Thematic Investing

Powering AI: The Politics of Data Centers

SIGNATURE

Today's politics will shape tomorrow's AI capacity. As data centers become politicized, local opposition, permitting delays, and regulatory fragmentation are determining where, when, and whether AI infrastructure gets built.

Thematic Investing
William Thompson
+1 212 526 8641
william.s.thompson@barclays.com
BCI, US

Hannah Greenberg
+1 212 526 2457
hannah.greenberg@barclays.com
BCI, US

Public Policy
Michael McLean[v]
+1 212 526 9393
michael.mclean@barclays.com
BCI, US

Sustainable Investing
Anika Agarwal[v]
+1 212 526 1322
anika.agarwal@barclays.com
BCI, US

AI's next bottleneck may be "permission." Data center opposition is broad and bipartisan, driven by concerns over electricity costs, water use, emissions, and local impacts. This is fueling pushback even in traditionally pro-growth U.S. states and pitting national AI leadership against local concerns. While the Trump administration continues to frame AI as an economic and national-security imperative and is pursuing multiple policy tools, many key decisions remain local through utility regulation, zoning, and permitting. Hyperscalers are responding with community engagement, but many concerns may prove structural. Bring-your-own-power helps mitigate cost-shifting concerns, though long-term growth still depends on grid expansion. This puts focus on compute efficiency, next-gen energy tech, and even space-based data centers.

Report highlights:

- **AI is a national priority, but "permission" remains local.** The Trump administration is pursuing federal land access, transmission expansion, and nuclear deployment, yet a growing patchwork of state and local laws, utility rules, and zoning and permitting requirements increasingly determines where data centers get built. Development continues to migrate toward more favorable jurisdictions, but bottlenecks often follow, while AI sovereignty and national-security considerations limit the ability to relocate capacity abroad.
- **Can hyperscalers win over local communities?** Permitting and political decisions made today will influence data center capacity later in the decade. Industry is emphasizing jobs, infrastructure funding, tax revenue, transparency, and sustainability, but many concerns appear structural rather than reputational. Electricity prices may continue to rise, water-use debates remain complex, and speed-to-power has increasingly taken priority over emissions.
- **Growing urgency for new solutions.** Bring-your-own-power models can accelerate deployment and reduce cost-shifting concerns, but they aren't a panacea as long-term growth still requires grid expansion. Compute and model-level efficiencies, advanced energy technologies, and distributed architectures could help ease constraints, while even space-based data centers are attracting attention.

Barclays Capital Inc. and/or one of its affiliates does and seeks to do business with companies covered in its research reports. As a result, investors should be aware that the firm may have a conflict of interest that could affect the objectivity of this report. Investors should consider this report as only a single factor in making their investment decision.

[v] This author is a registered US equity research analyst who is subject to US FINRA Rule 2241 and who may write debt research under FINRA Rule 2242.

Please see analyst certifications and important disclosures beginning on page 20.
Completed: 02-Sep-26, 12:41 GMT Released: 02-Sep-26, 12:45 GMT Restricted - External

CONTENTS

Executive Summary	3
Data Centers Face Growing Political Resistance	4
Can Hyperscalers Win Over Local Communities?	6
Emerging Conflicts Between Local & State Controls	9
Federal Influence Remains Powerful but Complicated	13
You Can Move the Data Center, Not the Constraints	17
Is this a Win for Behind-the-Meter?	18

Executive Summary

"Permission" may be emerging as another AI bottleneck. While power, capital, labor, and compute remain critical constraints, they aren't sufficient to maintain the industry's social license to operate. As data centers become increasingly politicized, local opposition, permitting delays, and regulatory fragmentation are emerging as meaningful risks to AI infrastructure deployment. Polling suggests resistance is broad and bipartisan, driven by concerns over electricity affordability, grid reliability, water and land use, emissions, community impacts, and transparency. As a result, even traditionally pro-development states are beginning to impose additional safeguards and oversight on data center development.

AI is a national priority, but many key decisions remain local. The Trump administration continues to frame AI infrastructure as a strategic economic and national-security imperative and is pursuing tools such as federal land access, transmission expansion, power-market reform, and nuclear deployment. Yet many of the most consequential decisions remain outside federal control, reflecting the fragmented nature of the U.S. utility, permitting, and regulatory landscape. Moreover, many actions taken thus far, including executive orders and agency guidance, may prove less durable than statutory changes.

Data centers remain primarily local issues, albeit with growing state involvement. Siting and development decisions are still largely governed through local planning, zoning, and permitting processes, though states and utility regulators are becoming increasingly involved through utility regulation, environmental permitting, grid planning, and ratepayer protections. The result is a growing patchwork of laws, utility rules, moratoria, zoning restrictions, and permitting requirements that increasingly determine when, where, and how AI infrastructure can be built.

Data centers can move, but the bottlenecks often follow. Development continues to migrate toward more favorable jurisdictions, yet power constraints, local opposition, and political scrutiny frequently follow. At the same time, AI sovereignty, data-residency, and national-security concerns limit the industry's ability to simply relocate capacity abroad.

Can hyperscalers win over local communities? The industry is increasingly emphasizing community engagement, infrastructure investment, tax revenue, workforce development, transparency, and sustainability. However, many concerns appear structural rather than reputational. Electricity prices may continue to rise, water-use debates remain complex, and speed-to-power has increasingly taken priority over emissions. The central question is whether communities believe the benefits outweigh the costs.

Today's political and permitting decisions may shape AI capacity later in the decade. Data centers typically take several years to develop, meaning today's decisions are likely to influence future capacity additions. To the extent permitting challenges and community opposition slow hyperscale deployment, scarce compute capacity could become even more valuable.

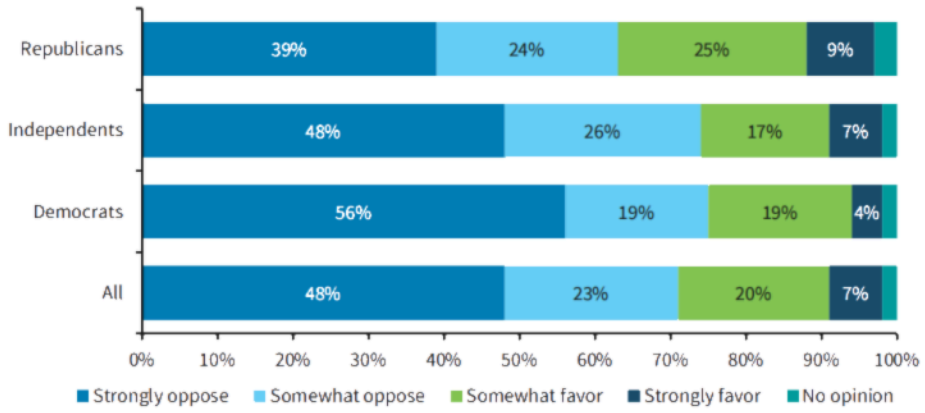
This is creating urgency for new solutions. Bring-your-own-power (BYOP) models can accelerate deployment and reduce cost-shifting concerns, but they aren't a panacea as grid expansion remains the long-term solution. Growing infrastructure constraints are increasing the importance of compute efficiency, advanced energy technologies, and more distributed computing architectures. They are also drawing attention to longer-term possibilities, including space-based data centers, as potential ways to reduce reliance on constrained terrestrial infrastructure.

Data Centers Face Growing Political Resistance

The politicization of data centers is emerging as a major risk to the AI buildout. While power, capital, compute, and labor are widely cited as the primary constraints on data center development, we have long argued that permitting and local opposition are risks to the industry's social license to operate, as outlined in Powering AI: More Data Centers, More Problems and discussed in the Barclays Brief: Cooling the AI buildout. Concerns over utility bills, grid reliability, emissions, water and land use, and broader community impacts are fueling NIMBYism—"not in my backyard" resistance—turning data center development into a growing political issue ahead of the November elections.

Polling suggests opposition to data center development in the United States is both broad and bipartisan. In a March Gallup survey, 71% of respondents said they would oppose the construction of a data center in their community, compared with 53% who opposed a nuclear power plant. While opposition was strongest among Democrats, majorities of Republicans and Independents also expressed concerns. More recent polling from Politico suggests sentiment has deteriorated further, including among self-identified MAGA voters.

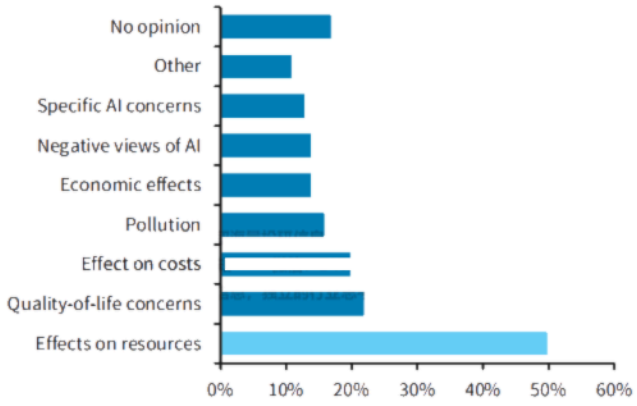
FIGURE 1. **Gallup survey: Overall, would you strongly favor, somewhat favor, somewhat oppose or strongly oppose the construction of a data center in your area to support artificial intelligence, or AI, technology in the U.S.?**



Polling dates: March 2-18, 2026
Source: Gallup

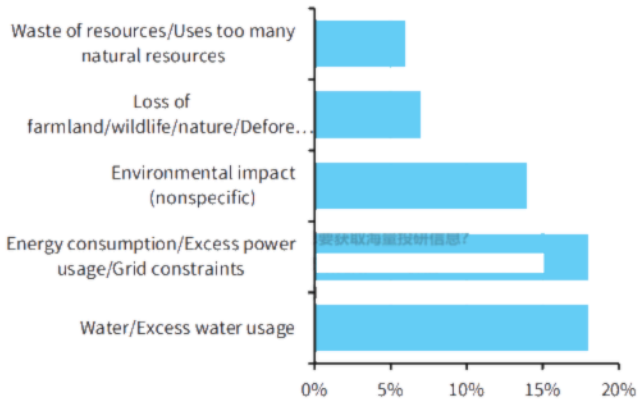
While the minority of respondents who support data centers primarily cite economic benefits (66%), opponents express a broad range of concerns. More than half point to resource-related issues, including electricity consumption (18%) and water use (18%). Other frequently cited concerns include impacts on local quality of life (22%); higher costs (20%), particularly utility bills; environmental impacts (16%) such as noise, air pollution, and water use; negative effects on local jobs and economies (14%); and broader concerns about AI itself (14%).

FIGURE 2. **Gallup poll: What are some of the reasons you would oppose the construction of a data center in your area?**



Polling dates: April 1-15, 2026. Totals for all mentions exceed 100% because respondents could give multiple responses.
Source: Gallup

FIGURE 3. **Breakdown of 'effect on resources' mentions**



Polling dates: April 1-15, 2026.

Source: Gallup

These results suggest that opposition is not driven by any single issue. Rather, it reflects a broader distrust of the industry and a perception that local communities bear the costs while the benefits accrue elsewhere. In some cases, confidentiality agreements, proprietary business information, and limited public disclosure during the development process have reinforced perceptions that communities lack sufficient information to fully evaluate potential project impacts. Furthermore, growing concerns about AI, privacy, data security, and the societal influence of large technology companies may be reinforcing these views. As a result, the economic and national-security arguments for data centers in the AI race against China (Deep impact: The global AI arms race) often resonate more strongly in Washington, Silicon Valley, and on Wall Street than in the communities being asked to host these facilities.

These results suggest that opposition is not driven by any single issue. Rather, it reflects a broader distrust of the industry and a perception that local communities bear the costs while the benefits accrue elsewhere. Many local officials and residents are also being asked to evaluate projects without clear public disclosure of expected electricity use, water consumption, transmission requirements, or backup-generation plans, potentially contributing to concerns about transparency and trust. Furthermore, growing concerns about AI, privacy, data security, and the societal influence of large technology companies may be reinforcing these views. As a result, the economic and national-security arguments for data centers in the AI race against China (Deep impact: The global AI arms race) often resonate more strongly in Washington, Silicon Valley, and on Wall Street than in the communities being asked to host these facilities.

Reflecting the growing political sensitivity of the issue, the National Republican Senatorial Committee recently warned AI companies that data center backlash could jeopardize a critical Ohio Senate seat. Even in traditionally pro-development states, policymakers are imposing additional safeguards and oversight as data center development becomes increasingly controversial. Texas Governor Greg Abbott (R) directed the state's utility regulator and grid operator to review proposed data center projects before advancing them through the interconnection process. Pennsylvania Governor Josh Shapiro (D) issued an executive order requiring developers to meet state infrastructure standards and obtain local approval before permit applications can be reviewed. In Michigan, Republican Senate candidate Mike Rogers has called for a one-year moratorium on new data center development. In Virginia—the world's largest data center market—Governor Abigail Spanberger (D) highlighted a State Corporation

Commission (SCC) order requiring data centers to pay for transmission infrastructure built exclusively to serve their facilities rather than allocating those costs to other ratepayers, reinforcing the state's focus on assigning grid expansion costs to large-load customers. Meanwhile, New Jersey Governor Mikie Sherrill (D) announced new energy and water-use reporting requirements for data centers, municipal guidance for negotiating community benefits agreements, and state technical assistance programs, completing all four pillars of her data center governance framework.

Why are data centers often viewed differently than other infrastructure projects? One explanation is that they combine many of the perceived costs associated with traditional infrastructure development, including significant electricity demand, water consumption, land use, transmission expansion, and, in some cases, emissions, without providing equally visible local benefits. Unlike manufacturing facilities, data centers typically support relatively few permanent jobs once operational. Unlike roads, hospitals, or schools, they offer limited direct utility to local residents. And unlike many renewable energy projects, data centers are increasingly linked to AI, a technology that remains the subject of public debate over its implications for employment, privacy, and broader societal outcomes. These concerns may contribute to growing public skepticism toward data center development.

Supporters, however, argue that focusing solely on direct employment understates the broader economic benefits of these projects. They point to construction spending, supplier networks, tax revenues, and induced economic activity as meaningful sources of local value creation. Proponents also contend that AI infrastructure may support advances in healthcare, scientific research, education, climate-related technologies, and broader economic prosperity. More broadly, some observers view opposition to data centers as part of a wider pattern of resistance to large-scale infrastructure development, sometimes characterized by the acronym BANANA ("Build Absolutely Nothing Anywhere Near Anyone").

Can Hyperscalers Win Over Local Communities?

Hyperscalers have focused on improving community engagement and making the economic case for data centers. In our view, hyperscalers appear to recognize that maintaining a social license to operate through improved community engagement may prove just as important as securing land, capital, or electricity. Microsoft, Amazon, Google, Meta, OpenAI, and others have launched initiatives emphasizing local job creation, workforce development, infrastructure investment, water stewardship, and commitments to fund the infrastructure needed to support their facilities.

For example, Microsoft's Community-First AI Infrastructure initiative is built around five core commitments:

- Ensuring that its data centers do not increase electricity costs for existing customers by paying for the infrastructure they require.
- Minimizing water consumption and replenishing more water than it uses.
- Creating jobs and workforce-development opportunities for local residents.
- Expanding local tax revenues that support schools, hospitals, libraries, parks, and other public services.
- Strengthening communities through investments in AI education, digital-skills training, and local nonprofit organizations.

The challenge, however, is that many of the concerns driving local opposition are structural rather than reputational. Even with credible commitments to fund grid upgrades and reduce water and emissions impacts, electricity prices may continue to rise in many parts of the U.S. for reasons largely unrelated to hyperscaler behavior. The debate over water use is more nuanced

than it appears, involving trade-offs between direct water consumption, electricity demand, embedded upstream water use, and power-sector emissions. And in the current environment, speed-to-power has increasingly taken priority over cost and emissions in the AI arms race, complicating the industry's ability to demonstrate progress toward its long-term sustainability goals.

Ultimately, the industry's challenge is not simply demonstrating that data centers create value, but convincing host communities that they receive a fair share of the value being created. To that end, hyperscalers are increasingly seeking to position AI infrastructure as a catalyst for local and national economic development. Supporters argue that the benefits extend beyond direct employment to include construction activity, long-term tax revenues, investments in energy infrastructure, workforce training, educational programs, digital-skills development, and support for local community organizations. They also contend that AI infrastructure underpins broader societal gains, including advances in healthcare, scientific research, education, productivity, economic growth, and national competitiveness.

Electricity Affordability Remains at the Center of the Debate

Data centers are not the sole driver of higher electricity bills. While data centers face criticism on multiple fronts, electricity affordability has emerged as one of the more politically salient issues. For many communities, the central question is not whether data centers consume large amounts of power, but who ultimately pays for the infrastructure required to serve them.

Average U.S. residential electricity prices in 1Q26 were 12% higher than in 1Q24 and rose even more sharply in many regions experiencing significant data center growth, including Washington (+24%), Virginia (+17%), Pennsylvania (+17%), Wyoming (+16%), Utah (+15%), Illinois (+14%), Louisiana (+14%), Ohio (+14%), and Georgia (+13%).

However, we would not attribute these increases solely to hyperscaler growth. Much of the rise reflects years of underinvestment in grid infrastructure, including transmission and distribution systems, as well as growing spending on grid hardening to address more severe weather events. Amazon has similarly argued in a post that a meaningful portion of rising electricity costs reflects the need to modernize the transmission and distribution infrastructure that delivers electricity to homes and businesses, while noting that hyperscalers generally pay for their own electricity use. Nevertheless, rapid growth in data center demand has complicated an already challenging environment for utilities and regulators to balance reliability, affordability, and investment incentives simultaneously.

We are also skeptical of the argument that data centers will automatically lower electricity prices by spreading fixed grid costs across a larger customer base. That thesis assumes sufficient excess generation and transmission capacity already exists. While this may be true in select regions (e.g., Nevada, where utility costs have declined) and during certain periods, most AI data centers require reliable 24/7 power and often increase demand when the grid is already constrained. Cooling requirements frequently cause data center demand to peak during the summer, precisely when electricity systems face the greatest stress from air-conditioning load. As a result, accommodating incremental data center demand often requires new dispatchable generation, transmission upgrades, and other infrastructure investments rather than the utilization of existing spare capacity.

Are hyperscalers willing to pay their share of these costs? Our sense is yes. In our view, the question is less whether hyperscalers are willing to pay and more whether local communities believe those payments adequately compensate for the perceived costs. Converting watts into monetizable AI tokens rather than traditional low-margin industrial output fundamentally changes the economics of power. A 1 gigawatt (GW) AI data center campus can cost roughly $50

billion to build, with approximately two-thirds of that investment tied to compute infrastructure. Annual compute depreciation alone may exceed $6 billion, compared with an annual electricity bill of roughly $560 million (assuming an $0.08/kWh average industrial rate for 1-GW at 80% utilization). Even funding dedicated generation at more than $2,000/kW would add only around $2 billion of incremental capital to a $50 billion campus. Given that a single gigawatt of AI capacity can potentially support tens of billions of dollars of annual revenue, hyperscalers appear increasingly willing to fund dedicated generation, transmission upgrades, and other supporting infrastructure if doing so accelerates deployment.

Several recent initiatives point in this direction. Hyperscalers have increasingly committed to funding the infrastructure required to support their facilities, while the Trump Administration's Ratepayer Protection Pledge calls on developers to build, bring, or buy the energy needed to support new data centers. Developers are also increasingly funding dedicated generation, transmission upgrades, and other supporting infrastructure directly. In fact, a number of projects have claimed to provide energy cost savings or broader grid benefits to existing ratepayers. Meta's Hyperion campus in Louisiana is being paired with substantial new power infrastructure investments to support its long-term electricity requirements and is projected to provide approximately $2.7 billion of customer savings to Entergy Louisiana customers. Sponsors of the PORTS-Pike Technology Campus in Ohio have committed to fund approximately $4.2 billion of grid and transmission upgrades.

As we discuss below, this dynamic is generally supportive of Behind-the-Meter (BtM) and Power-as-a-Service (PaaS) models because they can accelerate deployment while reducing the risk of shifting costs onto other ratepayers. Utilities are also becoming more disciplined, implementing data center-specific tariffs (e.g., AEP Ohio's "Data Center Tariff") and requiring large upfront payments, guarantees, minimum-bill provisions, and termination fees to ensure data center customers bear the full cost of the infrastructure they require. The goal is increasingly clear: protect existing ratepayers while ensuring AI infrastructure can still be built.

Hyperscaler Net-Zero Targets Remain Challenging

Hyperscalers have made ambitious net-zero commitments, but progress has been more gradual than many expected. In our view, the industry has underestimated both its continued reliance on fossil fuels and the complexity of achieving its stated net-zero goals. While hyperscalers routinely claim to use 100% renewable energy, we continue to emphasize the distinction between (a) annually matching electricity consumption through renewable power-purchase agreements and renewable energy certificates and (b) matching electricity demand with carbon-free generation every hour of every day on the same grid.

The latter is a significantly higher standard for claiming net-zero operations. Google remains the only hyperscaler providing comprehensive disclosure regarding its progress toward its 24/7 carbon-free energy goal for 2030. While the company has matched 100% of its annual electricity demand with renewable energy since 2017, its 24/7 carbon-free energy mix has remained in the mid-60% range, implying that approximately one-third of electricity demand is not matched by carbon-free generation on an hourly basis. Google has indicated that future progress may remain uneven due to challenges related to transmission buildout, interconnection queues, permitting delays, and other infrastructure constraints (Powering AI: Insights from AMZN & GOOG Sustainability Reports).

Water Use Is More Complicated Than It Appears

It's not simply a matter of "less direct water use is better." Data center water use has become a central concern for local communities. As we detailed in Powering AI: Demystifying Data Center Cooling & Water Use and summarized in our accompanying chart pack, there are many misconceptions regarding data center cooling and related water use.

For example, liquid cooling is often incorrectly associated with elevated water risk, when evaporative cooling—not liquid cooling—is the primary source of direct on-site water consumption. Importantly, cooling a data center involves two distinct steps: (i) removing heat from servers and (ii) rejecting that heat into the atmosphere. Liquid cooling is part of the first step, while evaporative cooling is often part of the second.

Evaporative cooling is a long-standing HVAC technique used in both commercial buildings and data centers. Warm air is passed through wetted media or fine water sprays, causing a portion of the water to evaporate. This phase change from liquid to vapor absorbs heat and enables more energy-efficient heat rejection. The process is analogous to human perspiration—as sweat evaporates, it absorbs heat and cools the body.

Why use evaporative cooling? It generally requires less energy to reject heat into the atmosphere, particularly in hot climates or during peak summer conditions. As a result, data center cooling and water-use decisions involve a trade-off among electricity demand, direct water use, indirect upstream water consumption, and power-sector emissions. Reducing evaporative cooling can lower direct water use but often increases electricity demand, which can in turn raise emissions and indirect upstream water consumption because thermoelectric power plants—including natural gas, coal, and nuclear facilities—can themselves be highly water intensive. Meta's embedded water consumption associated with purchased electricity, for example, was nearly 13 times its direct water use in 2024.

While hyperscalers have improved their public communication around liquid cooling—typically closed-loop systems in which coolant is continuously recirculated—there remains meaningful disagreement regarding the appropriate use of evaporative cooling. Meta and Microsoft have increasingly sought to avoid evaporative cooling in order to reduce direct water consumption, while Google and Amazon continue to utilize evaporative cooling, particularly in lower-risk regions, to optimize energy efficiency, reduce emissions, and lower overall water consumption when both direct and indirect water use are considered.

However, growing public scrutiny over water use concerns has already translated into policy action. Multiple states have introduced water-use disclosure requirements, dedicated water-permitting processes for data centers, and enhanced local review authority. This suggests that water use is likely to become an increasingly material variable in evaluating siting decisions, entitlement risk, and long-term operating exposure for AI infrastructure. As a result, we expect greater emphasis on water-reuse systems, water-replenishment programs, reclaimed-water infrastructure, and other technologies designed to reduce freshwater consumption (Water Tech: Wave of Opportunity).

Emerging Conflicts Between Local & State Controls

Data centers remain primarily local issues, albeit with growing state involvement. Across most of the United States, data center siting and development decisions are still governed through local planning, zoning, and permitting processes, with limited state involvement beyond standard utility and environmental regulation. In these jurisdictions, project outcomes are determined primarily by local elected officials, planning commissions, zoning boards, and community stakeholders.

However, in many regions experiencing significant data center development, increasingly influence projects through utility regulation, tax incentives, eco development programs, infrastructure planning, and environmental permitt become larger and more resource-intensive, tensions between state and loc increasing. Some states have sought to restrict development (e.g., New York'

人在草木间

statewide moratorium), impose additional safeguards and oversight (e.g., recent actions in Texas and Pennsylvania), or reconsider the scope of data center tax incentives and other economic-development programs (e.g., Ohio and Arizona's pauses on certain data center tax incentives), while others have limited local authority to encourage investment (e.g., West Virginia's state-preemption framework). The result is a growing patchwork of laws, utility rules, moratoria, zoning restrictions, and permitting requirements that increasingly determine when, where, and how future AI infrastructure can be built.

As a result, permitting risk is becoming increasingly geography-specific, making regulatory geography an important consideration in site selection, project timelines, development economics, and ultimately the pace of AI infrastructure deployment. To simplify this evolving landscape, we categorize state approaches into six primary governance models.

FIGURE 4. **Different regional approaches to data center governance**

Governance category	Examples	Key characteristics
State-led preemption	West Virginia	State legislation limits local authority over qualifying high-impact data centers and microgrid districts; economic development prioritized at the state level.
Direct state oversight	New York	State has taken a direct role in regulating data center growth through statewide permitting, environmental reviews, utility cost-allocation policies, and a temporary statewide moratorium on new hyperscale developments.
Mixed state-local oversight	Texas, Ohio, Pennsylvania	Local governments retain primary zoning authority, but states, utilities, and regulators increasingly shape outcomes through grid planning, interconnection requirements, cost-allocation mechanisms, reliability oversight, environmental permitting, and ratepayer-protection policies.
Primarily local control with increasing state involvement	Virginia, Georgia, Maryland, California	Local governments remain the primary decision-makers, but state policymakers are becoming increasingly involved in issues such as electricity affordability, utility cost allocation, environmental impacts, water use, tax incentives, and long-term infrastructure planning.
Local control with active moratoria / zoning restrictions	Parts of Virginia, Georgia, Maryland	Local governments are increasingly using moratoria, zoning amendments, permit pauses, and discretionary review processes to slow or reshape development.
Primarily local control	Most other states	Data center siting and zoning remain largely governed by local planning and zoning processes, with state involvement focused primarily on tax incentives, utility regulation, and economic development policy.
Federal / Tribal jurisdiction	Federal lands, Navajo Nation, other Tribal lands	Development is governed primarily by federal agencies or Tribal governments rather than traditional state and local zoning processes.

Source: Barclays Research

Primary local control

In much of the country, data center development remains governed largely through traditional local planning, zoning, and permitting processes. While this model generally provides the greatest degree of local autonomy, it can also create significant variation in permitting timelines, development standards, and project outcomes, particularly as data centers become more visible and politically contentious.

These jurisdictions often offer developers considerable flexibility, but they can also produce the widest range of permitting outcomes because approvals depend heavily on local politics, community sentiment, and the priorities of individual elected officials. As a result, even jurisdictions that currently operate under traditional local-control frameworks may face increasing pressure to adopt elements of the governance models outlined below.

State-led preemption

What is state preemption? State preemption occurs when a state government limits or overrides the ability of local governments to regulate a particular activity. In the context of data centers, preemption laws can restrict counties and municipalities from imposing their own zoning, permitting, siting, or development requirements, shifting authority from local governments to the state.

It is important to distinguish state preemption from by-right development. By-right development reduces permitting risk because projects that comply with existing local zoning requirements can proceed without discretionary approvals. State preemption reduces permitting risk by limiting the ability of local governments to restrict or block development. In short, by-right development works through local zoning certainty, while preemption works through state override of local authority.

At present, West Virginia represents one of the clearest examples of data center-specific state preemption in the United States. In 2025, lawmakers passed House Bill 2014, the Power Generation and Consumption Act, designed to encourage investment in high-impact data centers and associated microgrids. The legislation includes provisions such as public-records exemptions, the creation of a certified microgrid development program, special tax treatment for qualifying projects, and other economic-development incentives. The bill also preempts counties and municipalities from regulating qualifying high-impact data centers and certified microgrid districts, effectively shifting significant siting and development authority from local governments to the state.

Supporters argue that preemption reduces permitting uncertainty, accelerates project timelines, and improves a state's ability to attract large-scale investment. Critics contend that it reduces local control over projects that may have significant impacts on land use, infrastructure, water resources, and community character. Since the bill's enactment, debate has increasingly focused on whether additional oversight mechanisms and community protections are needed as large-scale projects move from concept to implementation. In response, state leaders recently announced West Virginia's Responsible Data Center Development Plan, a framework intended to complement HB 2014 through enhanced project screening, ratepayer protections, infrastructure planning, environmental stewardship, and local economic-benefit provisions while maintaining the state's broader pro-development approach.

Mixed state-local oversight

Texas, Ohio, and increasingly Pennsylvania represent a mixed state-local governance model in which states increasingly influence development through utility regulation, environmental permitting, grid planning, and ratepayer protections, while local governments retain substantial authority over zoning, siting, and land-use decisions.

In these states, the policy debate has increasingly shifted from whether data centers should be built to how the costs, reliability impacts, and infrastructure requirements associated with large-load development should be managed. Questions regarding cost allocation, ratepayer protections, grid reliability, infrastructure investment, and community benefits have become central policy concerns.

Direct state oversight

New York has taken a more direct state-led approach to regulating data center development. In July 2026, Governor Kathy Hochul (D) issued an executive order establishing the nation's first statewide moratorium on new hyperscale data center development. The order temporarily pauses certain state environmental permits for large projects while regulators develop a

comprehensive framework addressing grid reliability, environmental impacts, community protections, and cost allocation.

While New York remains the most prominent example of direct state intervention, proposals for statewide data center moratoria have also emerged in several other states. To date, however, most have failed, stalled in committee, been vetoed, or remain under consideration. Maine's legislature approved a statewide moratorium in 2026, but Governor Janet Mills (D) vetoed the legislation. Similar proposals in Minnesota, New Hampshire, Oklahoma, and South Dakota failed to advance, while additional proposals remain under discussion in Michigan, Pennsylvania, Vermont, and Virginia.

While direct state intervention can create greater regulatory certainty through uniform statewide standards, it can also increase policy risk if future administrations revise, weaken, or reverse those standards. Many of the actions taken thus far—including executive orders, agency guidance, and administrative policies—may prove less durable than statutory or constitutional changes. As a result, the long-term regulatory environment could remain subject to political cycles and changes in leadership.

Local Control with Active Moratoria & Zoning Restrictions

In parts of Virginia, Georgia, Maryland, and Michigan, local governments have moved beyond traditional zoning oversight and pursued moratoria, permit pauses, zoning revisions, overlay districts, and other measures designed to slow or reshape data center development.

This reflects a broader shift from viewing data centers primarily as economic-development opportunities to treating them as major infrastructure projects with significant resource, utility, and community impacts. Concerns regarding electricity demand, water use, transmission infrastructure, noise, ratepayer impacts, and community character have become increasingly prominent in local decision-making.

For developers, these markets can present elevated entitlement and permitting risks, as project approvals increasingly depend on discretionary review, political considerations, litigation risk, and community acceptance rather than traditional by-right development pathways.

Federal / Tribal jurisdiction

Data center development on federal lands and Tribal lands operates under a distinct governance framework that falls outside traditional state and local zoning structures. Projects located on federal property are generally subject to federal agency oversight, environmental review requirements, and federal permitting processes. Similarly, development on Tribal lands is governed primarily by Tribal governments exercising sovereign authority over land use, permitting, taxation, and economic development.

While these jurisdictions currently represent a relatively small portion of the U.S. data center market, they may become increasingly relevant as developers seek access to large land parcels, power infrastructure, renewable energy resources, and locations with fewer local zoning constraints. Federal policymakers have increasingly promoted the use of Department of Energy (DOE) sites for AI infrastructure development, identifying multiple federal locations with existing power and industrial infrastructure suitable for large-scale data center campuses. For example, the U.S. Department of Energy has identified several potential AI infrastructure sites on federal property, including up to 8 GW of compute capacity in Ohio, more than 1.2 GW in Kentucky, and 1 GW in South Carolina.

However, such projects typically involve unique regulatory, environmental, legal, and jurisdictional considerations that differ materially from traditional state and local development processes. DOE has also identified other federal sites, including Idaho National Laboratory, Oak

Ridge Reservation, and Savannah River Site, as potential locations for future AI data center and energy infrastructure development.

Primarily Local Control with Increasing State Involvement

In states such as California, Louisiana, Maryland, and parts of Virginia and Georgia, local governments remain the primary decision-makers for data center siting, zoning, and land-use approvals. However, state policymakers have become increasingly involved in debates surrounding utility cost allocation, environmental impacts, water use, tax incentives, and long-term infrastructure planning.

Unlike mixed-governance states where utility regulators and state agencies increasingly shape project outcomes through formal oversight mechanisms, local governments in these jurisdictions remain the primary gatekeepers for project approvals. State involvement is generally focused on broader policy issues—such as electricity affordability, grid reliability, environmental standards, and economic development—rather than direct intervention in individual projects.

Nevertheless, growing state engagement suggests that data center development is increasingly being viewed as a matter of statewide economic, environmental, utility, and industrial policy rather than solely a local land-use issue. Over time, these jurisdictions may evolve toward more formal state oversight frameworks as policymakers seek greater consistency in how large-scale data center projects are evaluated and regulated across local markets.

Federal Influence Remains Powerful but Complicated

AI is a national priority, and Washington is pulling multiple levers to accelerate deployment—but many of the most consequential decisions remain local. Despite the political backlash to data centers, the White House, Department of Energy (DOE), Department of Commerce, and other federal agencies have increasingly framed AI infrastructure as a strategic asset critical to U.S. technological leadership, economic growth, and national security. The White House's July 2025 AI Action Plan describes AI leadership as both an economic and national-security imperative, and dedicates an entire pillar to accelerating AI infrastructure development through permitting reform, grid expansion, workforce development, and secure data center deployment. Data center construction has also become a meaningful contributor to economic activity through capital investment and job creation.

President Trump has repeatedly emphasized the strategic importance of AI infrastructure, arguing that U.S. leadership in AI is essential to maintaining economic and national-security advantages over China. Similarly, EPA Administrator Lee Zeldin has warned against restricting data center development, citing the potential implications for U.S. competitiveness in the global AI race.

At the same time, the Administration has sought to address concerns regarding electricity affordability. President Trump stated on Truth Social that he "never want[s] Americans to pay higher Electricity bills because of Data Centers" and subsequently promoted a Ratepayer Protection Pledge signed by hyperscalers, utilities, data center developers, governors, and state legislators. The pledge calls on the industry to build, bring, or buy the energy required to power new data centers and to bear the full cost of the associated generation, transmission, and supporting infrastructure. However, the pledge remains voluntary. It neither imposes binding obligations nor addresses many of the issues driving local opposition, including concerns regarding water use, land use, emissions, noise, community impacts, and local control.

The Administration has also sought to influence regional electricity markets directly. In January 2026, President Trump and a bipartisan group of 13 governors urged PJM Interconnection—the

operator of the nation's largest power grid covering all or parts of 13 states in the Mid-Atlantic and Midwest regions of the U.S.—to reform its market rules and protect residential ratepayers from rising electricity costs associated with data center growth. The proposal included capacity-market and interconnection reforms designed to accelerate new supply, improve load forecasting, and ensure large-load customers bear a greater share of the costs associated with supporting rapid demand growth. While PJM's response was broadly supportive, it stopped short of adopting several elements outright, instead directing additional stakeholder review and analysis. The episode underscored both the growing pressure to accommodate AI-related load growth and the limits of federal influence over regional electricity-market governance.

While the federal government has multiple tools to accelerate AI infrastructure development—including federal reviews, access to federal land, interstate transmission oversight, wholesale-market rules, and nuclear deployment—its influence is complicated by the fragmented nature of the U.S. utility, permitting, and regulatory system. Many of the most consequential decisions affecting data center development, including transmission investment, interconnection rules, cost allocation, reliability standards, and ratepayer protections, remain governed through regional electricity markets and state regulatory processes. FERC oversees interstate transmission service, wholesale electricity markets, and RTO/ISO tariffs. Regional grid operators and transmission owners are responsible for planning, interconnection studies, market administration, and implementation. State utility commissions retain authority over retail rates, retail cost allocation, utility resource approvals, and, in many jurisdictions, generation siting. Local governments continue to control zoning, land use, building approvals, noise restrictions, setbacks, and community conditions, while state and local environmental agencies oversee air, water, and other environmental permits.

Nevertheless, the Trump Administration has pursued a range of federal initiatives within its authority to accelerate transmission development, reduce interconnection backlogs, reform large-load connection processes, expand access to federal land, and advance nuclear generation capacity.

Transmission Expansion and Interconnection Reform

One area where the federal government can exert meaningful influence is transmission infrastructure. Despite the industry's focus on power generation, transmission increasingly represents one of the most significant bottlenecks to AI infrastructure deployment. Electrons and molecules are not useful if they cannot be moved to where they are needed. In many regions, access to transmission capacity has become a more immediate constraint than access to generation itself.

Recognizing this constraint, OpenAI's economic blueprint calls for a *National Transmission Highway Act* designed to unlock the "three Ps"—planning, permitting, and payment—for the transmission, communications, and pipeline infrastructure needed to support a national AI buildout. The proposal draws comparisons to the 1956 National Interstate and Defense Highways Act, which helped create the modern U.S. interstate highway system. However, building an AI-era transmission network is considerably more complex than building highways in the 1950s. Transmission projects routinely cross multiple states, utility territories, and regulatory jurisdictions, requiring agreement among stakeholders with differing economic and political interests.

FERC Order 1920 requires transmission providers to conduct long-term regional transmission planning, evaluate a broader set of future reliability and economic needs, and develop more transparent frameworks for allocating the costs of new transmission infrastructure among beneficiaries. The rule is intended to improve long-term planning and address one of the

industry's most persistent challenges: determining who pays for infrastructure that benefits multiple regions, utilities, and customer classes.

Beyond transmission planning, FERC Order 2023 seeks to reform how new projects connect to the grid. Recent reforms have shifted from a "first-come, first-served" model toward a more streamlined "first-ready, first-served" approach designed to remove speculative projects and accelerate viable developments. While primarily focused on generation interconnection, these efforts indirectly benefit data centers by helping bring new generation online more quickly.

FERC has also directed each of the six regional grid operators under its jurisdiction to justify or reform the rules governing how data centers, manufacturing facilities, and other large-load customers connect to the electric grid. These proceedings increasingly focus on cost allocation, queue priority, reliability impacts, and the extent to which existing ratepayers should bear the costs associated with supporting large new loads.

However, such reforms go only so far. A truly national transmission buildout would likely require extensive infrastructure-permitting reform, resolution of longstanding cost-allocation disputes (i.e., determining who pays for new infrastructure and who benefits), and potentially significant changes to utility-market structures that have evolved over nearly a century.

These challenges help explain why previous bipartisan efforts at infrastructure-permitting reform have repeatedly struggled to gain traction. Policymakers generally agree that more infrastructure is needed but disagree on which projects should be prioritized, how environmental reviews should be conducted, the appropriate roles of fossil fuels and renewable energy, and how costs should be distributed among utilities, consumers, and large-load customers. Transmission projects also face many of the same permitting, land-use, and community-opposition challenges confronting data centers themselves.

Federal Land and Supporting Energy Infrastructure

In addition to making federal land available for AI infrastructure development, federal agencies can influence deployment timelines through permitting decisions affecting transmission lines, natural-gas pipelines, generation facilities, and other supporting infrastructure required for large-scale data center development.

Executive Order 14318 seeks to accelerate qualifying projects that require federal approvals, financing, or access to federal land. However, it does not preempt state utility regulation, local zoning, or most retail-rate decisions, limiting its ability to address many of the bottlenecks affecting data center development.

Complementing this effort, the DOE initially identified 16 potential AI infrastructure sites on federal property and subsequently advanced four—Idaho National Laboratory, Oak Ridge Reservation, Paducah Gaseous Diffusion Plant, and Savannah River Site—toward private-sector solicitations. These sites offer developers access to existing infrastructure, established security frameworks, and potentially faster permitting timelines, but they represent a targeted rather than nationwide solution to broader deployment constraints.

Notably, the DOE has advanced several large-scale AI infrastructure projects on federal land, including up to 8 GW of compute capacity (4.25 GW initial phase; option for add'l 3.75 GW) at OpenAI's proposed PORTS-Pike Technology Campus in Ohio, more than 1.2 GW at the former Paducah Gaseous Diffusion Plant site in Kentucky, and 1 GW at the Savannah River site in South Carolina, as detailed below.

FIGURE 5. **Major proposed US data center projects on federal land**

Projects	Involved Parties	Capacity	What we know				
			General	Timing	Power	Financing	Community Claims
PORTS-Pike Technology Campus (Pike County, Ohio)	SoftBank, SB Energy, OpenAI, Nvidia, AEP Ohio, U.S. Dept. of Energy, U.S. Dept. of Commerce	8 GW of IT (4.25 GW initial phase; option for add'l 3.75 GW) supported by 10+ GW of power	On remediated DOE land (former Portsmouth Gaseous Diffusion Plant) and surrounding private land. SB Energy will build, own and operate the DC under a 20-yr lease to OpenAI. NVIDIA will provide credit support and will be the exclusive AI compute infrastructure provider.	800 MW will start construction in 2026 with capacity online in phases beginning in 2028.	9.2 GW ($33bn) of gas generation funded in part through $550bn U.S.-Japan Strategic Trade and Investment Agreement and expected to be owned by the U.S. Gov't. SB Energy & Softbank will fund $4.2bn through AEP Ohio to pay for 100% of grid upgrades and new power transmission.	NVIDIA will invest $1.5bn in SB Energy and will provide credit support on land, power and shell buildout for 4.25 IT-GW, with an option to take the remaining 3.75 IT-GW.	Project sponsors state that the project will create 4,000 construction jobs and 300-400 permanent operating jobs, while strengthening grid reliability without shifting upgrade costs to existing customers. OpenAI and SB Energy have committed an $80mm community benefits fund.
DOE's Former Paducah Gaseous Diffusion Plant Site (McCracken County, KY)	Brookfield, NextEra Energy, Big Rivers Electric Power Corp., Jackson Purchase Energy Cooperative, Paducah Power System, U.S. Dept. of Energy	1.2+ GW of IT and 1.8 GW of power, supported by up to 4.6 GW of dedicated generation and storage.	Brookfield will lease DOE land and develop/operate the data center campus at the former Paducah enrichment site as part of DOE's American Energy Hubs initiative.	Development will occur in phases, with dedicated generation added as load ramps. Campus expected to be fully operational by 2032.	NextEra will own and operate up to 2 GW of gas generation and 2.6 GW of battery energy storage. Generation and storage are intended to be built specifically for the project, with excess electricity potentially available to the regional grid.	$100bn privately funded project. Brookfield will develop the DC campus, while NextEra will fund and build dedicated generation and transmission upgrades.	Project sponsors state the development will protect residential and small-business ratepayers from cost increases, improve regional grid reliability, create approximately 8,000 construction jobs and 600 permanent jobs, and represent one of the largest investments in Kentucky history.
Savannah River Site (Aiken County, SC)	Amentum, DOE/ NNSA, DC BLOX (consortium partner), U.S. Dept. of Energy	1 GW of IT, supported by ~2 GW of dedicated on-site generation	DOE/NNSA selected Amentum to enter negotiations for a phased lease to develop an AI DC and dedicated energy infrastructure at the Savannah River Site, one of four DOE locations selected for AI and energy development on federal land.	Amentum selected in July 2026. Project remains subject to lease negotiations, permitting, safety reviews, and other federal approvals.	~2 GW of on-site generation, consisting of gas bridging to nuclear energy. The project is intended to provide dedicated power for the DC while increasing power availability to the grid.	Financing details have not been publicly disclosed. Project is being pursued through a public-private partnership on federal land.	Project sponsors state the development will strengthen U.S. AI leadership, expand reliable energy generation, avoid shifting power costs to existing utility customers, and support national security objectives.

Source: US Department of Energy, company reports and Barclays Research

Advancing Nuclear Capacity

Federal policymakers have increasingly emphasized nuclear energy as one of the few scalable technologies capable of simultaneously supporting AI-driven electricity demand growth, grid reliability, emissions-reduction objectives, and national-security priorities. Recent federal efforts have focused on streamlining reactor licensing, supporting advanced reactor and small modular reactor (SMR) development, strengthening domestic nuclear-fuel supply chains, and encouraging the continued operation of the existing nuclear fleet.

However, significant challenges remain, including reactor construction costs, supply-chain constraints, workforce availability, licensing timelines, and local opposition to new nuclear infrastructure. Many of the advanced reactor and SMR technologies most frequently discussed in the context of AI infrastructure remain years away from commercial deployment at scale. As a result, while nuclear power may play an important role in supporting long-term AI infrastructure growth, it is unlikely to materially alleviate the industry's near-term power and interconnection constraints.

You Can Move the Data Center, Not the Constraints

Why can't data centers simply migrate to more favorable regions? They already are—but the underlying bottlenecks often follow. Historically, data centers were concentrated near major population centers because conventional cloud services required low latency—the time it takes for data to travel between a data center and the end user. AI is changing that equation. AI training workloads are largely latency-insensitive, while the performance of many AI inference workloads is constrained more by compute availability than by physical distance. In other words, adding a few milliseconds of network latency often matters far less than access to sufficient computing resources.

As a result, data center development has expanded beyond traditional hubs such as Northern Virginia into parts of the South and Midwest in pursuit of more favorable regulatory environments, lower-cost land, available transmission capacity, and economic incentives such as tax abatements and other development incentives. However, many of these regions are now experiencing the same challenges that drove development elsewhere, including growing interconnection backlogs, transmission constraints, and local opposition. In other words, data centers can move, but the underlying bottlenecks often follow.

Even with the greater flexibility afforded by behind-the-meter power, developers continue to face constraints related to natural-gas supply, pipeline availability, and the permitting of supporting energy infrastructure. In many cases, the challenge shifts from securing grid interconnection to securing fuel, transportation, and permitting infrastructure.

Furthermore, this migration is increasingly pushing development into regions that may lack sufficient skilled labor and supporting infrastructure to accommodate large-scale construction activity. Many emerging data center markets also face elevated exposure to extreme weather, water scarcity, and other climate-related risks. Areas such as West Texas offer abundant land, favorable regulatory environments, and access to energy resources, but they also expose operators to greater weather-related disruption risk and growing concerns regarding long-term water availability. As data center development becomes increasingly concentrated in these regions, operators may face higher insurance costs, tighter underwriting standards, and greater scrutiny of resilience, business continuity, and climate-adaptation planning. Consequently, some of the same economics that initially attracted development may become less compelling over time.

International expansion is not always a viable solution either. Moving AI infrastructure into jurisdictions such as Canada, Mexico, or other foreign markets introduces questions surrounding data sovereignty, national security, regulatory oversight, and legal jurisdiction. Governments and enterprises increasingly want sensitive data, AI models, and critical digital infrastructure to remain under domestic control, particularly as AI becomes more strategically important. Requirements related to data residency, cybersecurity, privacy, and sovereign AI initiatives can therefore limit the extent to which workloads can simply be relocated across borders in pursuit of lower costs or easier permitting.

Is this a Win for Behind-the-Meter?

This backdrop is generally supportive of Behind-the-Meter (BtM) and Power-as-a-Service (PaaS) models, though not without important limitations. By funding dedicated generation and related infrastructure, BtM can reduce concerns that data centers will shift grid-upgrade costs onto other ratepayers. It can also provide a meaningful speed-to-power advantage by reducing dependence on transmission upgrades and lengthy interconnection queues.

In the AI arms race, speed increasingly matters more than cost or emissions. A single gigawatt of AI capacity can potentially support tens of billions of dollars of annual revenue, making delayed deployment extraordinarily costly. While hyperscale campuses can typically be built in roughly two years, new utility-scale generation and transmission infrastructure often require substantially longer development timelines. As a result, the value proposition of BtM is increasingly centered on speed-to-power rather than simply the cost of electricity.

PaaS models may be particularly attractive because they allow hyperscalers to outsource power procurement, fuel management, generation operations, and reliability risk while preserving many of the benefits associated with dedicated infrastructure. Policymakers also appear increasingly receptive to these approaches. West Virginia's House Bill 2014 explicitly encourages high-impact data centers supported by dedicated generation and microgrids, while Texas Senate Bill 6 seeks to ensure large-load customers bear a greater share of the infrastructure costs associated with their demand. President Trump has similarly promoted "bring-your-own-power" approaches, arguing that data center developers should provide or procure the energy required to support their facilities rather than relying on existing ratepayers to fund new infrastructure.

However, BtM is not a panacea. Our sense is that hyperscalers still generally prefer grid-connected power where available given the economies of scale and diverse supply mix of the grid, both of which support the five nines of reliability (99.999% uptime) demanded by mission-critical workloads. By contrast, most BtM solutions remain heavily reliant on natural-gas generation and associated pipeline infrastructure, raising questions about whether they can economically achieve comparable levels of reliability. In our view, the incremental cost of eliminating fuel-supply and pipeline-disruption risks can be prohibitive, leading many BtM projects to target lower reliability thresholds.

BtM projects also remain exposed to permitting, fuel-supply, environmental, and community-acceptance risks. Developers must still secure air permits, pipeline capacity, fuel-supply agreements, and local support. The controversy surrounding xAI's Memphis facility illustrates how dedicated generation can reduce dependence on the grid while creating a different set of environmental, regulatory, and community-relations challenges.

Regulators may also become increasingly focused on ensuring BtM projects contribute to the broader costs associated with rapid infrastructure growth. Virginia's decision to tax both behind-the-meter and grid-supplied electricity for data centers illustrates this dynamic. The

state's rationale was that large data centers should contribute to the public costs associated with rapid infrastructure growth regardless of how electricity is procured, reflecting a broader policy focus on cost allocation rather than power-supply models.