

Can data centers keep pace with AI?

How different growth trajectories for AI adoption
could shape global data center demand



In the AI era, data centers have moved from a niche infrastructure asset to a principal engine of economic and competitive advantage. Our latest report explores why the question is no longer whether data center demand will grow, but whether the supporting infrastructure can scale quickly enough to support it.

The AI race is becoming an infrastructure race

To understand the implications, NTT Global Data Centers commissioned economic consultancy ThoughtLab to model three potential growth pathways for global data center capacity. These pathways were stress tested against real-world scenarios to identify pressure points, where innovation helps and what changes are needed to align AI ambition with infrastructure reality.

ThoughtLab



In an **accelerated** data center growth scenario (30% CAGR), explosive enterprise AI adoption and heavy investment drive an AI infrastructure supercycle, but also place considerable strain on supporting infrastructure resources, including grid connections, transformers, backup generators and construction labor.



With **slow growth** (~16% GAGR), AI demand is expected to increase, but external regulatory, supply chain and power constraints, as well as economic, geopolitical and business concerns limit growth.



In a **steady** (base case) scenario (23% CAGR), data center growth is strong, but only because the industry continually adapts to mounting regulatory restrictions, supply chain bottlenecks and trade disruptions through advancements in cooling, IT hardware and digital optimization.

However, across all scenarios, it is infrastructure, rather than AI demand, causing a capacity crunch.

The constraint is not one headwind, but the convergence of many: grid connections, power availability, transformers, chips, backup generators, suitable land, capital and specialized construction and operations labor.

Key findings:



Power is the central infrastructure constraint: US data center electricity requirements are expected to equal ~50% of planned capacity additions from 2026 to 2030. While Europe faces lower aggregate demand pressures, it also has greater complexity in grid integration and stability.



Capacity growth concentrates geographically: About 57% of new capacity is projected for the US, 25% for APAC, 8% for Europe and 11% for the rest of the world. This clustering of data centers is driven by access to power infrastructure, skilled labor, connectivity, and proximity to major population centers.



Layered regulatory environments add cost and complexity: The compliance environment spans energy reporting requirements, sector-specific rules, national security screening of foreign investment and data localization mandates. This raises costs and slows deployment, particularly for operators with multinational footprints.



Equipment bottlenecks extend beyond chips: In addition to processor shortages, power transformers and switchgear carry extended lead times. Backup generators face both supply constraints and increasing regulatory scrutiny.



Water consumption requirements can be overstated: While friction remains in water-stressed regions and outdated cooling systems can consume millions of gallons a day, newer technologies rectify this. Modern closed-loop, direct-to-chip, immersion and water-lean cooling technologies can reduce or eliminate water use.



Community consent is critical: Public sentiment toward data center construction worsened in 2025 and 2026, with opposition increasingly delaying or canceling projects.

What data center industry leaders should do now



Set consistent, ambitious Power Usage Effectiveness (PUE) and Water Usage Effectiveness (WUE) standards, while clearly communicating modern cooling performance and water efficient approaches to demonstrate industry commitments to policymakers, and dispel misconceptions held by communities and the media.



Accelerate innovation in processors, cooling, AI-driven workload optimization and efficient model architectures, to help position data center builders and operators as active shapers of more sustainable and performant AI infrastructure stacks.



Treat supply chain resilience as a core design and financing requirement, not an afterthought. Diversifying suppliers for critical components, negotiating long-term purchase agreements with manufacturers and holding inventories of long-lead items help manage risks. Standardizing design and equipment specifications across site prevents the need for redesign.



Demonstrate local value through early, transparent community engagement, and shared economic benefits. Visibly link data center revenues to local priorities (including schools, parks, jobs or transit) to address opposition. Joint communications between policymakers and operations, rather than closed door deals and NDAs, can help build citizen trust and dispel misconceptions.



Co-plan power and grid infrastructure with utilities, so data centers become grid-supportive resources rather than unmanaged loads. A coordinated approach to growth, aligned with investments in generation, transmission storage and demand side resources, helps prevent capacity shortages and project delays. Showing that facilities can share the costs and benefits of grid modernization can help improve public acceptance that sites aren't just adding system stress.

The bottom line

No single technological breakthrough will solve the challenge. Only with coordinated planning across infrastructure, energy, supply chains and policy can we determine whether AI growth can be sustained.

The winners will be those who secure not just sites and megawatts, but can orchestrate the full stack needed to turn capacity into reliable AI infrastructure.



Read the full report

Learn more about data center growth scenarios, critical stressors and actions industries can take to align AI ambition with infrastructure reality.

[READ THE REPORT HERE](#)

