AI MEETS THE PHYSICAL WORLD

ESG RESEARCH– DATA CENTERS (1/2)

AI investment at the largest hyperscalers[(a)] could increase eightfold by 2028 [1]. This surge is pulling attention toward physical infrastructure, and toward the environmental and social issues of building it out. Whether infrastructure can keep up with demand is only half the question. Data centers draw on energy, water, and land at a scale that carries real financial risks: **grid connection delays, suspended projects, local moratoriums.** For a responsible investor, analyzing these issues through a double-materiality lens is essential to support the build-out of sustainable AI infrastructure.

Risks run the length of the value chain, with semiconductor manufacturing and critical minerals upstream, and electronic waste and recyclability downstream. This study focuses on the infrastructure itself, because that is where the gap opens between capacity announced and capacity delivered. The International Energy Agency (IEA) estimates that **one in five projects may never be built within five years, held up by grid connection alone** [2].

ENERGY: THE GRID IS THE FIRST BOTTLENECK

A data center, first of all, a massive load on the grid. A new hyperscale campus can draw more than 100 MW [3], hence the bottleneck sits at grid level. Installed global data center capacity stood at 102 GW in 2025 and could double by 2030. The sector's electricity consumption could rise from 1.5% to 3% of global electricity, the equivalent of Japan's consumption today [4].

Legacy hubs are already saturating. Connection now takes 2 to 6 years, against 12 to 18 months previously, and up to 7 to 10 years in the densest markets such as Frankfurt. Northern Virginia's Data Center Alley has a queue of more than 30 GW, Amsterdam has suspended all development permits since 2019 and Dublin did so until 2025, where data centers were absorbing 22% of national electricity [2].

Two consequences follow:

- Expansion is shifting to where power is cheap, accessible and preferably low carbon, such as in the Nordic countries, Spain, or Texas ;

- Some operators are building their own generation to get off the grid entirely, behind-the-meter. Most onsite solutions sit in the United States, where gas is runs through simple or combined cycle turbines. The behind-the-meter option, although effective for developers, draws local over the nuisances it creates. It also strains the 100% renewable commitments hyperscaler advertises, commitment that rest mostly on buying low-carbon certificates across the year and not on matching supply in real time.

Estimated equipment lead times and grid connection times

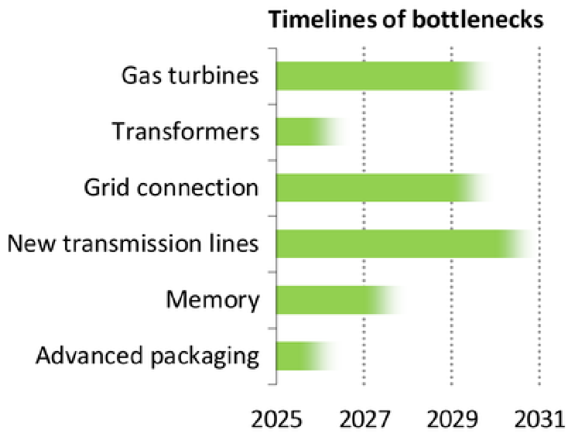

Source: IEA

LAND, NUISANCE AND LOCAL BENEFITS: ACCEPTABILITY IS NOW A MATERIAL RISK

US opinion has swung fast, and the shift is well documented. More than 7 in 10 Americans now oppose data center near their home, against barely more than 4 in 10 in the fall of 2025, and a majority say they would back a national moratorium [5]. In Q1 2026, opposition groups blocked or delayed some 75 projects worth more than $130 billion [6].

Among the reasons for rejection:

- Land. A modern hyperscale campus takes up more than 40 hectares (roughly 100 acres) [7] and has to sit close to electrical infrastructure. That siting constraint can drive the conversion of farmland and depress the value of neighboring residential property.

- Noise. Ventilation produces constant low-frequency vibration. Without acoustic treatment, the loudest equipment reaches 100 dBA at source [8].

- Local jobs. The number is frequently overstated. Most of the work is in construction, and modular prefabrication is moving even that off-site. Operating headcount is small and specialized: Equinix's PA13 site in Meudon, near Paris, runs 7,000 m² of data halls with about 30 permanent staff [9].

- The rising cost of tax exemptions for States. In the United States, these incentives were introduced to attract a nascent industry, and their cost have risen sharply with AI. In North Carolina, estimated annual forgone revenue went from $4 million in 2015 to $45–57 million today [10]. In 2026, 28 of the 38 states with such fiscal scheme filed amendments, and at least 9 considered repealing them outright [11]. The conditions attached to these exemptions could tighten, requiring developers to bear more the true costs associated with their projects, including grid connections, low-carbon energy procurement, water-use disclosure, and stricter noise and air pollution standards.

Spatial concentration and urban proximity of various facilities

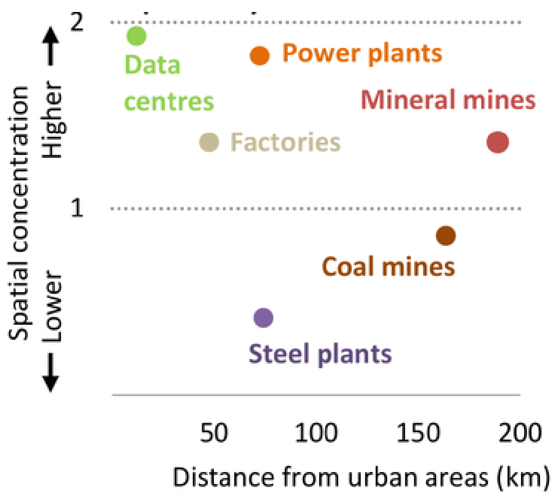

Source: IEA

Data centers combine heavy geographic clustering with siting close to populated areas.

FROM AIR COOLING TO LIQUID COOLING: WATER AS A PERMITTING ISSUE

While water is a second-order condition in siting decisions, it is one of the leading causes of local opposition. Data center cooling relies on two loops. An internal loop removes heat from the servers (using air, water or a refrigerant). An external loop rejects that heat to the outside air, either dry or by evaporation.

The central point of our review of cooling systems is that: how much water a data center consumes depends more on this method of rejecting heat to the outside (external loop) than on the technology used to cool the servers (internal loop).

Each loop raises a specific issue:

- A volume issue (external loop): most of a data center's indirect water footprint sits off-site. In 2023, about two-thirds of data center water consumption came from energy supply and electricity generation, against a quarter for cooling (see chart). On site, this creates a water-energy trade-off: evaporative heat rejection (cooling towers, adiabatic systems) is energy-efficient but consumes water, whereas a combination of free cooling (using outside air) and dry coolers [b] saves water but uses more energy. The choice of heat rejection method therefore depends on the water intensity of the regional power mix and on the level of local water stress. Notably, 43% of the world's data centers sit in areas of high or extreme water stress [12].
- A fluid management issue (internal loop), following the shift from air cooling to water cooling. Air works up to a density of roughly 15 kW per rack[d], whereas NVIDIA's latest chips reach 120–130 kW (Blackwell generation), rising to an announced 600 kW for the Rubin Ultra generation [13]. Liquid cooling [e] is now a default design assumption, and the market could grow around 22% a year [14]. But these technologies do not carry the same water pollution risks: blowdown water loaded with minerals and treatment chemicals in most liquid systems, through to the use of PFAS[f] in some liquid solutions that are less deployed. Beyond the adoption rate of liquid cooling, what matters is the nature of the fluids and the way operators manage discharges.

Water withdrawals and consumption by data centres in IEA's Base Case, 2023 and 2030

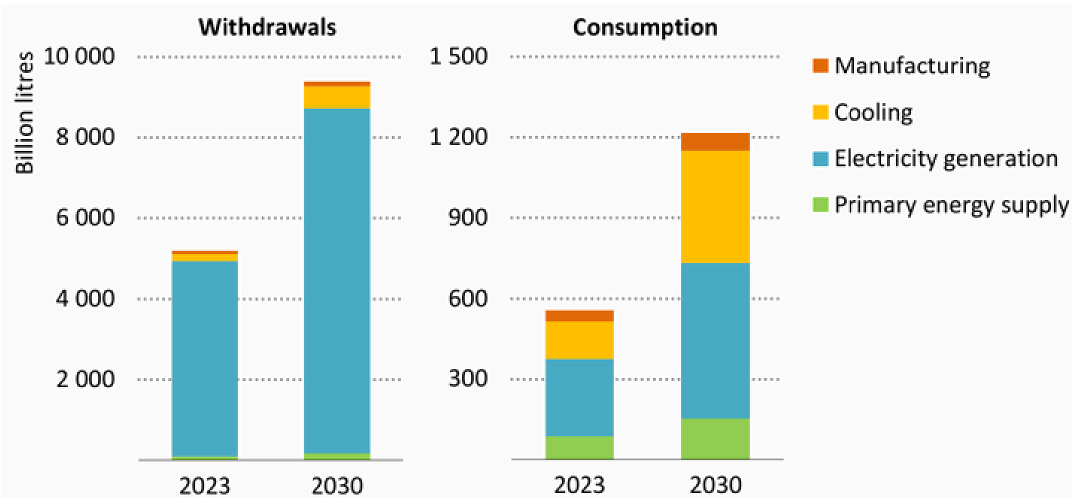

Source: IEA

Reduce and anticipate: the value chain's twin challenge

Energy, local acceptability, water. These constraints decide how much capacity actually gets built, and therefore how much value the sector creates. Other issues remain poorly documented: end-of-life equipment management and recyclability, pollution risks to water, air and soil linked to refrigerants and onsite fossil fuels, or the availability of a skilled workforce to run the new liquid cooling systems and keep pollution risks in check.

Challenge runs two ways for value chain players. Reducing the environmental and social impacts of data centers already running but also addressing these issues in as early as possible, in site selection and technologies choices, so as fewer announced projects end up disappearing.

Investors have a real part to play in meeting these challenges: analyzing the issues, supporting best practice and picking the companies that provide solutions. The second part of our study, to be published, focuses on that.

Contributors: Clémence Bourcet, Anaïs Cassagnes, Erwan Crehalet, Alice Descourtieux et Louis Reinhart

Completed August, 2026. The opinions, estimates or forecasts expressed on market trends reflect our judgment and may change without notice, as may any statements on financial market trends, which are based on current market conditions. Sycomore AM gives no undertaking that they will be realized. References to any particular type of securities and/or issuers are for illustrative purposes only and should not be construed as recommendations to buy or sell those securities. We recommend that you inform yourself carefully before making any investment decision.

NOTES AND SOURCES

[1] Based on announced annual capital expenditure by the largest AI builders (Google, AWS, Microsoft, Meta, Oracle and the neoclouds) in 2024 and forecast for 2028

[2] IEA, Energy and AI, April 2025

[3] Congress.gov, Data Centers and Their Energy Consumption: Frequently Asked Questions, May 2026

[4] IEA, Key Questions on Energy and AI, April 2026

[5] Heatmap and Embold Research, polls, August 2025 to May 2026, June 2026

[6] Data Center Watch, Q1 2026: Data Center Watch Report, 2026

[7] NTT Data, AI Stress Test Report: Can Data Centers Keep Pace with AI?, June 2025

[8] While it is the persistence of the noise that is the greater problem, levels can reach 105 to 110 dBA depending on the source, driven in particular by how well the equipment is acoustically insulated. Environmental and Energy Study Institute, Communities Are Raising Noise Pollution Concerns About Data Centers, March 2026; Data Center Knowledge, What are the 5 Main Causes of Noise in Data Centers?, March 2025

[9] Ville de Demain, Guide des Data Centers : Retombées Socio-Économiques, Fiscalité et Acceptabilité Locale, April 2026

[10] Good Job First, Data Center Tax Breaks Becoming Billion-Dollar Budget Sinkholes, April 2026

[11] Newsweek, Which States Are Pausing Data Center Tax Incentives—And Why, July 2026

[12] Kepler Cheuvreux, AI & water: the opportunity angle, June 2026

[13] NVIDIA Blog, Hotter Than a Hot Tub: The 45°C Breakthrough to Cool AI's Biggest Machines, June 2026

[14] Mordor Intelligence, DATA CENTER LIQUID COOLING MARKET SIZE & SHARE ANALYSIS, January 2026

TECHNICAL NOTES

(a) Hyperscaler: a company that designs, builds and operates its own data centers at very large scale, both for its own needs and for its cloud customers. The main ones are Amazon Web Services (AWS), Google, Microsoft (Azure), Meta and Apple.

(b) Wet cooling tower: water in the external loop, having picked up heat from the internal loop, runs down through a tower with air flowing across it. Part of that water evaporates and carries the heat away, cooling what remains before it returns to the circuit. In an adiabatic system, outside air passes through a wet medium or a fine mist. The water evaporates, which lowers the air temperature and improves its capacity to absorb heat from the internal loop.

(c) Free cooling: uses outside air to remove heat from the internal loop, doing away with the energy-hungry compression chiller and relying on heat exchangers and fans to transfer heat to the outside air whenever the ambient temperature is low enough to allow it. Dry cooler: circulates the liquid through a network of finned tubes swept by forced outside air.

(d) The rack is the cabinet that houses the servers, their power supply and their cabling.

(e) There are two types of liquid cooling. Direct-to-chip captures heat through cold plates, often copper, in direct contact with the chip. Immersion cooling submerges the whole server in a tank. Each can be single-phase (the fluid stays liquid) or two-phase (heat transferred as the liquid turns to gas — via a refrigerant for direct-to-chip, or a low-boiling-point liquid such as certain PFAS for immersion).

(f) Per- and polyfluoroalkyl substances, or PFAS, are chemicals that persist for extremely long periods in the environment and in living organisms