



# The Global Boom of AI Data Centers: Billions, Gigawatts, and the Race for Compute

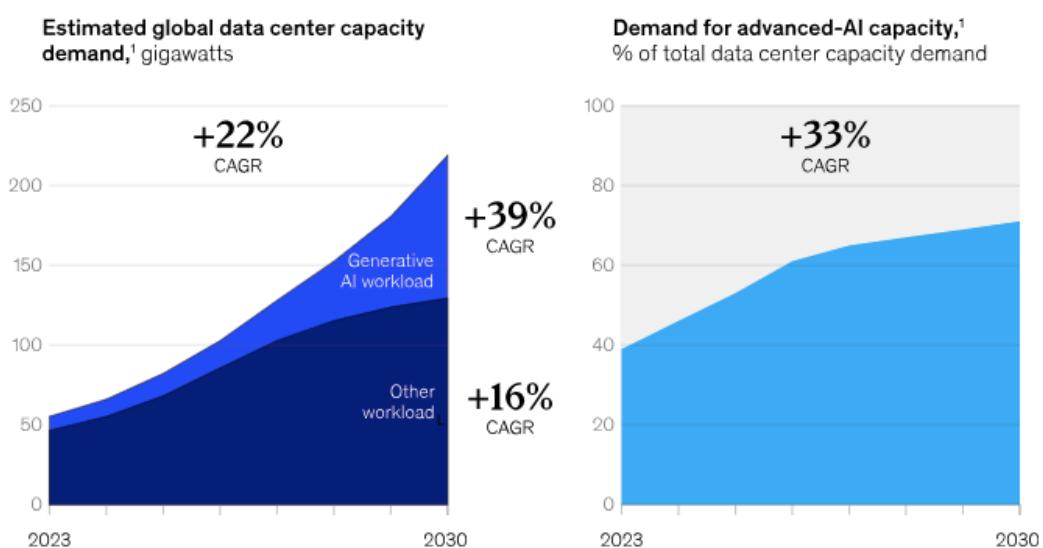
September 24, 2025

AI has moved beyond code, startups, and breakthroughs in research. The true frontier now is infrastructure—vast campuses of silicon and steel, devouring power and land at unprecedented speed. Compute has become the new oil, and nations are competing fiercely to host the next wave of hyperscale AI data centers. What began as experiments in cloud regions is now a geopolitical arms race: who controls the gigawatts of electricity, the racks of Nvidia GPUs, the networks and cooling systems that allow the largest models to exist. New Zealand is making its bid in the Pacific, the United States is running at industrial-revolution speed, Microsoft is pushing technical boundaries, while China, the Middle East, and Europe refuse to be left behind. The boom is indisputable, and its effects are reshaping economies and energy systems worldwide.

## Amazon's South Pacific Gamble

In New Zealand, Amazon Web Services has at last launched its Asia Pacific (New Zealand) Region, resurrecting a plan first announced in 2021. The commitment is more than NZ\$7.5 billion, with the promise of a thousand jobs per year and a lift of NZ\$10.8 billion to GDP. For Wellington, still nursing a sluggish economy, the timing is critical. Prime Minister Christopher Luxon admitted the country is a costly place to build, but foreign investment is the strategy for recovery. Customers like Xero and Kiwibank now gain a homegrown cloud option with lower latency, while planning reforms and visa changes are designed to make the country more attractive to future investors. The locations remain secret, but the intent is plain: even small economies want a seat at the AI infrastructure table.

**AI is the key driver of growth in demand for data center capacity.**



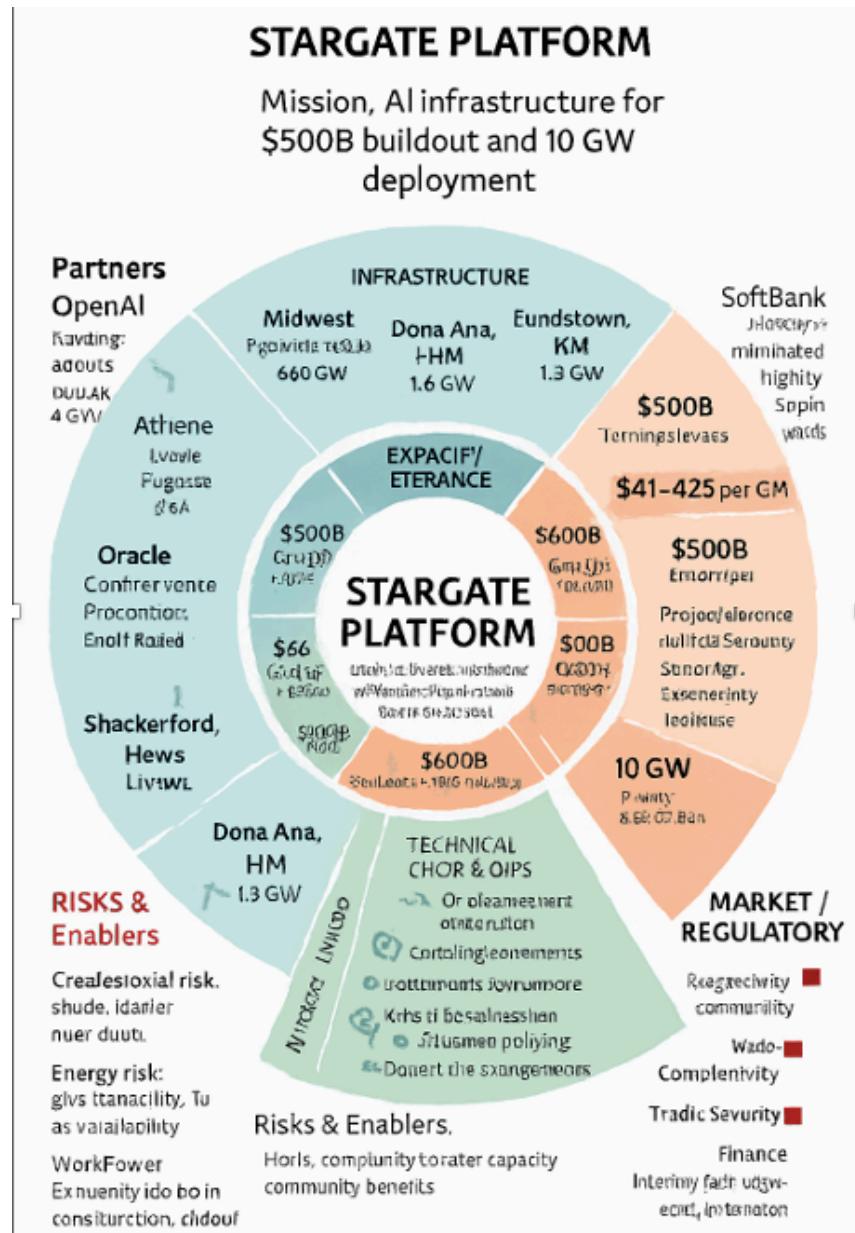
Demand for AI-ready capacity is the main driver of this potential deficit—as it must provide the high computational power and power density required by AI workloads. Our analysis suggests that demand for AI-ready data center capacity will rise at an average rate of 33 percent a year between 2023 and 2030 in a midrange scenario. This means that around 70 percent of total demand for data center capacity will be for data centers equipped to host advanced-AI workloads by 2030. Gen AI, currently the fastest-growing advanced-AI use case, will account for around 40 percent of the total

**[McKinsey & Company]**

## Stargate: The U.S. Mega-Build

No country is moving faster than the United States. OpenAI, Oracle, and SoftBank are executing Stargate, an infrastructure program of a scale rarely seen outside energy or aerospace. Five new sites are being added to the flagship in Abilene, Texas, and CoreWeave-linked builds, bringing total planned capacity to nearly 7 gigawatts and over \$400 billion in capital over three years. The \$500 billion, 10-gigawatt target—announced in January at the White House—is already on track to be achieved by the end of 2025, ahead of schedule.

The geography is deliberate: Shackelford County in Texas, Doña Ana County in New Mexico, a still-unnamed Midwest site, expansions near Abilene, plus Lordstown, Ohio and Milam County, Texas for SoftBank. Abilene is already alive with Oracle Cloud Infrastructure and Nvidia GB200 racks. Lordstown is scheduled to be operational next year. At least 25,000 direct jobs and tens of thousands of indirect roles will ripple out from these projects. The financial model is wholesale, anchor-tenant heavy, tied to long-term PPAs and public-private agreements. The outcome is more compute density than the world has ever seen, built at a pace that reflects both ambition and necessity.



## Microsoft's Wisconsin Super-Campus

Microsoft is competing with scale and engineering. In Wisconsin, the company is building what it calls the most advanced AI data center in the world: three massive buildings over 315 acres, 1.2 million square feet, hosting hundreds of thousands of Nvidia GB200 GPUs arranged in two-story racks for latency gains. Cabling stretches 120 miles underground, piping nearly 73 miles, with a closed-loop water cooling system that recycles more than 90 percent of its

input. Microsoft insists its water use is lower than a golf course in summer, a calculation designed to calm local fears. It also agreed to pay higher utility rates itself to avoid passing energy costs onto residents.

This is just one of three “Fairwater” sites, with another in Georgia and a second Wisconsin campus announced for \$4 billion. The land was acquired from Foxconn’s failed \$10 billion project. What Foxconn promised in manufacturing, Microsoft is delivering in compute.



Microsoft said its \$3.3 billion Wisconsin data center will come online in early 2026 and initially be used to train OpenAI models. Photo: Microsoft

## China's Parallel Push

China is racing in parallel, though with less transparency. Beijing has prioritized AI infrastructure as part of its “East Data, West Compute” strategy, designed to balance power-hungry AI clusters with renewable energy-rich western provinces. Massive GPU farms are being built in Inner Mongolia, Ningxia, and Gansu, while coastal hubs like Shanghai and Shenzhen push for lower-latency deployments closer to financial and industrial centers. Baidu, Alibaba Cloud, and Tencent are all investing heavily, with Huawei pushing its Ascend series as a domestic alternative to Nvidia chips under U.S. export controls. Chinese clusters are measured in gigawatts too, though exact figures are tightly held. The strategy is national: ensure compute sovereignty, insulate from sanctions, and make AI an engine of economic and military modernization.

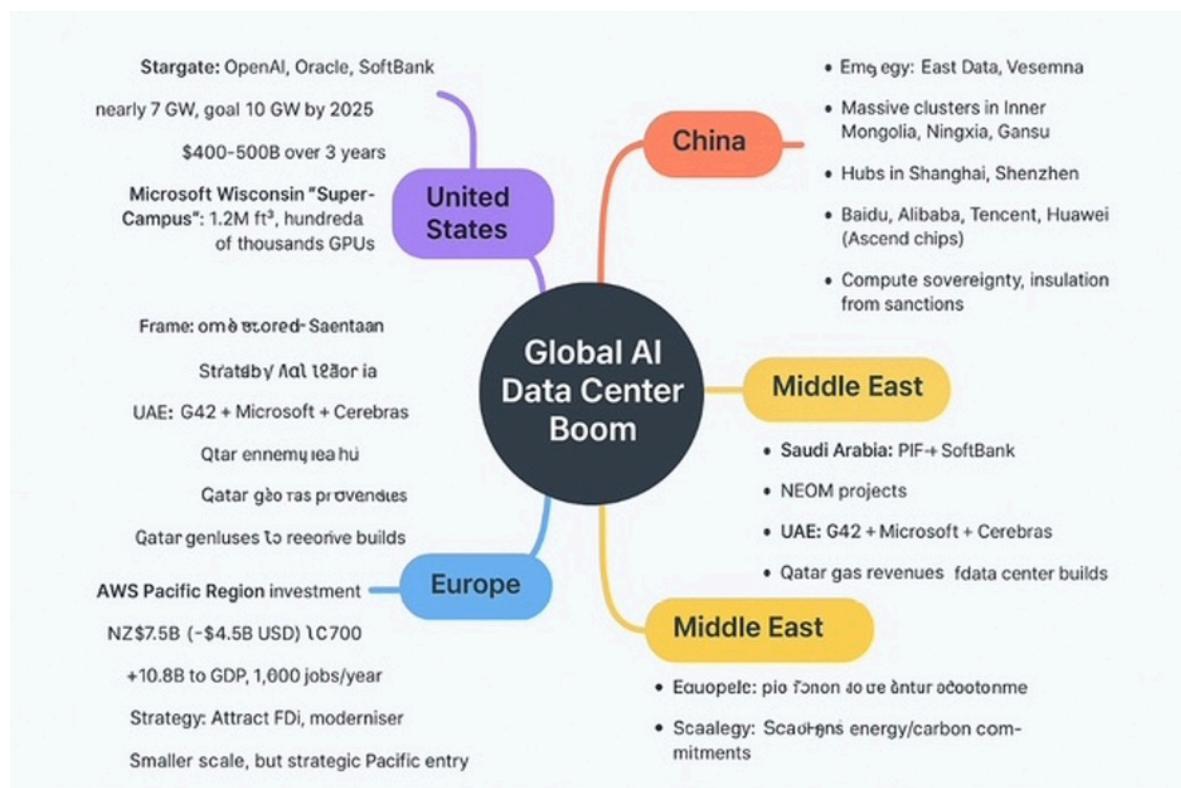
## Middle East: Oil into Compute

In the Gulf, sovereign wealth funds are converting petrodollars into AI power plants. Saudi Arabia’s Public Investment Fund has teamed with SoftBank and U.S. firms to site hyperscale

campuses, while NEOM's green energy projects are being paired with AI infrastructure plans. The UAE, through G42 and partnerships with Microsoft and Cerebras, is positioning itself as a hub for large-scale training runs, blending its role as an oil exporter with ambitions as an AI compute exporter. Qatar, flush with gas revenues, has followed suit. Here, data centers are framed not just as tech projects but as part of national diversification strategies. For states whose wealth has long been tied to hydrocarbons, AI data centers are becoming a bridge into the post-oil economy.

## Europe's Balancing Act

Europe is caught between ambition and regulation. France is the most aggressive, with President Emmanuel Macron openly courting Microsoft and Amazon while financing French startups like Mistral AI. Paris and Marseille are emerging as hubs, driven by subsea cable landings and proximity to Africa and the Mediterranean. Germany lags, slowed by high energy costs and regulatory caution, though Frankfurt retains importance as a European data hub. The UK, post-Brexit, is trying to reinvent itself as an AI-friendly destination, with projects tied to Oxford, Cambridge, and London, and with Microsoft and Google Cloud both expanding capacity. Yet Europe's challenge remains energy: electricity prices, carbon commitments, and strict environmental regulation make scaling multi-gigawatt campuses far harder than in the U.S. or China.



## Conclusion

From Wellington to Wisconsin, Shanghai to Riyadh, the AI boom is no longer about apps or models, but about the physical infrastructure to host them. The U.S. leads in scale and speed,

building multi-gigawatt campuses with Nvidia racks shipping by the truckload. China is pursuing sovereignty with its East Data, West Compute model, circumventing export restrictions. The Gulf states are buying their way into the race with sovereign capital, while Europe cautiously balances ambition with regulation. Even smaller economies like New Zealand are retooling laws to attract foreign tech giants.

The new geopolitics of AI is built in concrete, powered by gigawatts, cooled by water and air, and financed in the hundreds of billions. Whoever controls the data centers controls the future of artificial intelligence—not in theory, but in steel, silicon, and electricity.