GLOBAL INFRASTRUCTURE BRIEFING · JULY 2026

Inside the World's Data Centers

Ownership, scale, energy, geography, and the trillion-dollar infrastructure race reshaping the physical backbone of the internet and AI.

~$430–540B
Global market revenue, 2026

~11,000+
Data centers worldwide, 174+ countries

~945 TWh
Projected annual electricity use by 2030

~$725B
"Big Four" hyperscaler capex, 2026

~200 GW
Global installed capacity projected by 2030

$500B
Stargate's four-year U.S. infrastructure target

Data centers have quietly become one of the largest capital-construction categories on the planet. What used to be anonymous warehouses full of servers are now gigawatt-scale industrial campuses, financed like power plants, sited like factories, and watched as closely as any commodity market. This briefing pulls together the latest available statistics — as of mid-2026 — on who owns the world's data centers, where they sit, how big they are, how much energy and water they consume, and where the money is flowing next.

SECTION 01

The Market at a Glance

Estimates of the global data center market vary by methodology, but they converge on a similar order of magnitude and trajectory. Most industry research firms place 2026 global data center market revenue somewhere between **$425 billion and $540 billion**, up from roughly $270–390 billion in 2025, with compound annual growth rates in the 9–11% range projected out to the early-to-mid 2030s. On that trajectory, the market is commonly projected to reach **$900 billion to $1.1 trillion by 2033–2035**.

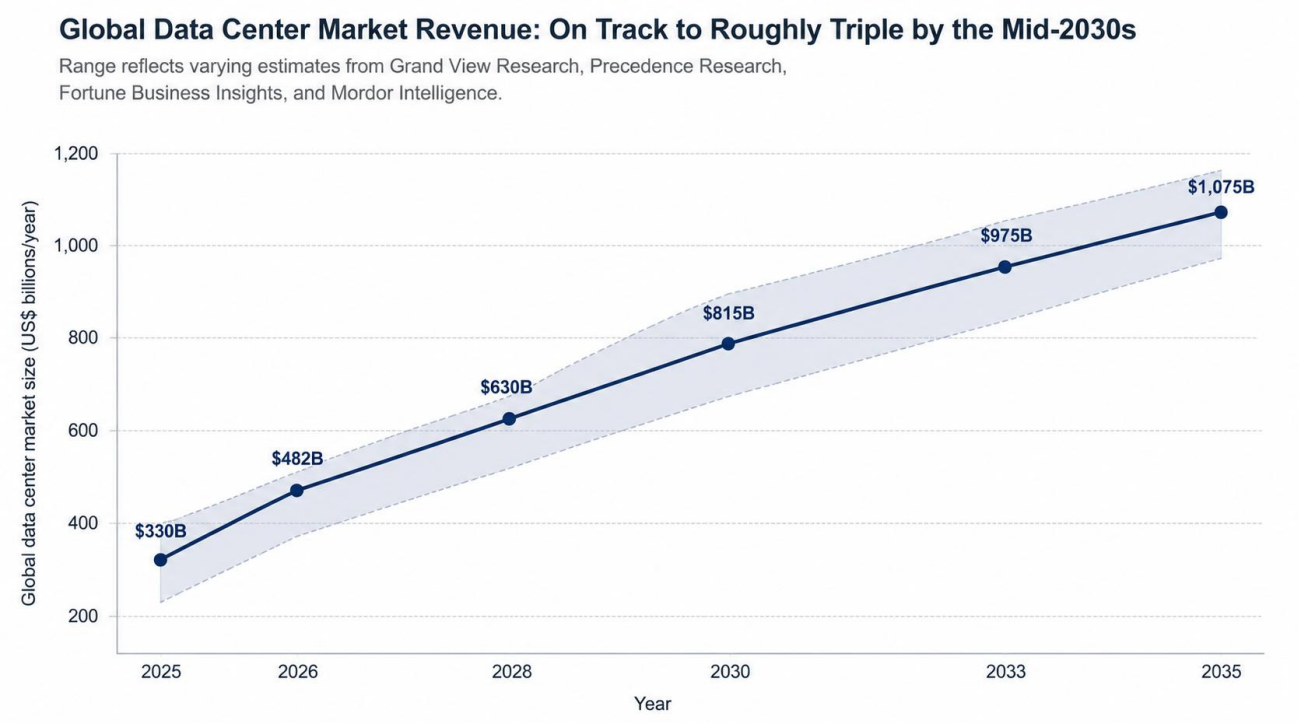


Figure 1. Estimate range across major research firms (Grand View Research, Precedence Research, Fortune Business Insights, Mordor Intelligence). Midpoints shown; real figures will depend on how fast the AI build-out actually proceeds.

Behind that revenue figure sits a much larger physical build-out. Real estate advisory firm JLL describes the sector as being in an "infrastructure investment supercycle" requiring as much as **$3 trillion by 2030** for data center real estate alone, with roughly 100 GW of new capacity expected to come online between 2026 and 2030 — effectively doubling global capacity to around **200 GW**. Tenants are then expected to spend a further $1–2 trillion fitting that space out with servers, chips, and networking gear. Broader estimates that include compute hardware and power infrastructure put total capital spending as high as **$7 trillion-plus by 2030–2031**.

By facility count, data center trackers such as Data Center Map list roughly **11,000–12,000 data centers spread across 174–179 countries**, though these figures capture only the facilities operators choose to register and undercount both very small edge sites and large state-run or military facilities, particularly in China.

SECTION 02

Who Owns the World's Data Centers

Ownership breaks down into a handful of overlapping categories:

The hyperscalers

Amazon (AWS), Microsoft (Azure), Google (Cloud), and Meta together with Oracle now drive the majority of new capacity, and are increasingly building and operating their own campuses rather than leasing. Industry analysts estimate hyperscalers' self-built share already accounts for roughly half of global data center IT capacity among the largest operators, and project that hyperscalers will control around two-thirds of global capacity by 2031 — up from a market historically dominated by enterprise-owned, on-premises infrastructure.

Colocation and REIT giants

These companies don't run cloud services themselves; they build and lease data center space to hyperscalers, enterprises, and each other. Equinix is widely regarded as the world's largest data center REIT, with a global footprint that different counts place between roughly 260 and 270-plus facilities, spanning 70 to 77 markets across six continents. Digital Realty operates a similarly vast footprint — north of 300 facilities spread across more than 50 metro areas in 25-plus countries — and carries an enterprise value above $80 billion. NTT Global Data Centers ranks third-largest globally, with operations in 20-plus countries and a $10 billion expansion program running through fiscal year 2027. Other major names include CyrusOne, Iron Mountain, and DigitalBridge, the latter overseeing a diversified portfolio spanning data centers, cell towers, and fiber networks.

AI-native infrastructure builders

A newer category of owner has emerged directly out of the AI race: OpenAI (through the Stargate program), xAI (through its Colossus complex), and specialist developers such as Crusoe and CoreWeave that build and sometimes co-own capacity on a hyperscaler's or AI lab's behalf.

Financial owners

Private equity, infrastructure funds, and pension funds have become major landlords of digital infrastructure. Blackstone bought QTS Data Centers for $10 billion in 2021; that single deal has since grown into a data center and digital infrastructure business the firm now values at more than $70 billion, with over $100 billion of further development in the pipeline. It is now preparing to list a dedicated data center investment vehicle, Blackstone Digital Infrastructure Trust, and has been courting sovereign wealth funds as anchor investors. BlackRock entered the sector via its 2024 acquisition of Global Infrastructure Partners and has since been acquiring power utilities to secure generation capacity. Brookfield, DigitalBridge, Apollo, KKR, Blue Owl Capital, and pension investors such as Canada's CPPIB round out a fast-growing field of infrastructure capital targeting the sector.

China's state and private operators

China Telecom, China Mobile, and China Unicom operate much of the country's state-linked capacity, while Alibaba Cloud, Tencent Cloud, and ByteDance's Volcano Engine lead on the private/commercial side.

~2/3
Share of global data center capacity hyperscalers are projected to control by 2031

$70B+
Value of Blackstone's data center portfolio, built from a $10B QTS acquisition in 2021

300+
Facilities each operated by Equinix and Digital Realty, the two largest colocation REITs

SECTION 03

On Land, Underwater, and Everywhere in Between

The overwhelming majority of data center capacity — well over 99% — still sits on land, in everything from purpose-built concrete campuses to repurposed factories. Speed has pushed some operators toward unconventional shells: xAI's original Colossus cluster in Memphis occupies a former Electrolux appliance factory, while part of the Stargate program is converting a shuttered auto plant in Lordstown, Ohio into a modular data center fabrication site.

> *"Underwater data centers are moving from experiment to commercial reality — almost entirely in China." — seawater cooling avoids the freshwater draw that land-based evaporative cooling requires.*

China operates the world's first underwater data center to reach commercial service, positioned off Lingshui in the southern island province of Hainan. Developed by Shenzhen HiCloud Data Center Technology together with marine engineering firm Highlander, it lowers sealed, pressure-resistant steel cabins — each weighing roughly 1,300 tonnes and holding 24 server racks with 400–500 servers — onto the seabed about 35 meters down, where surrounding seawater removes heat that would otherwise require energy-intensive chillers on land. Operators report cooling energy savings approaching 90%, translating into an overall efficiency edge of more than 30% versus a similarly sized facility on land. Hainan's development plan calls for scaling the site to as many as 100 underwater cabins over five years, and the facility has already been expanded to take on AI workloads.

A second, newer Chinese project pushes the concept further. In May 2026, HiCloud brought online what it describes as the world's first wind-powered underwater data center, submerged roughly ten meters off Shanghai's Lin-gang coast within a cluster of more than fifty offshore wind turbines. Built to hold around 2,000 servers at up to 24 MW capacity, the roughly $228 million facility draws an estimated 95% of its electricity from that adjacent wind farm and is projected to spare an estimated 26,000 tonnes of freshwater a year that conventional cooling would otherwise require.

Western interest in underwater computing dates back further but hasn't reached commercial scale. Microsoft's own pilot, Project Natick, concluded in 2020 after validating the underlying concept, and two American start-ups, Subsea Cloud and NetworkOcean, have pursued underwater deployments since, though NetworkOcean reportedly ran into permitting obstacles attempting a pilot in San Francisco Bay. For now, seawater-cooled computing at commercial scale remains a distinctly Chinese phenomenon, driven in part by a desire to conserve the freshwater that land-based cooling in water-stressed provinces would otherwise consume.

SECTION 04

How Big Do They Get? Scale and Footprint

A "typical" data center covers roughly 100,000 square feet. Hyperscale AI campuses have blown past that baseline by one or two orders of magnitude. Research group Epoch AI notes that the largest AI campuses under construction are approaching roughly a fifth the footprint of Manhattan, with OpenAI's flagship Stargate site in Abilene, Texas expected to match the physical footprint of Central Park once its full set of buildings is complete.

Table 1. Selected AI-infrastructure megaprojects illustrate the scale now considered normal in the industry.

PROJECT	OWNER / DEVELOPER	LOCATION	REPORTED SCALE
PORTS Technology Campus	SB Energy (SoftBank)	Piketon, Ohio, US	Up to 10 GW planned; ~2,700 acres; ~$67B program
Hyperion	Meta	Louisiana, US	Up to 5 GW planned; ~2,250 acres; ~4M sq ft; ~$27B JV
Stargate Abilene	OpenAI / SoftBank / Oracle (built by Crusoe)	Abilene, Texas, US	~1.2 GW; ~4M sq ft; ~875 acres
Prometheus	Meta	New Albany, Ohio, US	~1 GW target — "world's first gigawatt-capable data center"; five-building supercluster
Colossus 1 & 2	xAI	Memphis, Tennessee, US	~2 GW combined; 555,000 GPUs; ~$18B
Project Rainier	Amazon	New Carlisle, Indiana, US	~910 MW–1 GW; 1,000+ acres; built for Anthropic's workloads
China Telecom–Inner Mongolia	China Telecom	Hohhot, Inner Mongolia, CN	~600 MW planned; 1,000+ acres; cold climate + wind power
Citadel Campus	Switch	United States	~650 MW — largest currently operational single data center

Two things stand out. First, the unit of comparison has shifted from square footage to gigawatts — power capacity, not floor space, is now the headline number developers compete over. Second, several of the very largest projects (Hyperion, PORTS, Prometheus) are not single buildings but multi-building "campuses" or "superclusters" built in phases, with tent-structure and prefabricated modules increasingly favored over conventional construction to compress build timelines from years to months.

SECTION 05

Power Hungry: What the Numbers Say About Energy Use

The most authoritative running estimate comes from the International Energy Agency, which put global data center electricity consumption at roughly **415 terawatt-hours (TWh) in 2024** — about 1.5% of global electricity use — following half a decade in which that demand had been compounding at close to 12% annually. The pace then jumped further in 2025: data center power demand climbed 17% for the year, well outpacing the roughly 3% growth seen across global electricity demand as a whole, with AI-specialized facilities expanding fastest of all.

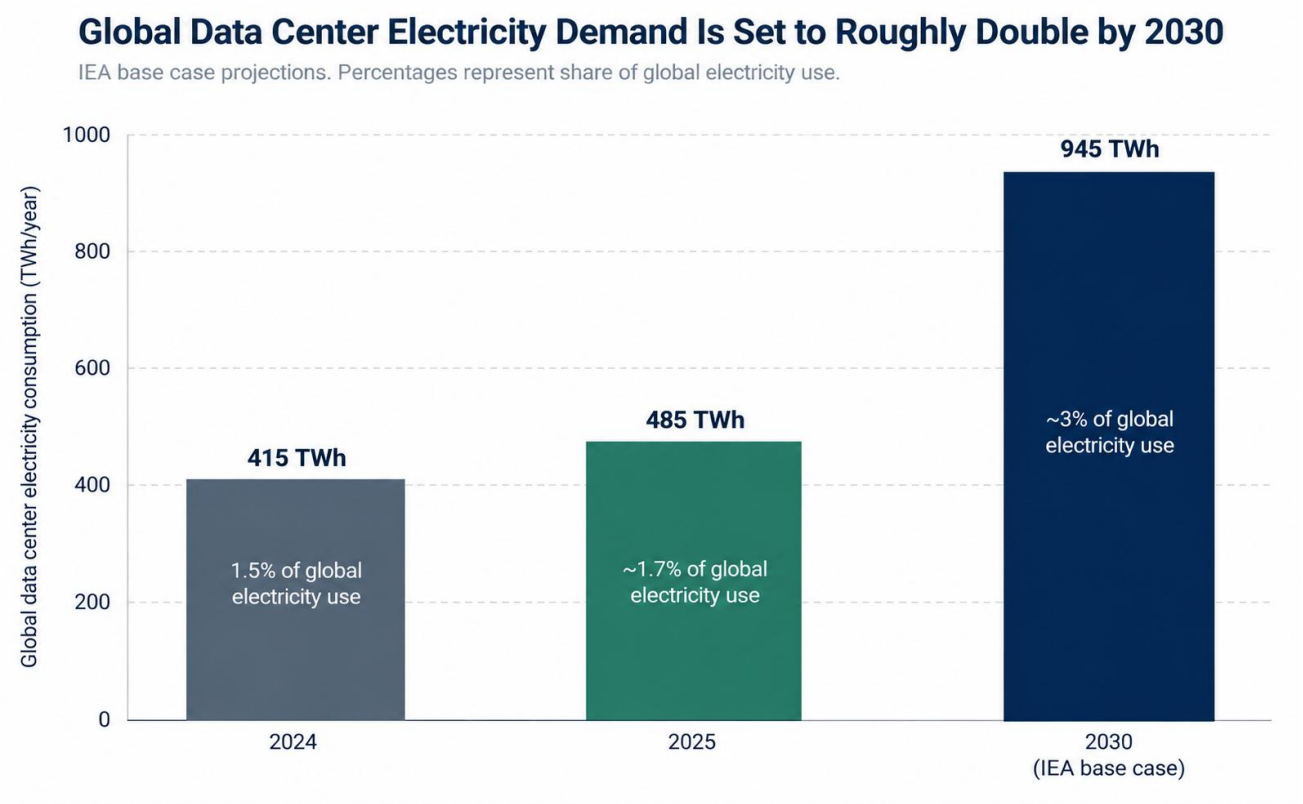


Figure 2. IEA base-case projection. AI-optimized servers — until recently a minority share of data center load — are expected to account for most of the additional growth through 2030.

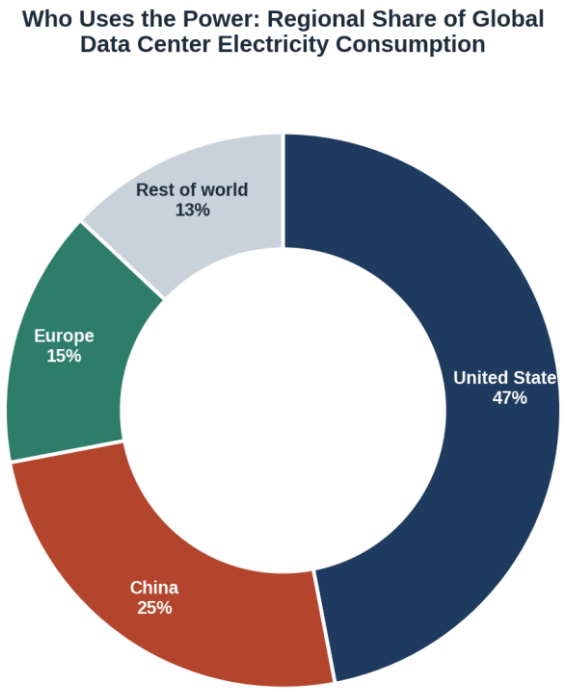


Figure 3. IEA-based regional breakdown (Carbon Brief analysis); shares are approximate and shift year to year.

Demand is concentrated at the national level too. By 2023, data centers were already drawing an estimated 4.4% of total U.S. electricity; Lawrence Berkeley National Laboratory's own modeling puts that figure on a path toward somewhere between 6.7% and 12% by 2028. A handful of hub markets are far more exposed: data centers already account for more than a fifth of Ireland's national electricity consumption, a share the IEA has suggested could approach 32%. By contrast, the International Data Center Authority's 2026 Global Energy Report notes data centers represent only about 0.5% of Middle Eastern electricity consumption today — a gap regional planners are actively trying to close.

Water tells a related story. Direct, on-site cooling water use by U.S. data centers is estimated at around 17.4 billion gallons in 2023, projected to reach 38–73 billion gallons by 2028; because most U.S. electricity still comes from thermal power plants that themselves consume large volumes of water, the indirect water footprint of running those data centers — from power generation alone — is estimated at roughly ten times the direct figure. Among hyperscalers that disclose water metrics, Google reported consuming about 10.9 billion gallons globally in 2025 (up 34% year-over-year), while Amazon and Microsoft report water-use-effectiveness ratios of roughly 0.12 and 0.30 liters per kilowatt-hour of electricity, respectively — figures improving on a per-unit basis even as absolute consumption keeps climbing with overall growth.

SECTION 06

Where They Are: Geographic Distribution

By raw facility count, the United States remains the clear global leader, though exactly how far ahead depends on which tracking database is used — estimates for 2026 range from roughly 4,400 to more than 5,400 U.S. facilities, representing somewhere between 40% and 46% of the global total. Northern Virginia remains the single largest data center cluster on Earth and continues to add more new capacity than any other market worldwide — over 1,100 MW in a single recent year, according to real estate firm CBRE — followed within the U.S. by Texas (led by the fast-growing Dallas–Fort Worth market), Atlanta, Chicago, and Phoenix.

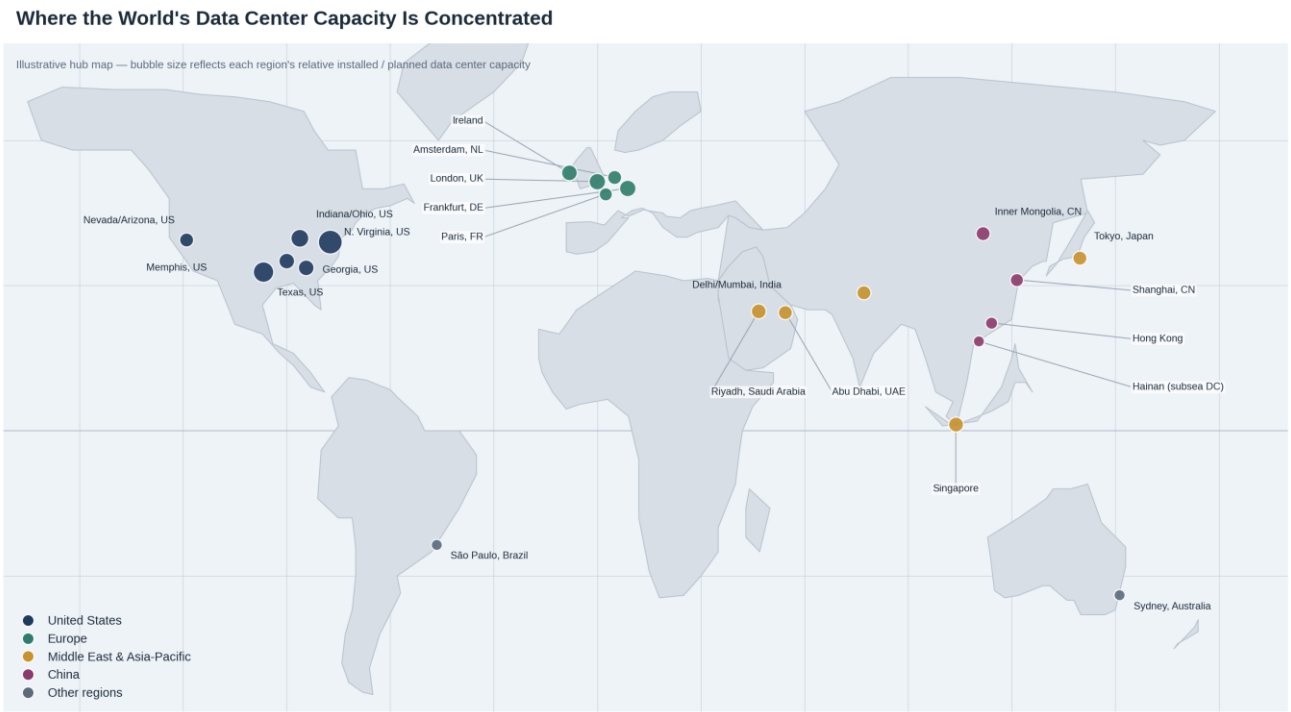


Figure 4. Major data center hub regions by approximate installed and planned capacity. This is a simplified, illustrative hub map rather than a precise geographic reference — bubble size reflects each region's relative scale, not an exact facility count.

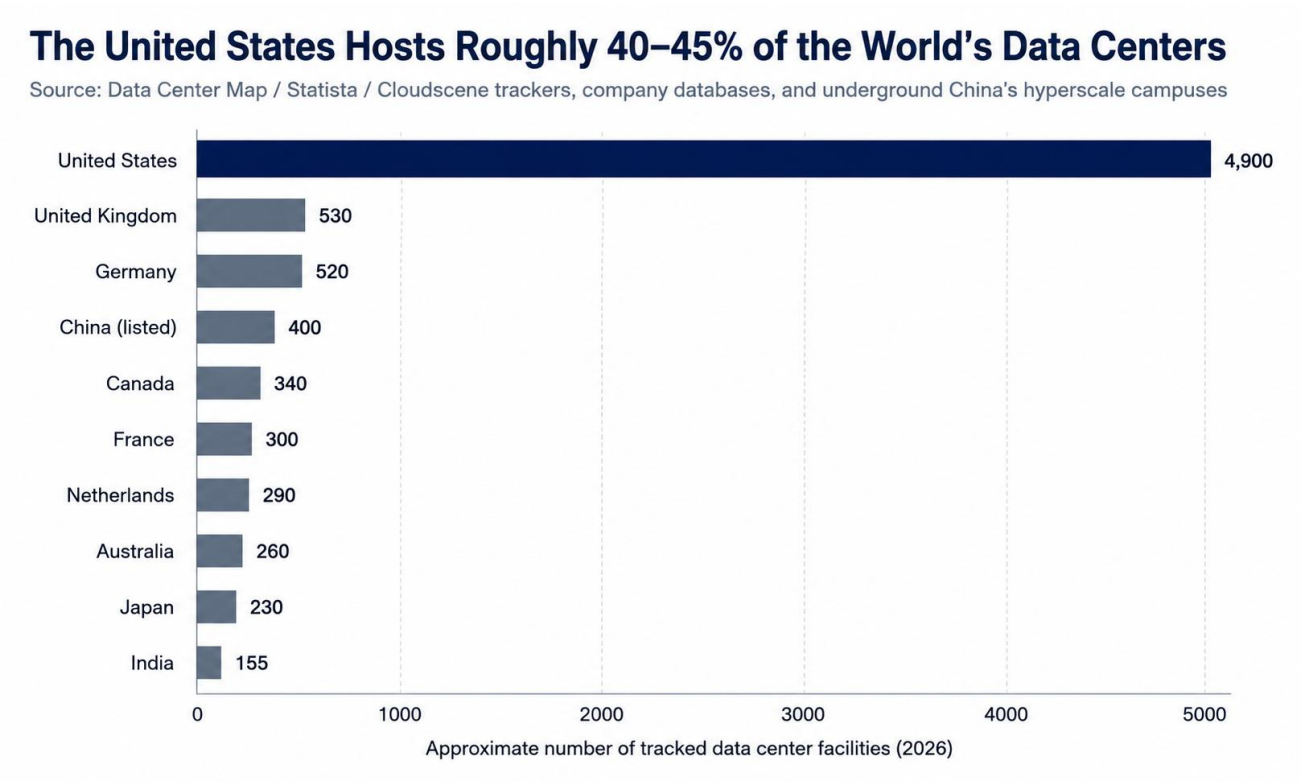


Figure 5. Approximate facility counts by tracking database (Data Center Map, Statista, Cloudscene). Counts vary by source and structurally undercount China, where capacity is concentrated in fewer, much larger hyperscale and state-run campuses.

Outside the U.S., the United Kingdom (roughly 510–555 facilities, concentrated around London) and Germany (roughly 510–530, concentrated around Frankfurt) rank as the next-largest markets by site count, followed by Canada. China's officially tracked facility count looks comparatively small, but this understates its real footprint: China's installed data center capacity is estimated at around 40 GW, delivered through a smaller number of very large state-linked and hyperscale campuses rather than many small commercial sites, plus a national buildout backed by roughly $295 billion (2 trillion yuan) in public investment between 2026 and 2030.

The fastest-growing new frontier is the Middle East. Saudi Arabia and the United Arab Emirates are both racing to become regional AI-compute hubs, backed explicitly by sovereign strategy: Saudi Arabia's Public Investment Fund has stood up HUMAIN as a sovereign AI and data center platform targeting 6 GW of capacity by 2034, while the UAE hosts Stargate UAE — OpenAI's first data center deployment outside the United States — alongside a $15.2 billion capacity expansion led by Microsoft and Abu Dhabi's G42. Regional analysts project Middle East data center construction growing at a compound rate above 25% a year through the early 2030s. Asia-Pacific overall is projected to post the fastest regional growth rate worldwide, driven by data-localization rules and national "sovereign AI" ambitions in markets from India to Japan.

SECTION 07

The Latest Trends Reshaping the Industry

Power, not chips, is now the binding constraint

Until recently, the tightest constraint on AI infrastructure was simply getting hold of enough GPUs. By 2026 that bottleneck has moved decisively downstream: securing a grid connection is now the harder problem, with interconnection approvals in major markets commonly stretching beyond two years and, in the most congested queues, taking five to seven. The U.S. Department of Energy estimates the country will need roughly 100 GW of new peak generating capacity by 2030, about half of it driven by data centers, and Gartner projects that some 40% of AI data centers will be power-constrained by 2027. The practical response has been a wave of behind-the-meter power: on-site gas turbines, battery storage (including large-scale deployments of Tesla Megapacks), and direct plant co-location deals that bypass the public grid entirely.

Liquid cooling has gone from niche to default for AI workloads

Air-cooled facilities typically operate at a Power Usage Effectiveness (PUE) of 1.4–1.8; direct-to-chip liquid cooling brings that down to roughly 1.05–1.15, and full immersion cooling can reach 1.02–1.08. By Goldman Sachs' reckoning, roughly three-quarters of the AI servers shipped in 2026 will run on liquid rather than air cooling, a shift driven by GPU racks that now commonly draw 60–100-plus kilowatts — well beyond what air can dissipate. New efficiency metrics — Power Compute Effectiveness and Water Usage Effectiveness — are emerging alongside PUE to capture how much useful AI computation a facility gets per unit of energy and water.

Nuclear power has moved from pilot programs to signed contracts

Across a string of deals with TerraPower, Oklo, Vistra, and Constellation Energy, Meta alone has lined up rights to as much as 6.6 GW of nuclear generating capacity — more than any other tech company has committed to so far. Microsoft signed a roughly $16 billion, 20-year agreement with Constellation to restart Three Mile Island's Unit 1 reactor, targeting 2027. Google has agreements covering roughly 500 MW with Kairos Power and up to 1,800 MW across three sites with Elementl Power. Amazon has a 1,920 MW power purchase agreement with Talen Energy tied to the Susquehanna nuclear plant plus an equity stake in small-modular-reactor developer X-energy. Trackers following the space count roughly a dozen major hyperscaler-nuclear deals totaling near 10 GW of committed capacity, with collective pledges surpassing $50 billion by early 2026.

Renewable power-purchase agreements have jumped to gigawatt scale

Corporate clean-power contracts mostly ran well under 500 MW as recently as a couple of years ago; Google's February 2026 agreement with TotalEnergies, covering 1 GW of new Texas solar capacity, is widely cited as the deal that reset industry expectations toward gigawatt-scale procurement. Rising demand has also pushed prices higher: North American solar and wind PPA prices rose roughly 13% and 24% respectively over the year to early 2026, even as buyers increasingly pursue firm, round-the-clock "24/7 carbon-free energy" rather than annual-average offsets.

Edge computing and modular, speed-first construction round out the picture

Rather than routing everything through centralized hubs, operators are also pushing smaller facilities out closer to end users to cut latency for applications such as autonomous vehicles and real-time AI inference — a decentralization trend some researchers expect to be worth upwards of $300 billion in its own right by late 2026. Meanwhile, the largest hyperscale campuses are increasingly built from prefabricated tent-style structures and factory-made power/cooling modules, compressing construction timelines that once took a year and a half or more down to a matter of months.

SECTION 08

The Investment Supercycle: What's Being Spent, and Where

The center of gravity is the "Big Four" U.S. hyperscalers. By mid-2026, Amazon, Microsoft, Alphabet (Google), and Meta had collectively guided their 2026 capital expenditure to roughly **$725 billion — up about 77% from approximately $410 billion in 2025.**

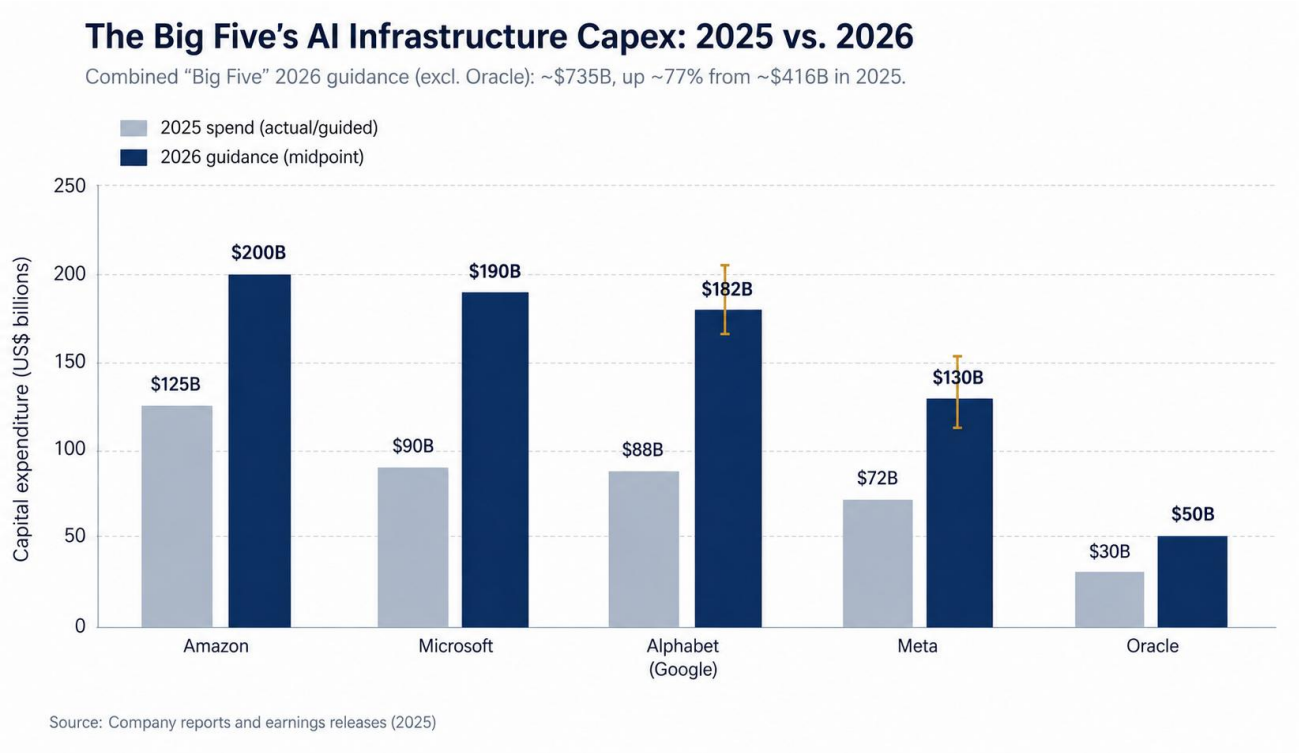


Figure 6. Gold bars mark the disclosed 2026 guidance range for each company; guidance rose over the course of the year as reported in quarterly earnings calls.

Add Oracle, guiding to roughly $50 billion of its own AI-infrastructure spending, and the top five U.S. providers' combined 2026 outlay approaches **$700–770 billion** — a figure that, as several analysts have pointed out, now exceeds the GDP of many mid-sized countries.

Beyond the hyperscalers' own balance sheets, the highest-profile dedicated infrastructure vehicle is **Stargate**, the joint venture announced by OpenAI, SoftBank, and Oracle in January 2025 with the UAE's MGX fund as a founding equity partner, targeting **$500 billion in U.S. AI infrastructure investment and 10 GW of capacity over four years**. By late 2025 the venture had already grown to nearly 7 GW of planned capacity, backed by upwards of $400 billion pledged for the following three years, with new sites added in Michigan and Ohio alongside the Abilene flagship and an international arm launched through Stargate UAE.

China is pursuing a parallel but distinct build-out. Its National Development and Reform Commission has earmarked roughly **$295 billion (2 trillion yuan) in public funding** for a nationwide AI-computing buildout between 2026 and 2030, built largely around state operators China Mobile and China Telecom and reliant on domestic chipmakers such as Huawei. On top of that public figure, China's leading private technology companies are investing heavily in parallel: Alibaba has committed roughly $52–69 billion to cloud and AI infrastructure over three years, Tencent has put a multi-year figure of roughly $70 billion (500 billion yuan) behind its own infrastructure build-out, on top of separate plans to double its dedicated 2026 AI spending versus the year before, and ByteDance is targeting more than 1.2 GW of new AI data center capacity with roughly $23 billion in 2026 capital spending. Combined, China's four largest internet companies — Alibaba, Tencent, ByteDance, and Baidu — are reported to be planning around $84 billion in AI infrastructure investment through 2027.

In the Middle East, Saudi Arabia's pipeline of publicly and privately backed AI-infrastructure commitments is estimated above $20 billion as of early 2026, anchored by the PIF-backed HUMAIN platform, telecom-linked operators Center3, Mobily, and DataVolt, and a separate $5.3 billion Amazon Web Services cloud-region investment tied to the kingdom's Vision 2030 program. The UAE's parallel push includes the Microsoft–G42 expansion and Stargate UAE.

SECTION 09

Toward Efficiency: How the Industry Plans to Do More With Less

Faced with power, water, and land constraints arriving faster than new grid capacity can be built, operators are converging on a common efficiency playbook:

1.02–1.15
Typical PUE for liquid or immersion-cooled facilities, vs. 1.4–1.8 for air cooling

~76%
Share of AI servers shipped in 2026 expected to be liquid-cooled (Goldman Sachs)

≤1.3
Maximum PUE Germany will require for new data centers from 2027

- **Switching to liquid and immersion cooling**, which can cut PUE by roughly a third to a half compared with conventional air cooling, while also enabling far higher rack densities on the same physical footprint.
- **Recovering waste heat rather than discarding it.** Some operators are capturing the heat their liquid-cooled racks give off instead of simply venting it: municipal district-heating networks in Nordic cities including Stockholm and Helsinki now draw part of their supply from nearby data centers, turning server exhaust heat into a second, saleable form of energy.
- **Tightening regulation.** Germany's Energy Efficiency Act will require new data centers to achieve a PUE of 1.3 or better starting in 2027 and mandates waste-heat recovery wherever technically and economically feasible — a rule that effectively pushes new German facilities toward liquid cooling by default.
- **Hardware-level redesign.** Nvidia's newest Vera Rubin GPU racks, unveiled in January 2026, tolerate cooling-water inlet temperatures up to 45°C — warm enough that many sites can skip mechanical chillers altogether, a change expected to meaningfully cut both energy and water use in next-generation AI facilities.
- **Seawater and free cooling as structural efficiency plays.** China's underwater facilities report cooling energy savings approaching 90% by using ambient seawater instead of mechanical refrigeration, while conventional land-based sites increasingly rely on "free cooling" — pulling in naturally cold outside air — wherever climate allows, a major reason cold-climate and coastal locations remain attractive for new builds.
- **Chasing genuinely round-the-clock clean power**, not just annual offsets. Hyperscalers are shifting away from simple renewable energy certificates toward power purchase agreements and nuclear contracts intended to match electricity consumption with carbon-free generation hour by hour rather than on an annual average basis.

It's worth noting the limits of this progress: even as per-unit efficiency improves, most hyperscalers' absolute emissions and water use are still rising because growth is outpacing efficiency gains — a tension that has drawn increasing scrutiny from investors and environmental groups alike, and one likely to remain a defining storyline for the industry through the rest of the decade.

SECTION 10

The Money Behind the Machines: Major Capital Holders

Several distinct pools of capital are funding this build-out, each with recognizable names behind them:

The hyperscalers themselves

By far the largest source of capital, self-funding their buildouts from operating cash flow and debt: Amazon (CEO Andy Jassy), Microsoft (CEO Satya Nadella, CFO Amy Hood), Alphabet/Google (CEO Sundar Pichai, with Thomas Kurian leading Google Cloud), Meta (CEO Mark Zuckerberg), and Oracle (chairman and CTO Larry Ellison, a driving force behind the Stargate venture).

AI labs and their financial backers

Now building infrastructure directly rather than only renting it: OpenAI (CEO Sam Altman) co-leads Stargate alongside SoftBank (CEO Masayoshi Son) and Oracle, with the UAE's MGX fund as founding equity partner; xAI, owned by Elon Musk, has self-funded the Colossus complex in Memphis at a pace and scale that has made it one of the most closely watched private AI infrastructure bets in the industry.

Private equity and alternative-asset managers

Have turned data centers into a mainstream institutional asset class: Blackstone (CEO Stephen Schwarzman) via its QTS platform and forthcoming listed data center vehicle; BlackRock (CEO Larry Fink) via its Global Infrastructure Partners arm and a related string of power-utility acquisitions; alongside Brookfield, DigitalBridge, Apollo, KKR, and Blue Owl Capital.

Sovereign wealth funds

Have become anchor investors rather than passive participants: the UAE's MGX fund helped seed Stargate itself; Saudi Arabia's Public Investment Fund backs the kingdom's HUMAIN platform; and Qatar Investment Authority has partnered with Brookfield on data center development in Qatar.

Publicly traded colocation REITs

Equinix, Digital Realty, and NTT foremost among them, give ordinary investors direct, dividend-paying exposure to the sector, and are themselves frequent counterparties and landlords to the hyperscalers above.

In China

Capital is a blend of state and private: the national buildout sits with state operators China Mobile and China Telecom, while private capital is led by Alibaba (CEO Eddie Wu), Tencent (CEO Pony Ma), and ByteDance.

OUTLOOK

Where This Leaves Us

Taken together, the numbers describe an industry still in the early stages of a build-out that its own participants openly compare to the construction of national power grids or interstate highway systems — but compressed into years rather than decades. The defining tensions for the rest of the decade look set to be the same ones already visible in 2026: whether power grids, water supplies, and local communities can absorb growth running at double-digit annual rates; whether the trillion-dollar bet that AI demand will justify this spending pays off; and whether efficiency gains from liquid cooling, nuclear power, and — at the margins — underwater data centers can outrun the sheer scale of new demand being built on top of them.

SELECTED SOURCES

Sources

Market size and capacity: Grand View Research, Fortune Business Insights, Precedence Research, Mordor Intelligence, Statista, JLL Global Data Center Outlook, CBRE Global Data Center Trends, Data Center Map.

Energy and water: International Energy Agency ("Energy and AI," "Electricity 2026," and related IEA analysis), Lawrence Berkeley National Laboratory, Carbon Brief, Our World in Data, Environmental and Energy Study Institute, Global Water Intelligence.

Company and project data: OpenAI, SoftBank Group, company SEC filings (Digital Realty, Blackstone), DatacenterDynamics, Data Center Knowledge, Epoch AI, SemiAnalysis, Wikipedia (Colossus data center).

Investment and markets: financial and trade press including CNBC, Yahoo Finance, Tom's Hardware, and industry newsletters (Futurum Group, ValueAddVC, Orennia).

Regional coverage: Arab News, South China Morning Post, People's Daily, Al Jazeera, and regional data-center trade publications.

Figures reflect publicly available estimates as of July 2026; because data centers are being built, financed, and re-priced on a near-weekly basis, exact numbers — especially capital expenditure guidance and in-progress project specifications — should be treated as directional rather than final. The world map in Figure 4 is a simplified, illustrative hub map, not a precision cartographic reference.