


Data Centers: Compute, the Environment and AI

By Paul Von Steenburg



Few topics have generated as much press of late as the data center buildout. The reporting runs in two directions at once: data centers are the indispensable backbone of the AI revolution or they are an [environmental problem that is draining reservoirs, straining the grid and driving up power bills](#). Having underwritten this sector for years, I hear a version of this from investors regularly, and, as with most things, the reality is more nuanced than either headline.

WHAT IS A DATA CENTER ANYWAY?

It is worth starting with what a data center actually is, because the term covers far more variety than the headlines suggest. At its simplest, a data center is a building full of computers plus the power and cooling to keep them running. For decades, most organizations ran their own small server rooms on-site. That model was inefficient, with those server rooms mostly sitting idle, built for peak demand but rarely used to capacity. Over the past twenty years, computing has migrated out of thousands of these small rooms and into a smaller number of large, specialized facilities run by cloud providers. The data room in the basement became a data center in the cloud.

It is also worth remembering that a data center is downstream of demand, not a source of it. These buildings exist because of what households and businesses already do every day—stream, search, shop, bank, and now prompt AI. That does not make their footprint imaginary. But it does mean the building is an effect of consumption we have collectively chosen, which is worth keeping in mind before treating the facility itself as the problem to be solved.


NOT ALL DATA CENTERS ARE THE SAME

Another point worth understanding, as it is often misunderstood: “data center” describes very different business models. At one end sits the pure landlord who builds the shell and brings power to the site, then leases it to a large technology company that installs its own computers and contracts for the electricity and water. At the other end, a company builds and operates the entire facility itself. In between, there is a range of hybrids.

Ownership of the buildings is dispersed across specialist developers, publicly traded REITs such as Digital Realty and Equinix, [private investment funds](#), and the technology giants themselves. The tenants are far more concentrated. A handful of “hyperscalers”—principally Amazon, Microsoft, Google and Meta—account for an outsized share of demand, alongside a newer group of AI-focused “neocloud” providers. Importantly, even these giants do not own everything they use; a large share of their capacity is leased rather than self-built.

OWNING THE BUILDINGS AND FILLING THEM ARE TWO DIFFERENT THINGS

Numbers in percent (%)



Hyperscalers run nearly half of global capacity but lease a large share of it.


Source: Synergy Research Group (global capacity by operator, Q4 2025); JLL (North America under construction pipeline). Shares approximate.

THE EFFICIENCY STORY

The consolidation of many users’ moving on-premise computing into a smaller number of large facilities had a consequence that runs counter to the headlines. The common assumption is that all this construction simply means more energy use but for most of the last decade, the opposite was true. A large, shared facility pools the demand of many users across the same hardware, so the machines do useful work rather than sit idle. The result was a striking decoupling: between 2010 and 2018, the computing done in the world’s data centers rose by roughly 550 percent while the electricity they consumed rose just 6 percent. Consolidating computing is, in effect, the equivalent of commuters moving from individual cars to mass transit. The footprint became more visible because it was concentrated in large buildings, but the energy required to perform a given task kept falling.

The industry measures facility efficiency with a simple ratio called Power Usage Effectiveness, or PUE—the total energy a data center draws divided by the energy its computers actually use. A value of 1.0 would mean every watt reaches the servers; 1.5 means the building consumes half—as much as the IT load itself, the rest going to cooling and power conversion. The typical on-premise server room runs around 1.6, while the hyperscalers run near 1.1. A hyperscale facility wastes less than a tenth of its power on overhead; a typical enterprise site wastes more than a third.

THE EFFICIENCY GAP: HYPERSCALE VS ON-PREMISE



The efficiency gap in facility overhead: hyperscale cloud (~1.1) versus the enterprise average (~1.6).

Source: Uptime Institute 2025; Google and AWS disclosures.

And PUE, if anything, understates the advantage. It captures only facility overhead—cooling and power loss—not how fully the servers themselves are used. The larger efficiency gain comes from utilization: pooled hyperscale hardware runs far busier than the idle servers of a company server room, so the true gap in useful computing per unit of energy is even wider.

The trend, however, is now reversing. AI is far more power-hungry than ordinary web tasks, and it is adding a heavy new layer of demand on top of the steady growth in everyday cloud use. The efficiency dividend that kept energy use flat for a decade has largely been spent.

THE ENVIRONMENTAL DEBATE: TWO QUESTIONS, NOT ONE

Much of the concern around data centers blends water and energy into a single worry. It is more useful to separate them, because they have very different answers.

On water, the concern is more tractable than it appears. The traditional way to cool a data center is to evaporate water, running it past the hot equipment and letting evaporation carry the heat away, which consumes the water and requires continual replacement. A facility cooled this way can draw large volumes from the local supply, which is what is fueling most community opposition. The newer approach is “closed-loop” cooling, in which the cooling fluid is sealed inside and sheds its heat to the outside air, much like a radiator in a car. The same water circulates again and again rather than evaporating away. A building designed this way consumes almost no water for cooling on an ongoing basis, although of course, it must be filled upon its commissioning.


This is not aspirational; the largest operators are already committing to it. Microsoft announced that every data center it has designed since August 2024 uses a “zero-water” evaporation design, eliminating the more than 125 million liters of water a typical facility previously consumed each year. As the company’s engineering lead put it, “once the system is filled during construction, it will continually circulate water between the servers and chillers to dissipate heat without requiring a fresh water supply.” Water, in short, is largely a design choice, and the operators that consume it are moving quickly to design it out. Our own partners confirm that all future projects are being designed with closed-loop cooling.

Energy is the harder problem. At the level of a single facility the industry has become remarkably efficient, wasting little power, particularly when compared to inefficient on-premise facilities. But efficiency per task is no longer the whole story, because the sheer volume of computing is now growing faster than efficiency can offset.

WHAT AI CHANGES

The numbers are stark. U.S. data centers consumed roughly 192 terawatt-hours of electricity in 2024 - about 4.7 percent of the nation's total, already more than double the 76 terawatt-hours they used in 2018. The Department of Energy's national laboratory projects that figure will reach roughly 649 terawatt-hours by 2030, or close to 12 percent of all U.S. electricity, with a scenario range running as high as 843. This is no longer an efficiency story; it is an [infrastructure-investment story](#).

U.S. DATA-CENTER POWER USE HAS SURGED — AND IS SET TO SURGE FURTHER



Source: Lawrence Berkeley National Laboratory, 2025 Update: U.S. data-center electricity use roughly tripled over 2018–24 and is projected to more than triple again by 2030.

The binding constraint is not only how much electricity is needed, but how quickly it can be delivered. New supply depends on transformers, gas turbines and grid interconnections, all of which carry long lead times and multi-year queues. In many markets, getting power to a site is now the gating item. Meeting this demand will require substantial capital in both generation and transmission, and it has revived interest in firm, around-the-clock power: [nuclear restarts](#) and [next-generation reactors](#) alongside natural gas and renewables.

All of this—the generation, the transmission, the grid upgrades have to be paid for, which raises a question as politically charged as any in the debate: who pays the bill?

WHO PAYS?


Here the trend is moving in a constructive direction: increasingly, the cost of new infrastructure is being placed on the large users driving the demand rather than on ordinary households. Virginia—home to the world's densest concentration of data centers—offers the clearest template. In late 2025, its State Corporation Commission approved a new rate class for Dominion Energy, effective January 2027, that places data centers in their own category and requires them to commit to long-term contracts and to pay for the bulk of the transmission and generation capacity they reserve, whether or not they use it. The state has since gone further, enacting the nation's first statewide consumption tax on data-center energy use and ordering that data centers bear the full cost of transmission lines built specifically to serve them. The large user, not the household, is increasingly being made to pay for the infrastructure its demand requires, and it is a framework now taking shape in other states and municipalities.

The tax incentives themselves are also drawing public outcry. Yet it is worth remembering that those exemptions are what drew these facilities to Virginia rather than Texas or Arizona in the first place; without them, the local tax base simply would not exist. In Loudoun County, which has the densest data-center market in the world, data centers now fund roughly 40 percent of the county's general fund, which has allowed Loudoun to cut its residential property tax rate nearly every year for the past decade, to the lowest in Northern Virginia. Absent that revenue, the county's own budget arithmetic implies the residential rate would need to rise by roughly 90 percent to hold services constant.

BEYOND THE HEADLINES

It is worth noting that these large consumers of energy are not passive consumers of dirty power; they have become the largest buyers of clean energy on the planet. In 2025, Amazon, Microsoft, Google and Meta together accounted for roughly 49 percent of all corporate clean-power purchase agreements signed globally – and nearly 87 percent in the United States. The two single largest corporate clean-energy buyers in the world last year were both data-center operators: Meta and Amazon, each contracting more than 10 gigawatts, with nuclear representing approximately a quarter of their activity as they reach for firm, around-the-clock power. This does not erase the problem that demand is currently outrunning clean supply, and their emissions have been rising as a result; but it is a very different picture from the one the headlines paint.



HYPERSCALERS HAVE BECOME THE WORLD'S LARGEST BUYERS OF CLEAN POWER



Source: BloombergNEF, 2025: The hyperscalers that fill data centers are now the largest corporate buyers of clean power in the world.

The grid these tenants draw on is itself getting cleaner: over the past decade U.S. electricity generation has grown even as the mix has shifted toward renewables and away from coal.

MORE ELECTRICITY OVERALL — AND A CLEANER MIX



Source: U.S. Energy Information Administration (EIA): U.S. electricity generation has risen while the mix has shifted toward renewables and away from coal.

AI has genuinely changed the energy equation, and the sector's rising power demand is a real challenge that governments, utilities and the technology companies will need to manage together.

DATA CENTER INVESTMENT: WHAT IT MEANS FOR INVESTORS

For investors, the implication follows from these dynamics. As power, permitting, and now the **cost of capital** increasingly ration the buildout, the advantage accrues to **owners of well-located, powered and contracted capacity**, particularly assets on the better end of the efficiency and water spectrum, leased to creditworthy tenants, in the markets best equipped to absorb them.



This page intentionally left blank

Important Notes

Certain information contained herein has been obtained from or is based on third-party sources and, although believed to be reliable, has not been independently verified. Such information is as of the date indicated, if indicated, may not be complete, is subject to change and has not necessarily been updated. No representation or warranty, express or implied, is or will be given by The Common Fund for Nonprofit Organizations, any of its affiliates or any of its or their affiliates, trustees, directors, officers, employees or advisers (collectively referred to herein as "Commonfund") or any other person as to the accuracy or completeness of the information in any third-party materials. Accordingly, Commonfund shall not be liable for any direct, indirect or consequential loss or damage suffered by any person as a result of relying on any statement in, or omission from, such third-party materials, and any such liability is expressly disclaimed.

All rights to the trademarks, copyrights, logos and other intellectual property listed herein belong to their respective owners and the use of such logos hereof does not imply an affiliation with, or endorsement by, the owners of such trademarks, copyrights, logos and other intellectual property.

To the extent views presented forecast market activity, they may be based on many factors in addition to those explicitly stated herein. Forecasts of experts inevitably differ. Views attributed to third-parties are presented to demonstrate the existence of points of view, not as a basis for recommendations or as investment advice. Market and investment views of third-parties presented herein do not necessarily reflect the views of Commonfund, any manager retained by Commonfund to manage any investments for Commonfund (each, a "Manager") or any fund managed by any Commonfund entity (each, a "Fund"). Accordingly, the views presented herein may not be relied upon as an indication of trading intent on behalf of Commonfund, any Manager or any Fund.

Statements concerning Commonfund's views of possible future outcomes in any investment asset class or market, or of possible future economic developments, are not intended, and should not be construed, as forecasts or predictions of the future investment performance of any Fund. Such statements are also not intended as recommendations by any Commonfund entity or any Commonfund employee to the recipient of the material. It is Commonfund's policy that investment recommendations to its clients must be based on the investment objectives and risk tolerances of each individual client. All market outlook and similar statements are based upon information reasonably available as of the date of this material (unless an earlier date is stated with regard to particular information), and reasonably believed to be accurate by Commonfund. Commonfund disclaims any responsibility to provide the recipient of this material with updated or corrected information or statements. Past performance is not indicative of future results. For more information, please refer to [Important Disclosures](#).

Published October 2026



commonfund

Darien, CT 06820

Tel (203) 563-5490

New York, NY 10022

Tel (646) 348-9201

San Francisco, CA 94111

Tel (415) 433-8800

Beijing, China

Tel +86 10 5759 3208

Munich, Germany

Tel +49 892 5007620

601 Merritt 7

Tel 888-TCF-Main

Norwalk, CT 06851

Tel (203) 563-5000

www.commonfund.org