment, especially when it comes to securing and installing the latest GPUs. Purpose-built data centers like Flexential deliver the reliable power, high density, and flexibility needed to support rapid AI growth.

Did you know?



44% of IT leaders cite infrastructure constraints as the biggest hurdle to scaling generative AI across their organizations.



The latest GPUs draw significantly more power than typical compute infrastructure. Without a thoughtful design, you hit density and thermal limits fast."

– Adam Knapp, DigitalOcean



Our centralized superpod lets us run experiments of any scale, all in one place—without scattering resources across multiple sites."

– Kenneth Chaney
University of Pennsylvania

From designing dense, energy-efficient data center environments to tackling the realities of bandwidth and data locality, these leaders are setting the pace for what's possible in AI-driven business and research.



Talent, planning, and scalability

AI infrastructure isn't just a technical challenge. It's also a people and planning challenge, too. Most organizations realize that new hardware is only part of the equation. Success comes down to having the right skills, a flexible strategy, and the ability to keep pace with rapid change.

The State of AI Infrastructure Report shows that 86% of leaders are concerned about finding or developing enough AI talent. With demand for specialized skills outpacing supply, both enterprise cloud providers and research institutions are taking creative steps to fill the gap. This means investing in current teams, partnering with outside experts, and building an environment where learning and adaptation are the norm.

Universities are uniquely positioned to develop the next wave of AI talent, but they're also focused on democratizing access, making advanced AI tools available to more departments and researchers. This approach doesn't just grow expertise; it opens the door to broader innovation across campus.

For enterprise teams, the path looks a little different but requires just as much flexibility. DigitalOcean, for example, combines upskilling with strategic hiring and close collaboration with technology partners. The ability to quickly respond to changing infrastructure needs—rather than relying on old 12- to 18-month planning cycles—gives them a real edge.

Continuous improvement is built into the operating rhythm. Weekly reviews focus on both operational excellence and customer experience, so teams can move fast, learn from what works (and what doesn't), and keep scaling without getting stuck.

“You really have to be approaching it from every angle. We invest in our team, acquire talent, and partner with specialists who bring deep expertise in AI workloads.”

– Adam Knapp, DigitalOcean



The new requirements: Security, networking, and sustainability

Scaling AI brings new challenges that go far beyond deploying the latest hardware. Security, network performance, and sustainability are now at the top of every IT leader's agenda, often for the first time at this scale.


53% of organizations report excessive latency in AI workloads,
and 51% say AI expansion has increased their attack surface.

– [State of AI Infrastructure Report](#)

AI workloads demand high-performance, low-latency networking and robust security from day one. Leaders are rearchitecting their environments for real-time east-west data traffic, tight access controls, and strong isolation. They're also leaning on software-defined networking (SDN) to dynamically manage traffic flows and keep GPU clusters operating efficiently, even as needs change.

Universities are also building a strong governance foundation, forming committees to tackle AI model transparency, bias detection, and ethical standards. By sharing best practices across departments, they're helping set the pace for responsible AI use both in academia and beyond.

At the same time, sustainability is under the microscope. The pressure to maximize performance per watt and avoid overbuilding is real, especially as 79% of leaders say they feel more pressure than ever to adopt sustainable infrastructure. Facilities that offer high density, energy efficiency, and flexible power options—like Flexential's—make it possible to run the world's most advanced AI workloads without sacrificing long-term goals.



79% of IT leaders say they feel more pressure to adopt sustainable infrastructure as AI expands.

“

We designed strong isolation and hardened controls from the start. Our security team is engaged at every level to make sure customer data is protected across the stack.”

– Adam Knapp, DigitalOcean

Ready to build your AI advantage?

AI is changing how businesses and research organizations operate, but it's not just about adopting new technology. Success depends on having resilient infrastructure, a strong network of partners, and teams that can adapt quickly.

Key lessons from our panelists:

- Don't underestimate the basics. Power, cooling, and data locality are just as critical as GPUs.
- Focus on people and culture as much as hardware. Talent development and collaboration make the difference.
- Build for change. The AI landscape shifts fast, and flexibility is essential for long-term results.
- Make sustainability part of your plan from day one. The most successful organizations are those that plan ahead.

"If your infrastructure can't adapt quickly, you're going to struggle to meet changing customer demands or jump on new opportunities when they present themselves."

– Adam Knapp, DigitalOcean

Curious how your organization stacks up or want to see what's next for AI infrastructure?

Watch the full on-demand webinar for more real-world stories, strategies, and data-driven guidance from industry leaders.

Watch the Webinar

