

The Increasing Fight Over Data Centers

Demand, Power, Water, and the Foreign Signal Riding Inside an Authentic Backlash

A Policy and Infrastructure Assessment for Communities, Officials, and Voters

Chris Mark, MBA, DSc (candidate)

DOI: 10.5281/zenodo.20359392

May 2026

Running head: THE REAL FIGHT OVER DATA CENTERS

Abstract

This paper examines the economics, infrastructure, and politics of the United States' artificial-intelligence data-center buildout and the public backlash it has provoked. It argues that two propositions are simultaneously true: that the local opposition rests on authentic and material grievances — rising electricity rates, water consumption, siting conflicts, and the socialization of grid-upgrade costs onto ordinary ratepayers — and that a documented foreign-influence apparatus, including Chinese and Russian state media and an aligned domestic nonprofit network, has amplified that opposition in service of the strategic objective of slowing American AI infrastructure. Distinguishing amplification from origination is the analytic core of the assessment. Applying the constraint-identification logic of Goldratt's Theory of Constraints, the paper identifies the binding constraint on the buildout as the speed at which new electrical capacity can be built and interconnected, and shows that much of the public debate misfires by attacking secondary constraints such as fiber connectivity. It assesses the genuine grid, water, siting, and connectivity questions against current evidence; distinguishes a data center's gross revenue contribution from the net benefit a community actually retains after incentives and socialized costs; and reviews a solution set — cost causation, behind-the-meter generation, transmission and interconnection reform, and demand flexibility — already in deployment. It situates the present anxiety within the recurring historical pattern of technological panic, and concludes with a Socratic framework of questions that enables voters and residents to evaluate a specific proposal on its enforceable terms.

Keywords: data centers, artificial-intelligence infrastructure, electricity grid, cost causation, foreign influence operations, Theory of Constraints, demand flexibility, water and cooling, ratepayer impact, technology anxiety

Executive Summary

The United States is in the middle of the largest, fastest electrical-load expansion in its modern history, driven by artificial intelligence. The public reaction has been intense, local, and increasingly organized — and into that authentic reaction a hostile foreign-influence apparatus has poured amplification. Both of those statements are true at once, and most public commentary fails by collapsing them into one.

This paper makes six arguments:

First, the demand is real and permanent. AI and data growth are not a fad that can be wished away, and developers are responding to genuine market signals, not manufacturing a crisis.

Second, the grid strain is real but misnamed. Data centers are not "overloading" the grid in the sense of melting wires; they are outrunning the grid's ability to *build*. Borrowing the constraint-identification discipline of Goldratt's Theory of Constraints, the binding constraint is construction speed — interconnection queues, transformer lead times, transmission permitting — not raw capacity.

Third, the public anger is fundamentally a *fairness* signal. Under legacy rate design, the capital cost of grid expansion gets socialized across all ratepayers, so an ordinary household can see its bill rise to serve a facility it never wanted. This cost-shift, more than any other single factor, is the engine of the backlash.

Fourth, a documented foreign-amplification layer exists — Chinese and Russian state media, a CCP-aligned U.S. nonprofit network, and foreign-tied funding streams have all converged on the goal of slowing the American AI buildout. But amplification is not origination. Adversaries did not invent the grievance; they found a real one and poured accelerant on it. The strongest evidence is open-source; the weakest claims rest on interested advocacy reports and should be read with that bias in mind. Beneath the specific material complaints sits an even broader and more

exploitable substrate — a rising, diffuse anxiety about AI and technology itself — that foreign actors, domestic operators, and opportunists alike have every incentive to inflame, and which the data center, as the most physical face of "AI," readily attracts.

Fifth, every major objection — power, siting, fiber, and water — has known engineering and policy solutions that are already being deployed. This is a hard problem, not an unsolvable one.

Sixth, the resolution is siting and structuring discipline: put facilities where power and industrial land already are, make large loads pay for the infrastructure they trigger, design for low water consumption, and write community agreements that actually pencil out. Do that, and the legitimate grievance dissolves — and the foreign amplification loses the real fault line it depends on.

The paper closes with a practical set of questions for voters and residents facing a data-center proposal in their own community.

1. The Premise: The Genie Is Out of the Bottle

Any honest analysis has to start by conceding what the loudest opponents will not: the demand is real, large, and durable. AI training and inference, cloud services, streaming, and the broader digitization of the economy are not speculative. Companies are building data centers because customers are paying for compute, not because they are conjuring artificial need.

The scale is genuinely without precedent in the history of the electrical system. The International Energy Agency projects that U.S. electricity demand will grow by nearly 2% annually through 2030, with roughly half of that increase driven by data centers alone [1]. Grid Strategies and others have estimated that more than 150 gigawatts of additional capacity will be

required within five years [2]. For perspective, individual utilities are now receiving interconnection requests for *more power than their entire historical peak demand* [3].

This matters because it reframes the entire debate. "Stop building data centers" is not a policy; it is a refusal to engage with a structural reality. The serious questions are not *whether* this infrastructure gets built, but *where, how it is cooled and powered*, and above all *who bears the cost*. Those are the questions this paper addresses.

A useful discipline for answering them comes from an unlikely source: the operations-management literature. In *The Goal* (1984) — the business novel by the physicist Eliyahu Goldratt that became a staple of MBA and operations curricula — Goldratt set out the Theory of Constraints: the principle that the total output of any system is governed by a single binding constraint, or bottleneck, and that effort spent improving anything *other* than that constraint produces little real gain. The first task, therefore, is always to identify the true constraint and to stop mistaking secondary limits for it. That lens organizes everything that follows. As the next sections show, the binding constraint on the data-center buildout is neither raw power, nor fiber, nor water taken in isolation, but the *speed at which new electrical capacity can be built and interconnected* — and a great deal of the public debate misfires by attacking constraints that only appear to bind. Goldratt taught the principle through a Socratic mentor named Jonah, who never hands the protagonist the answer but instead asks the questions that lead him to find the constraint himself — a method this paper deliberately adopts in its closing section for the citizen facing a proposal.

2. The Grid: Strain, Not Overload

It is true that data centers are straining the power grid. But "overload" is the wrong mental model. The grid is not failing under live load; it is failing to *build* fast enough to keep pace with

demand that scales at software speed while infrastructure scales at the speed of concrete, copper, and permitting.

The numbers are stark. In Texas, CenterPoint Energy reported a roughly 700% increase in large-load interconnection requests, growing from 1 GW to 8 GW between late 2023 and late 2024 [3]. Utilities including ComEd, PPL, and Oncor are reporting more gigawatts of data-center applications than their historical maximum peak demand [3]. ERCOT projects state electricity demand could rise dramatically over the coming decade, requiring thousands of miles of new transmission [4].

The supply side cannot move at that pace, and the reasons are physical, not financial:

- **Interconnection queues.** Lawrence Berkeley National Laboratory research finds interconnection wait times have more than doubled over the past fifteen years, with projects now averaging roughly five years in queue before reaching commercial operation [5]. In high-demand corridors like Northern Virginia, delays can reach seven years, and a majority of queued projects are eventually withdrawn [6].
- **Equipment lead times.** High-voltage transformers carry two-to-four-year lead times; new transmission lines can take a decade to permit and build [7].
- **Planning conservatism.** Grid planners evaluate new large loads against worst-case scenarios — peak demand, a contingency such as the loss of a major line or generator, and the data center running at full load — which makes interconnection slow and cautious by design [8].

The takeaway: the strain is real, but it is a timing-and-construction problem with known engineering levers, not an existential limit. The grid is a slow-building physical system absorbing demand growth it was never architected for. In the language of the constraint lens introduced above, this build-and-interconnect speed — not generation capacity in the abstract — is the system's true bottleneck, and it is the thing every serious policy response must be measured against.

3. Who Pays: The Cost-Causation Problem

Here is the crux of the public anger, and it is more subtle than "data centers are bad." A reasonable person observes that data centers pay their electricity bills and concludes the utility

therefore has every incentive — and obligation — to upgrade. That intuition is half right, and the half it misses is the entire fight.

A regulated utility recovers infrastructure costs through its rate base, and historically those costs have been *socialized* across all ratepayers. So even when a data center faithfully pays for the electrons it consumes, the *capital cost* of the new substations and transmission lines built to serve it can be spread onto every household's bill. Paying for power is not the same as paying for the grid build-out that the new demand triggers, and under legacy cost allocation those have been separate buckets — with the second one shared by everyone.

Policy analysts call the fix "cost causation," and the fact that it does not happen automatically is the grievance in a sentence. As one think-tank analysis put it, data-center load growth is already straining transmission infrastructure in the PJM region, with costs spreading across all ratepayers regardless of who caused them; the single most important ratepayer-protection step is requiring large new loads to pay for the grid infrastructure their demand creates rather than socializing it [9].

There is a second mechanism that raises bills even when capital costs are allocated perfectly: the *wholesale price effect*. When a large new load appears, the grid dispatches more expensive "peaker" plants to meet it, and in a wholesale market the higher clearing price is paid by everyone. Academic work documents that AI data centers raise marginal generation costs and can cause sharp increases in locational marginal prices, particularly in transmission-congested areas [10]. Independent reporting confirms the lived result: a Bloomberg analysis found wholesale electricity costing as much as 267% more in a single month than five years earlier in areas near significant data-center activity [11], and Bank of America economists have linked rising household utility payments directly to AI-driven demand [12].

This is why "they pay for their power, so just upgrade the grid" does not settle the matter. The obligation to serve is real, and utilities do build — but the question of *who is charged for the build* is precisely what the national fight is about.

4. The Foreign Amplification Layer

A genuine domestic grievance of this magnitude is exactly the kind of fault line that hostile state actors seek to widen. The evidence that they are doing so is real — but it must be stated with care, because the topic attracts both overstatement and reflexive dismissal.

4.1 What is documented

State media is operating in the open. *China Daily* published a piece on March 12, 2026, headlined "AI boom sends electricity bills in US skyrocketing," attributing surging power demand and consumer strain to data centers and AI [13]. A compilation by the Bitcoin Policy Institute documents a wider pattern: a CGTN segment in October 2025 blaming data centers for energy-price spikes across the West Coast, Mid-Atlantic, and New England, and Russia's RT promoting data-center bans in several U.S. states [14].

There is also a domestic content layer. In August 2023, *The New York Times* published a front-page investigation tracing a financial network from Chicago to Shanghai that used American nonprofits to push Chinese government positions, centered on Shanghai-based U.S. expatriate Neville Roy Singham and a cluster of interlocking nonprofits [15]. Singham denies the characterization, calling the reporting "misleading and innuendo-laden" and rejecting any claim that he received funds or instructions from China [16]. In a separate, parallel matter, a chargesheet reported on May 3, 2024, in India's NewsClick prosecution named the "Chinese State" as the network's "ultimate paymaster" — an allegation under India's anti-terror law that the outlet contests as an attack on independent journalism [17]. Senators Marco Rubio and Lindsey Graham jointly

asked the U.S. Attorney General to investigate the network for potential FARA violations in August 2023 [18].

Finally, there is a funding layer. An American Energy Institute report, independently carried by multiple outlets, found that twelve U.S. organizations opposing data-center development had collectively received more than \$39 million from foreign donors, primarily Swiss, British, and Danish [19].

4.2 The analytic discipline this requires

Three caveats keep this evidence honest and prevent it from sliding into conspiracy.

First, **the strongest single source is an interested party.** The most detailed account is from the Bitcoin Policy Institute, an organization aligned with AI-infrastructure interests; its primary-source citations check out where they can be verified, but its interpretive leaps deserve the skepticism owed any advocacy brief. The same applies to the American Energy Institute funding report — it is a pro-energy advocacy compilation, not a neutral government finding.

Second, **no enforcement action has followed.** Despite years of scrutiny, no entity in the network has been charged under FARA, none has lost tax-exempt status, and all continue to operate [14]. Critics, including the press-freedom group FAIR, characterize the entire framing as a McCarthyite revival that conflates sympathy for China with agency of it [20]. Convergence on a policy outcome is documented; centralized orchestration is not.

Third, and most important, **the amplification works precisely because the grievance is real.** The *China Daily* article does not fabricate; it quotes an American economist and a Harvard Law electricity expert and cites Lawrence Berkeley National Laboratory data — the same underlying story that Bloomberg, Bank of America, and the IEA were independently reporting [11][12][1]. The foreign hand is in the *selection and volume*, not the manufacture. That is what

makes it both effective and hard to counter: every individual claim is sourceable to a legitimate American outlet, even as the aggregate serves a foreign objective.

4.3 The strategic logic

The asymmetry is the tell. While Beijing's media warns American audiences that data centers are an economic and environmental menace, the Chinese state is reportedly subsidizing a substantial share of its own domestic AI operators' energy costs [14]. An adversary that genuinely believed data centers were ruinous would not be racing to build them faster. The messaging is not a sincere policy position; it is a targeted effort to slow a rival's buildout while accelerating one's own. Notably, the domestic pushback against a federal data-center moratorium has been bipartisan, with some legislators warning that such a measure would effectively hand the AI lead to China [14].

The correct conclusion is neither "the opposition is a foreign plot" nor "foreign influence is paranoia." It is that an authentic domestic backlash is being amplified by adversaries who did not create it — and that the most effective counter is not to discredit residents' concerns but to *resolve* them, which simultaneously removes the fault line the amplification depends on.

4.4 The Deeper Substrate: Technology Anxiety

Beneath the specific, material grievances — bills, water, noise — sits a broader and more diffuse unease about artificial intelligence itself, and it is rising. Pew Research finds that just 10% of U.S. adults feel more excited than concerned about AI's growing role in daily life, against roughly half who feel more concerned — up from 37% concerned in 2021 [43]. The public is far more likely to expect AI to harm them personally (43%) than to benefit them (24%) [44]. A Fox News poll in April 2026 put voter concern at 66%, up from 56% in 2023, with the sharpest increases among women and voters without a college degree [45]. The fears are broad rather than

narrow: a Reuters/Ipsos survey found large majorities worried that AI could be used to incite political chaos, that people might substitute AI companionship for human relationships, and that AI would strain the electrical grid [46].

This matters to the data-center debate because a data center is the most *physical, local, and visible* manifestation of "AI" that exists — a structure one can drive past, that hums, that draws water and power. It therefore becomes a lightning rod for anxieties that have nothing to do with the specific building: job displacement, loss of human connection, surveillance, and a general sense of losing control over the pace of change. The intangible dread acquires a tangible target. And, crucially, this layer is far harder to address than the material objections. One can answer "what is the cooling design" with a specification; one cannot answer a generalized fear of obsolescence with a permit. That non-falsifiability is exactly what makes it more durable and more manipulable than concrete grievances — and there is a self-reinforcing loop in the data, since the public's single largest AI concern is misinformation and manipulation itself [47], the very mechanism through which the anxiety is exploited.

The pattern is neither new nor unique to AI. Every transformative technology has generated a wave of anxiety about its trajectory, and the canonical illustration is the so-called Great Horse Manure Crisis of the late nineteenth century. The underlying problem was genuine: cities ran on horses — New York alone stabled on the order of 100,000, producing millions of pounds of manure daily — and the resulting filth was a real public-health hazard, spreading disease and overwhelming sanitation. The first international urban-planning conference, convened in New York in 1898, reportedly broke up early because the delegates could see no solution [48]. The episode survives in popular memory through a widely repeated prediction, attributed to *The Times* of London in 1894, that within fifty years the city's streets would lie buried under nine feet of

manure. That specific prophecy is of disputed provenance — its modern circulation traces largely to a 2004 essay, and *The Times* itself disowned the attribution in 2018 [48]. The embellishment is part of the lesson: even the archetypal story of technological panic has been dramatized and endlessly recirculated, which is precisely how such narratives behave.

The lesson cuts in both directions, and honesty requires drawing both edges. Against the apocalyptic extrapolations — that data centers will "consume the grid" or "drain the water" on a straight-line projection of today's trend — the horse-manure episode is a caution that linear extrapolation of a present trajectory routinely ignores efficiency gains, substitution, and adaptation. The predicted manure apocalypse never arrived; the problem was rendered moot by an unforeseen technology, the automobile and the electric tram. The equivalent curve-benders here are concrete and already discussed in this paper: liquid cooling that cuts water use by roughly 85%, and flexibility that can unlock on the order of 100 GW of existing grid capacity (Sections 5 and 9). But the analogy also rebukes the glib techno-optimist who deploys it as pure reassurance, because the automobile "solved" manure by introducing pollution, traffic fatalities, and sprawl — externalities arguably larger than the one it cured. Every solution breeds new problems. The honest reading is therefore neither doom nor complacency: the anxiety is real and recurring, the specific catastrophe seldom unfolds as forecast, and the issue remains exploitable in the interim regardless of how it ultimately resolves.

Finally, this substrate is exploitable by far more than foreign adversaries. Domestic political operators have every incentive to ride it — data centers are shaping up as a midterm issue — as do engagement-driven media and assorted opportunists. The danger in responding by dismissing the anxiety as irrational is that doing so widens the roughly fifty-point gap between AI experts' optimism and public concern that surveys consistently document [47], and it is precisely

that sense of being talked down to by distant elites that makes the dread so easy to weaponize. The remedy is the same as for every other grievance in this paper: transparency, honest engagement, and governing the technology visibly rather than imposing it. A data-center hearing conducted on straight facts about that one facility is, in a small way, an act of draining the larger anxiety — a demonstration that the technology can be deliberated and governed rather than merely inflicted.

5. The Solution Set: Five Levers

Because the demand is permanent, the realistic objective is to make supply move closer to the speed of demand while ensuring costs fall on those who cause them. Five levers do most of the work, and the most powerful is the least discussed.

5.1 Cost causation

Bind the large load to the infrastructure it triggers, through large-load tariffs and negotiated service agreements, rather than socializing costs across residential ratepayers. Texas enacted a structural version in 2025 with Senate Bill 6, which redefines the interconnection process for large loads such as data centers exceeding 75 MW — requiring a transmission screening study fee, proof of site control, disclosure of duplicate requests, and a financial commitment to cover transmission infrastructure costs [10]. This is the fairness fix and the precondition for public support of everything else. It reallocates cost; it does not, by itself, accelerate construction.

5.2 Bring your own generation (behind-the-meter)

If the public grid cannot build fast enough, the load builds its own supply on-site and bypasses the interconnection queue — gas turbines, fuel cells, solar-plus-storage, and increasingly nuclear. Industry forecasts suggest a meaningful share of data centers will operate substantially off-grid by 2030 [21]. This is the cleanest expression of "cause the load, build the supply, leave

the commons alone." Its costs are local emissions where on-site gas is used, and a grid that does not get modernized as a byproduct.

5.3 Transmission and interconnection reform

The queues are the actual bottleneck. Reforming federal interconnection rules, reconductoring existing lines with advanced conductors, and deploying grid-enhancing technologies extract more capacity from infrastructure already in the ground [22]. This is necessary, slow, and politically grinding — the unglamorous decade-long work that underpins everything else.

5.4 Flexibility — the underappreciated lever

This is the genuinely counterintuitive finding, and it reframes the problem. The grid is built for its single worst moment — peak demand plus a contingency. Most of the year, substantial headroom sits unused. A data center is treated as a problem largely because planners assume it runs at 100% and can never reduce. Relax that assumption even slightly and the arithmetic changes dramatically.

A 2025 Duke University analysis ("Rethinking Load Growth") found that roughly 76 GW of new load could be integrated into the existing U.S. power system if those loads could curtail for as little as 0.25% of their maximum uptime [23]. A quarter of one percent. And it is not theoretical: in a May 2025 demonstration in Phoenix, software-based workload orchestration cut a 256-GPU cluster's power consumption by 25% during peak grid demand, a reference design that could in principle unlock on the order of 100 GW of existing U.S. grid capacity without new transmission or generation [24]. AI workloads are unusually suited to this — training runs can pause or shift, and compute can be relocated geographically. Researchers have shown that spatially flexible data-

center demand can provide transmission-like operational value and substantially reduce renewable curtailment [25].

The caveat is incentive design: operators will not curtail out of civic virtue, so tariffs and market products must pay them for flexibility [24]. Which loops back to lever 5.1 — flexibility only materializes if rate design rewards it. In constraint-theory terms, flexibility is the purest example of *exploiting* the bottleneck: extracting maximum throughput from the existing grid before incurring the years and billions required to *elevate* it by building anew. Goldratt's sequence is exploit first, elevate second — and most of the current debate skips straight to elevation while leaving enormous latent capacity on the table.

5.5 The framing shift

Underlying all four levers is a change in posture: data centers moving from passive consumers to grid stakeholders that co-invest in infrastructure, supply some of their own power, and provide flexibility in exchange for faster interconnection [26]. The capital at stake is enormous — Goldman Sachs has estimated grid spending on the order of \$720 billion may be needed through 2030 to accommodate data-center growth [24] — which is precisely why capturing that spend as genuine grid modernization, rather than absorbing it as crisis, is the larger prize.

6. Where Should They Go? The Siting Question

If the demand must land somewhere, "not here" from everywhere is not an answer. The encouraging reality is that many communities actively compete to host these facilities — you simply hear less about them, because a town quietly signing a deal is not news while a packed, furious hearing is.

6.1 The communities that want them

The clearest demand comes from rural communities with eroding tax bases and few alternatives. A south Georgia community that lost its largest employer and most of its tax base has been actively recruiting a data center as a survival strategy [27]. Port Washington, Wisconsin, wrote a special zoning code to facilitate construction, with its mayor framing the real threat not as the data center but as the absence of development [27]. The economic logic is specific: data centers generate substantial property-tax revenue without requiring a large workforce, and therefore without stressing housing supply or school systems the way a factory or warehouse would [27]. Oregon, Washington, Louisiana, Colorado, and parts of rural Pennsylvania compete on exactly this basis [28][29].

6.2 The engineering logic of good siting

The ideal site combines abundant, affordable power — ideally near generation, to avoid building long transmission spurs — with a cool climate, water or air-cooling-friendly conditions, and distance from dense residential areas. The most elegant version co-locates the facility with existing generation, such as data centers built adjacent to nuclear plants. Put the load where the power already is, and much of the grid problem in Section 2 evaporates.

6.3 Why it nonetheless feels like pure obstruction

Two dynamics explain the impression that no one wants them anywhere. The first is a collapse in siting discipline under the gold rush. Proposals increasingly place facilities much closer to residential areas, shifting from industrial parks to land in the middle of communities [30]. A data center in an industrial zone next to a substation is uncontroversial; the same building beside a subdivision, with backup generators and constant low-frequency noise, draws a fight every time.

The second is benefit allocation. Communities are turning on their own leaders for approving badly structured deals: officials in Festus and Independence, Missouri, lost their seats

over data-center support, and in rural North Carolina a candidate unseated a four-term incumbent after a nearby approval [29]. One Illinois mayor noted that five data centers provide roughly \$1.6 million a year in taxes while residents question whether that sum justifies the noise and grid strain [31]. When tax incentives are too generous and cost-causation is not enforced, even a willing community ends up feeling fleeced.

The fix is siting and structuring discipline: industrial-zoned land, near existing power, with an enforceable cost-causation deal and a community-benefit agreement that genuinely pencils out. That matches real demand to the places that actually want the investment — and drains the grievance the amplification depends on.

7. The Economic Question: Gross Revenue versus Net Benefit

The first question any community asks is also the one most often answered misleadingly: how much money does a data center actually bring in? The honest answer is that there is no single figure, and that the number a developer presents — the *gross* — is rarely the number the community *keeps*. The gap between the two is where most local conflict over data centers actually lives, and understanding it is the difference between a deal that funds a school district for a generation and one that quietly raises every resident's electricity bill.

7.1 The gross figures are real, and for a small jurisdiction can be large

Property and sales taxes are the genuine prize, and for a rural county with an eroding tax base they can be transformative. A single facility serving a county of roughly 180,000 residents was projected to generate on the order of \$72.5 million in sales-tax revenue from its IT-equipment purchases alone [50]. Virginia, the largest market in the world, draws substantial annual property-tax revenue from each facility, and local construction firms, restaurants, and suppliers benefit materially from the multi-year build [51]. Industry analyses put the broader ripple effects high: the

Data Center Coalition cites roughly \$26 returned in revenue for every \$1 a center spends on services [42], and a frequently cited estimate holds that each direct data-center job supports more than six jobs elsewhere in the economy [52]. These figures sit at the optimistic, interested-party end of the range and should be read as such — but the underlying property-tax stream is real and, for the right jurisdiction, genuinely meaningful.

7.2 The jobs are mostly temporary

Expectations about employment need resetting, because this is where pitches most often mislead. Construction is a genuine, multi-year boom — an industry economic model shows a large build generating thousands of total jobs and billions in output during the construction phase [53]. But permanent operations are remarkably thin: a working data center typically employs between 15 and 100 people depending on its size [52]. The same model that showed thousands of construction jobs showed only about fifty permanent operating positions once the facility opened [53]. A community sold primarily on lasting jobs is being sold the wrong benefit; the durable value of a data center is tax base, not payroll.

7.3 What erodes the gross — incentives and socialized costs

Two deductions separate the headline figure from the net benefit. The first is the incentive package given away to win the project. The most common is a sales-tax exemption on the enormously expensive equipment, cooling, and power gear [52] — which can erase much of the very sales-tax revenue cited as a benefit. In extreme cases the subsidy per permanent job is staggering: watchdog analyses of New York deals have documented packages amounting to well over \$1 million per job [54], and one proposed package working out to roughly \$6.4 million per job [55]. The second deduction is the share of grid-upgrade and wholesale-price costs that lands on local ratepayers (Section 3). In the territory of one Pennsylvania utility, an analysis found the

ratepayer cost of new capacity driven by data-center demand running about \$2.18 billion in a single year — a sum it judged greater than data centers' total contribution to tax revenue [56]. Region-wide, capacity costs in the PJM grid jumped from roughly \$2.2 billion to \$14.7 billion in one year, with data centers responsible for nearly two-thirds of the increase, even as residential electricity rates rose about 32% nationally between mid-2020 and mid-2025 [57].

7.4 How to read the number

The practical lesson is to insist on the *net* figure, computed honestly: gross property and sales tax, minus the value of every abatement and exemption granted, minus the defensible share of grid and infrastructure cost that will fall on existing ratepayers rather than the operator. A well-structured deal — a modest incentive, plus enforceable cost-causation so the operator pays for the grid impact it creates — preserves most of that large gross number as real community benefit. A badly structured one inverts it, and the town ends up subsidizing a facility that raised its own residents' bills. The swing between those outcomes is wide enough that several states which once competed hardest to attract data centers — Virginia, Georgia, and Oklahoma among them — are now reconsidering their incentive programs entirely [57]. The figure that belongs at the center of any community debate is therefore not the developer's headline, but the net — and the questions in Section 10 are designed to extract it.

8. The Connectivity Question: Fiber

A common technical objection holds that rural communities cannot host major data centers because they lack the fiber bandwidth such facilities require. This is half right, and the half it gets wrong explains the rural trend rather than undermining it. It is true that a site chosen primarily for its power often lacks pre-existing carrier-grade connectivity. Industry practitioners note plainly that because such locations are selected for power access, fiber capacity frequently does not exist

nearby, with distances to major metros running hundreds of kilometers [32]. So a remote gigawatt campus does not arrive with connectivity in place.

But fiber is the *soft* constraint and power is the *hard* one. A transmission line or substation is a multi-year ordeal of permitting, land acquisition, and long equipment lead times; reaching a long-haul backbone is a matter of months and millions, not years and billions. Power availability guides where infrastructure is built; fiber connectivity is the engineering follow-on [33]. Critically, hyperscale operators do not rely on local telecommunications infrastructure at all — they build or buy dedicated dark fiber to a backbone route. Major carriers have reported over a billion dollars in AI-related long-haul network deals in a single year, with multiples more in the pipeline, and demand for connectivity in emerging hubs has spiked by thousands of percent [34].

The deeper reason rural works is the nature of the dominant workload. AI *training* is latency-tolerant: it is batch work that can run for days, largely indifferent to distance from population. Remote locations function well for it because latency plays a minimal role unless training data must be fetched from far away [35]. The activity consuming the most power and provoking the most local opposition is therefore precisely the activity that does not need to sit near anyone.

The important caveat is the division of labor. *Inference* — serving a model to end users in real time — is latency-sensitive and does require metro proximity [33][35]. So the fiber-bandwidth concern is valid, but for the inference/edge layer near cities, not for the training megacampuses driving the rural land rush and the grid fights. A newer wrinkle reinforces that connectivity is rising in strategic importance: operators increasingly split single training runs across multiple campuses linked by long-haul fiber, making inter-data-center connectivity a first-order design pillar [36]. For practitioners, the real questions are not "is there fiber" but whether there is *lit*, *route-diverse*

capacity — a single long-haul cut can strand a gigawatt campus — and whether the latency budget supports the distributed-training fabric.

In short: rural sites usually lack fiber on day one, but that is the financeable, fast-to-solve part of the equation, and the workload that wants those sites tolerates the distance. The binding constraint remains power. Fiber follows power; it does not gate it — in the constraint vocabulary of Section 1, it is a *subordinate* resource, not the bottleneck, and treating it as the bottleneck is exactly the kind of misdiagnosis the constraint lens is meant to prevent.

9. The Water and Cooling Question

Water is the most misunderstood element of the entire debate, because public discussion conflates two fundamentally different quantities and ignores that the cooling technology is changing rapidly.

9.1 Withdrawal versus consumption

"Data centers use X gallons" almost always blurs *withdrawal* — water drawn and largely returned — with *consumption* — water actually lost, primarily to evaporation. The villain in the water fights is one specific design: **evaporative cooling**, which consumes water to the atmosphere and is popular because it trades water for electricity. A 100-megawatt facility on evaporative cooling can use anywhere from 300,000 to more than one million gallons per day, and more in extreme heat [37].

9.2 The technology is moving away from the thirsty design

For high-density AI workloads, the engineering trend is decisively toward water-light cooling. **Closed-loop liquid cooling** recirculates fluid in a sealed circuit, requiring a large initial fill and then very little ongoing water; one operator states that a closed-loop facility uses less water than five typical households [38]. **Air-cooling** can approach zero process water — one major

operator has committed that an upcoming Texas facility will limit water use to critical campus operations like kitchens [39]. The lever is real and large: in one Central Texas case, switching a proposed facility from evaporative to air-cooled design reduced its projected water needs by roughly 85% [37]. The "data centers guzzle water" narrative is accurate for the cheap legacy design and increasingly inaccurate for the new one. The decisive questions are therefore *which cooling design* and *is it consumptive*, not "do they use water."

9.3 Source matters as much as quantity

Where the water comes from determines whether a facility competes with residents. The fights that escalate involve facilities drawing from aquifers that feed residential wells, or buying treated municipal drinking water — which is why the sharpest disputes in Texas have arisen in areas of high water stress [40][41]. Drawing on a large surface-water body is directionally better. But a managed water-supply *reservoir* is not a free heat sink: consumptive draw still competes with the municipal supply the reservoir exists to provide, simply upstream of the tap rather than at it, and that competition intensifies during drought drawdowns. Open-loop "once-through" cooling, which returns most water but warmer, raises a different regulated concern — *thermal discharge* — because elevated temperatures affect dissolved oxygen, fish, and algae in the receiving water, and are governed under the Clean Water Act. On a recreational lake, thermal impact is its own dispute independent of consumption.

9.4 Power-rich, water-poor sites

A site can have ample power but insufficient water — the West Texas profile of cheap land, abundant power, and minimal water. This is rarely a hard physical stop, because cooling is a tradeoff between water and power: air and dry-cooling use almost no water but consume more electricity and lose efficiency in heat, which is acceptable where power is plentiful. As a Purdue

researcher framed it, an identical facility "could consume the same amount of water in Texas and Iowa," but Iowa's cooler, wetter climate offers free cooling that Texas heat denies — so the Texas site either consumes more water or spends more power to avoid it [41]. A power-rich, water-poor site is therefore pushed toward liquid or air cooling and a higher power bill, not disqualified. The genuinely unbuildable sites are those short on both.

9.5 Regulatory direction

States are beginning to act. The Texas Commission on Environmental Quality has moved to review water-use plans for new data centers above a capacity threshold; California has advanced legislation requiring projected water-use reporting before operations begin; and Minnesota requires developers to consult the state environmental agency on supply adequacy [40][41]. The trajectory is toward disclosure and water-source scrutiny as standard conditions of approval.

10. Considerations for Voters and Residents

A data-center proposal is neither a gift nor a curse by default. It is a deal, and the terms determine which it becomes. The objective for a resident or voter is not to be reflexively for or against, but to interrogate the specific project well enough to tell a well-structured one from a badly structured one. The following framework distills the analysis above into questions worth bringing to a hearing.

The approach here is deliberately Socratic, in the spirit of Goldratt's Jonah: it does not hand residents a verdict, because a verdict imposed from outside invites resistance and rarely fits the particular facts of a particular town. Instead it supplies the questions that let a community locate the real constraint for itself, test the developer's claims, and judge the deal on its enforceable terms. The questions are grouped by the constraint each one probes.

10.1 On cost — who pays for the grid?

- Will the operator pay the full capital cost of the transmission and substation upgrades its demand requires, under an enforceable large-load tariff or service agreement — or will those costs enter the general rate base and appear on residential bills?
- Is there a written cost-causation commitment, or only a voluntary corporate pledge with no oversight mechanism? Voluntary pledges can be exited at will and depend entirely on state-level negotiation.
- What is the projected effect on local wholesale and retail electricity prices, and who has modeled it — the developer, or an independent party?

10.2 On the deal — does the benefit pencil out?

- What is the *net* annual tax revenue after any abatement or incentive, and over what period? A large tax break can erase the headline benefit.
- Is there a community-benefit agreement, and what does it concretely fund (schools, roads, water infrastructure)?
- How many *permanent* jobs result, as distinct from temporary construction employment? Operational headcount for data centers is typically small.
- Who pays for road, public-safety, and infrastructure impacts during multi-year construction?

10.3 On water — design and source

- Is the cooling design evaporative (consumptive), closed-loop liquid (low ongoing use), or air-cooled (near-zero process water)?
- What is the projected *consumption* — not just withdrawal — in gallons per day, including peak summer figures?
- What is the water *source*: municipal potable supply, groundwater shared with residential wells, or surface water under a specific contract? What happens to that draw during a declared drought restriction?
- If once-through cooling is proposed, is there a thermal-discharge permit, and what is the projected temperature impact on the receiving water?

10.4 On siting — location and livability

- Is the parcel zoned industrial and sited near existing power and transmission, or is it residential-adjacent land chosen for a fast grid connection?
- What is the projected noise profile at the nearest homes, and what mitigation is enforceable?
- What backup-generation technology is proposed (diesel, gas), and what are the local air-quality and emissions implications?
- How close are schools, parks, and residences, and what buffer is committed in writing?

10.5 On power and reliability

- Will the facility bring some of its own generation (behind-the-meter), or rely entirely on the existing grid?

- Has the operator agreed to *demand flexibility* — curtailing during grid emergencies — in exchange for interconnection, and is that commitment enforceable?
- Does the local utility have the capacity to serve this load without degrading reliability for existing customers, and who has verified that?

10.6 On process — transparency and authority

- Was the proposal disclosed early, or did the community learn of it only when the developer needed water, power, or a tax break?
- Does the local government have the staff and expertise to evaluate the technical claims, or is it relying on the developer's own figures?
- What enforcement and clawback provisions exist if the operator fails to meet its commitments?

10.7 A note on the information environment

Residents should be aware that this issue sits inside a contested information space. Some of the loudest "ban them everywhere" messaging is amplified — though not originated — by foreign state media and aligned networks whose interest is in slowing American AI infrastructure, not in protecting any particular town's water or rates (Section 4). At the same time, some of the loudest "the opposition is foreign-funded" messaging comes from industry-aligned advocacy groups with a direct financial stake in discrediting local concerns. The practical defense against both is the same: focus on the *specific terms of the specific project*, demand independent verification of the developer's claims, and judge the deal on its enforceable commitments rather than on either side's framing. A legitimate local concern does not become illegitimate because an adversary amplifies it, and a real influence operation does not make residents' water bills imaginary.

11. Conclusion

The American data-center buildout is not a crisis to be stopped or a gift to be accepted uncritically. It is a large, permanent shift in how the country uses electricity, arriving faster than the grid was built to absorb, and landing in communities that are right to ask hard questions. The grid strain is real but is a construction-speed problem with known fixes. The public anger is real

and is fundamentally about fairness — about a legacy cost structure that lets the price of serving these facilities fall on people who never wanted them. The foreign amplification is also real, and is effective precisely *because* the underlying grievance is genuine; the answer is to resolve the grievance, not to dismiss it.

Every major objection — power, siting, connectivity, water — has an engineering or policy solution already in deployment. The path that satisfies the demand while protecting residents is consistent across all of them: site facilities where power and industrial land already exist, make large loads pay for the infrastructure they trigger, design for minimal water consumption, secure flexibility commitments, and structure community deals that genuinely pencil out. That combination matches real demand to the communities that want the investment, protects those that do not, and removes the authentic fault line that hostile amplification depends upon.

For the resident facing a proposal at the next county hearing, the lesson is neither fear nor enthusiasm but diligence. The questions in Section 10 are answerable. Insisting that they be answered — in writing, with independent verification, before approval — is how a community converts a contested megaproject into a fair one, or correctly declines a bad one.

References

1. International Energy Agency, *Electricity 2026*, Executive Summary. <https://www.iea.org/reports/electricity-2026/executive-summary>
2. Grid Strategies, load-growth forecast (November 2025), as cited in Novogradac analysis of interconnection and transmission. <https://www.novoco.com/notes-from-novogradac/resolving-the-interconnection-queue-bottleneck-along-with-transmission-expansion-is-critical-for-timely-us-energy-deployment-to-meet-demand>
3. CenterPoint Energy large-load interconnection data and utility application figures, as reported by Camus Energy, "Why Does It Take So Long to Connect a Data Center to the Grid?" (July 2025). <https://www.camus.energy/blog/why-does-it-take-so-long-to-connect-a-data-center-to-the-grid>
4. Introl, "PJM Grid 6GW Shortfall: 2027 Data Center Power Crisis" (February 2026), citing ERCOT demand projections. <https://introl.com/blog/pjm-grid-crisis-6gw-shortfall-data-center-power-2027>

5. Lawrence Berkeley National Laboratory interconnection-queue research, as summarized in Hanwha, "Data Center Grid Limitations: The Power Bottleneck" (February 2026). <https://www.hanwhadatacenters.com/blog/data-center-grid-limitations-the-power-bottleneck/>
6. "Part 4: How Interconnection Queues Are Stalling Data Center Growth," Acres (July 2025). <https://landvalues.acres.com/part-4-interconnection-queues-stalling-data-center-growth>
7. "Data Center Power in Months, Not Years — The Behind-the-Meter Solution," Datapower Supply (May 2026). <https://www.datapowersupply.com/post/data-center-power-in-months-not-years-the-behind-the-meter-solution>
8. Camus Energy (see [3]) on utility "green book" worst-case interconnection evaluation.
9. Searchlight Institute, "Seizing the Data Center Buildout for Grid Modernization" (March 2026). <https://www.searchlightinstitute.org/research/seizing-the-data-center-buildout-for-grid-modernization/>
10. "Electricity Demand and Grid Impacts of AI Data Centers: Challenges and Prospects," arXiv preprint (2025), on wholesale price effects and Texas SB6. <https://arxiv.org/html/2509.07218v3>
11. Bloomberg News investigation into electricity costs near data centers, as summarized by CNET/AOL, "The AI Data Center Boom Is Driving Up Electricity Costs" (2026). <https://www.aol.com/articles/ai-data-center-boom-driving-214400830.html>
12. Bank of America (David Tinsley) analysis of utility costs and AI demand, as reported by Fortune/Yahoo (2025). <https://finance.yahoo.com/news/ai-boom-making-utility-bills-154103282.html>
13. "AI boom sends electricity bills in US skyrocketing," *China Daily*, March 12, 2026. <https://global.chinadaily.com.cn/a/202603/12/WS69b1f569a310d6866eb3d4f5.html>
14. Sam Lyman, "Foreign Influence in the Campaign against American AI," Bitcoin Policy Institute, May 2026. <https://www.btcpolicy.org/articles/foreign-influence-in-the-campaign-against-american-ai> (*Advocacy source aligned with AI-infrastructure interests; interpretive claims should be read with that bias in mind. The CGTN and RT specifics rest on this compilation and were not independently confirmed for this paper.*)
15. Mara Hvistendahl, David A. Fahrenthold, Lynsey Chutel, and Ishaan Jhaveri, "A Global Web of Chinese Propaganda Leads to a U.S. Tech Mogul," *The New York Times*, August 5, 2023. <https://www.nytimes.com/2023/08/05/world/europe/neville-roy-singham-china-propaganda.html>
16. "Never received funds from any government or political party: Neville Roy Singham," *Deccan Herald*/PTI, October 2023 (Singham's statement to *The Hindu*). <https://www.deccanherald.com/india/never-received-funds-from-any-government-or-political-party-neville-roy-singham-2730697>
17. "Chinese State was the ultimate paymaster of NewsClick, says Delhi Police chargesheet," *The Hindu*, May 3, 2024. <https://www.thehindu.com/news/national/chinese-state-was-the-ultimate-paymaster-of-newsclick-says-delhi-police-chargesheet/article68133189.ece> (*A contested allegation under India's UAPA; NewsClick rejects it.*)

18. Sens. Marco Rubio and Lindsey Graham, joint letter to the U.S. Attorney General requesting a FARA investigation, August 2023. <https://www.rubio.senate.gov/rubio-graham-to-biden-doj-we-must-not-allow-the-ccp-to-succeed-in-information-war-against-our-country/>
19. American Energy Institute foreign-influence report (April 2026), as reported by The Center Square. https://www.thecentersquare.com/national/article_ee075c72-b681-4618-96e5-75de408ef28c.html (*Pro-energy advocacy compilation from tax filings and grant databases; the directness of the link to data-center opposition specifically is disputed.*)
20. Joshua Cho, "NYT Reveals That a Tech Mogul Likes China—and That McCarthyism Is Alive and Well," FAIR, August 17, 2023. <https://fair.org/home/nyt-reveals-that-a-tech-mogul-likes-china-and-that-mccarthyism-is-alive-and-well/>
21. Bloom Energy 2026 Power Report, as cited in Introl, "PJM Grid Crisis" (see [4]).
22. RMI, "The Interconnection Queue Continues to Be a Barrier to American Economic Competitiveness" (March 2026). <https://rmi.org/interconnection-reform-ai-data-centers-generator-queues/>
23. Duke University, "Rethinking Load Growth" (2025), as summarized in Knowledge Problem, "Data Centers, Flexibility, and the Architecture of the Grid" (2026). <https://knowledgeproblem.substack.com/p/data-centers-flexibility-and-the>
24. Introl, "Data Centers as Grid Stabilizers" (January 2026), on the NVIDIA/EPRI Phoenix flexibility demonstration and Goldman Sachs grid-spending estimate. <https://introl.com/blog/data-centers-grid-stabilizers-flexible-power-nvidia-epri-2026>
25. "Grid Operational Benefit Analysis of Data Center Spatial Flexibility," arXiv preprint (2025). <https://arxiv.org/pdf/2511.08759>
26. "2026 Predictions: AI Sparks Data Center Power Revolution," Data Center Knowledge (April 2026). <https://www.datacenterknowledge.com/operations-and-management/2026-predictions-ai-sparks-data-center-power-revolution>
27. "Data centers are expensive, unpopular — and could be a tipping point in the midterms," NPR/OPB (April 2026). <https://www.npr.org/2026/04/20/g-s1-117729/data-center-disputes-local-midterms>
28. "A blueprint for rural data center growth," Data Center Dynamics (March 2026). <https://www.datacenterdynamics.com/en/marketwatch/a-blueprint-for-rural-data-center-growth/>
29. NPR/OPB (see [27]) on electoral backlash in Festus and Independence, Missouri, and rural North Carolina.
30. "Data centers as 2025 newsmaker of the year," on the industrial-to-residential siting shift in Northeast Pennsylvania, via Yahoo News (2026). <https://www.yahoo.com/news/articles/data-centers-2025-newsmaker-200200098.html>
31. "How do data centers benefit the places where they're built? Local mayors give mixed reviews," WJBC (May 2026). <https://www.wjbc.com/2026/05/11/how-do-data-centers-benefit-the-places-where-theyre-built-local-mayors-give-mixed-reviews/>

32. "Power Drives the AI Data Center Boom, but Connectivity Cannot be Overlooked," Data Center Knowledge (May 2026).
<https://www.datacenterknowledge.com/networking/power-drives-the-ai-data-center-boom-but-connectivity-cannot-be-overlooked>
33. "The Backbone of AI: Performance, Proximity and Bandwidth," CoreSite (August 2025).
<https://www.coresite.com/blog/the-backbone-of-ai-performance-proximity-and-bandwidth>
34. "Data Center Bandwidth Soars 330%, Driven by AI Demand," Data Center Knowledge (November 2025), citing Zayo long-haul deal figures and Memphis/Salt Lake City demand growth. <https://www.datacenterknowledge.com/networking/data-center-bandwidth-soars-330-driven-by-ai-demand>
35. "AI and Latency: Why Milliseconds Decide Data Center Winners," Data Center Knowledge (December 2025). <https://www.datacenterknowledge.com/infrastructure/ai-and-latency-why-milliseconds-decide-winners-and-losers-in-the-data-center-race>
36. Corning, "AI Data Center Design Strategies," on distributed long-haul training architectures. <https://www.corning.com/optical-communications/worldwide/en/home/the-signal-network-blog/ai-data-center-design-strategies.html>
37. "How One Texas County Struck a Deal With Its Data Centers," *Texas Monthly* (April 2026), on Medina County and the evaporative-to-air-cooled redesign. <https://www.texasmonthly.com/news-politics/data-center-water-shortage-medina-county/>
38. "Texas Data Center Developers Play Offense on Water," Inside Climate News (April 2026), on closed-loop systems (Skybox). <https://insideclimatenews.org/news/10042026/texas-data-center-developers-play-offense-on-water-claiming-huge-cuts-in-usage/>
39. "Google promises its upcoming Texas data center won't use any water for cooling," New Atlas (February 2026). <https://newatlas.com/environment/google-data-center-texas-water-cooling/>
40. "How Data Centers Are Shaping Texas's Water Crisis," Datacenters.com (July 2025), on evaporative cooling, TCEQ review, and water-stress geography. <https://www.datacenters.com/news/texas-water-crisis-the-impact-of-data-center-developments-on-a-fragile-resource>
41. "Data centers are thirsty for Texas' water, but state planners don't know how much they will need," *The Texas Tribune* (September 2025), including Purdue researcher Yi Ding on Texas-versus-Iowa consumption and other states' disclosure rules. <https://www.texastribune.org/2025/09/25/texas-data-center-water-use/>
42. "Texas Republicans have a data center problem," *The Texas Tribune* (May 2026), documenting Rep. Will Metcalf's recruitment of a data center for the Conroe Technology Park. <https://www.texastribune.org/2026/05/07/texas-republicans-data-centers-rural/>
43. Pew Research Center, "How the US Public and AI Experts View Artificial Intelligence" (April 2025), on rising public concern (10% more excited than concerned) and the expert-public gap. <https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/>

44. Pew Research Center, "What the data says about Americans' views of artificial intelligence" (2026), on harm-versus-benefit expectations and concern rising from 37% in 2021.
<https://www.pewresearch.org/short-reads/2026/03/12/key-findings-about-how-americans-view-artificial-intelligence/>
45. "Fox News Poll: Broad anxiety about AI doesn't extend to jobs," Fox News (April 1, 2026), on 66% voter concern (up from 56% in 2023) and demographic shifts.
<https://www.foxnews.com/politics/fox-news-poll-broad-anxiety-about-ai-doesnt-extend-jobs>
46. Reuters/Ipsos survey (August 2025), as reported by eWeek, on the breadth of AI fears (political chaos, AI companionship, electricity strain).
<https://www.eweek.com/news/americans-fear-ai-job-loss-survey/>
47. Stanford HAI, *2026 AI Index Report*, Public Opinion chapter, on expert-versus-public gaps; and Statista Consumer Insights, on misinformation and manipulation ranking as the public's leading AI concern. <https://hai.stanford.edu/ai-index/2026-ai-index-report/public-opinion>
48. On the Great Horse Manure Crisis: Stephen Davies, "The Great Horse-Manure Crisis of 1894," Foundation for Economic Education (2004), for the popular framing and the 1898 urban-planning conference; and the historiographical caveat that the widely quoted *Times* "nine feet of manure" prediction is of disputed provenance and was disowned by *The Times* in 2018 (see the documented history at https://en.wikipedia.org/wiki/Great_horse_manure_crisis_of_1894). Davies essay: <https://fee.org/articles/the-great-horse-manure-crisis-of-1894/>
49. Eliyahu M. Goldratt and Jeff Cox, *The Goal: A Process of Ongoing Improvement* (North River Press, 1984), the founding text of the Theory of Constraints and a standard reference in operations-management and MBA curricula; the constraint-identification logic and the Socratic "Jonah" method referenced throughout this paper derive from it.
50. "Measuring the Data Center Boom: Facts and Statistics (2026)," Programs.com, on a projected \$72.5 million in sales-tax revenue from a facility's IT-equipment purchases.
<https://programs.com/resources/data-center-statistics/>
51. "Data center tax breaks are on the chopping block in some states," Stateline (February 2026), on Virginia property-tax revenue and the difficulty of measuring net benefit versus incentive cost. <https://stateline.org/2026/02/24/data-center-tax-breaks-are-on-the-chopping-block-in-some-states/>
52. "States Rethink Data Center Tax Incentives as Costs Soar," MultiState (February 2026), on the sales-tax-exemption incentive model, the 15–100 permanent-employee range, and the industry estimate that each direct job supports more than six others.
<https://www.multistate.us/insider/2026/2/4/states-rethink-data-center-tax-incentives-as-costs-soar>
53. IMPLAN, "The Rise of Data Centers" (2025), an economic-impact model contrasting large construction-phase employment and output with roughly fifty permanent operating jobs.
<https://support.implan.com/hc/en-us/articles/36262006944155-The-Rise-of-Data-Centers>

54. "This Data Center Is Getting a \$77 Million Tax Break to Create One Job," New York Focus (April 2026), on per-job subsidy magnitudes. <https://nysfocus.com/2026/04/20/data-center-tax-break-jpmorgan-chase>
55. "Data center wants subsidies of \$6.4 million per job," Investigative Post (February 2026). <https://investigativepost.org/2026/02/04/data-center-subsidies-of-6-4-million-per-job/>
56. "Why Data Centers Will Be Economic Development Duds," Ohio River Valley Institute (February 2026), on the Duquesne Light ratepayer-cost analysis exceeding data centers' tax contribution. <https://ohiorivervalleyinstitute.org/why-data-centers-will-be-economic-development-duds/> (*Advocacy source skeptical of data-center incentives; its underlying utility and PJM figures are independently reported — see [57].*)
57. Dany Bahar and Greg Wright, "New evidence on data center employment effects," Brookings (May 2026), on PJM capacity costs rising from \$2.2 billion to \$14.7 billion in a single year (about two-thirds attributable to data centers), the ~32% rise in residential electricity rates from July 2020 to July 2025, and the reconsideration of incentive programs in Virginia, Georgia, and Oklahoma. <https://www.brookings.edu/articles/new-evidence-on-data-center-employment-effects/>

This paper synthesizes publicly reported information as of May 2026. Where load-bearing claims rest on advocacy-aligned sources, that interest is noted in the references so readers can weigh the evidence accordingly. It is analysis, not legal or investment advice.