Accenture Research Report

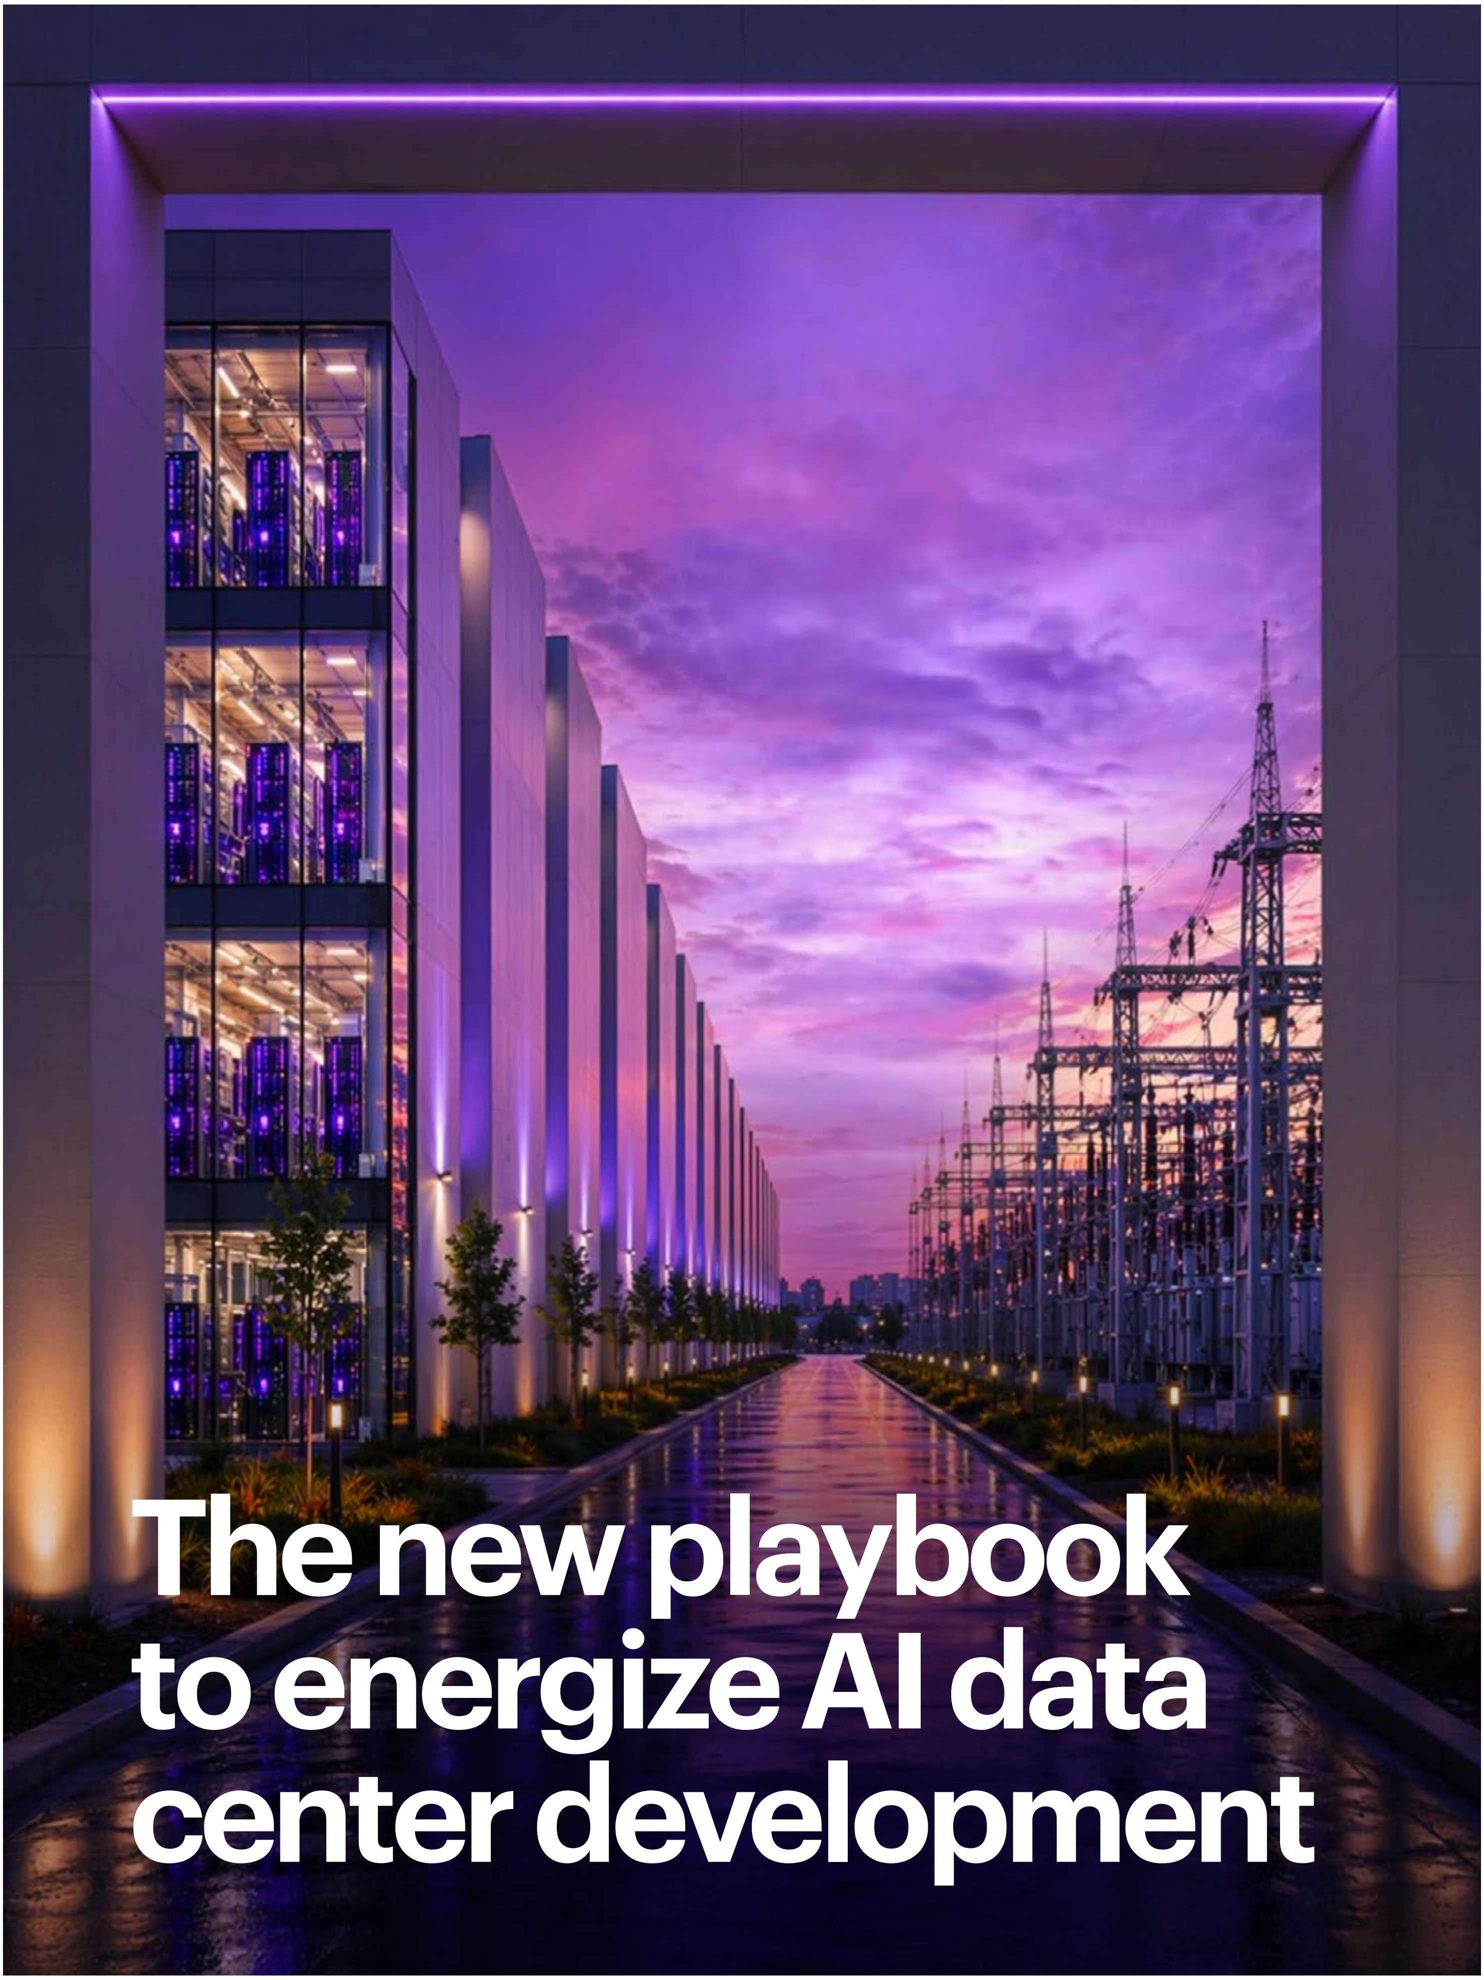

The new playbook to energize AI data center development

Authors

Neeraj Vadhan

Senior Managing Director,
Supply Chain & Engineering,
Data Center Services Lead

Neeraj leads the Data Center Services practice of Accenture, drawing on years of experience leading complex global technology and operations engagements for major clients in Silicon Valley and beyond. He holds four patents in applying artificial intelligence to large scale operational platforms.

David Quirk

Chief Executive Officer
DLB Associates

David is a visionary entrepreneur with over 25 years of experience. He is a licensed Professional Engineer in 49 states, Certified Energy Manager, LEED Accredited Professional and has chaired ASHRAE TC9.9. As co-owner of several companies, he leads 1,000+ professionals focused on innovation and transformation in the built environment.

James Hebblethwaite

Managing Director
Software & Platforms and
Industry & Enterprise, United
Kingdom and Ireland Lead

Along with leading Accenture's AI Infrastructure offering, James serves as UK & Ireland lead for Industry & Enterprise Reinvention, helping clients tackle the biggest challenges in this hypergrowth sector. He has deep expertise in B2B, digital infrastructure, operational excellence, supply chain and shared services.

Bethany Patton

Managing Director,
Industry & Enterprise, Utilities

Bethany has nearly 20 years of experience in power generation and utilities. Before Accenture, she worked as a mechanical design engineer for a large utility company, and in operations at a nuclear power facility, with experience in long-term capital project management and operations innovation.

Caroline Larose

Senior Manager,
Supply Chain
& Engineering

Caroline leads client work at the intersection of sustainability, energy transition and infrastructure transformation. In the data center sector, she helps operators solve power and infrastructure challenges, assessing stranded capacity, flexible load strategies and power productivity trade offs while balancing speed, resilience, cost and environmental performance.

Paul Johnson

Senior Manager
Software & Platforms
Research Lead

Paul specializes in research and analysis focusing on digital infrastructure, cloud and AI demand. He works closely with industry leaders and cross functional teams to shape perspectives on the future of compute and data center ecosystems, connecting technical innovation to commercial impact and long-term value.

Contents

Introduction:

The data centers that get power (and the ones that don't)

Picture this: Two data center developers submit separate power requests to their utilities on the same day. The first files a 500-megawatt (MW) request, a rough site map and a target go-live date. The second brings a documented load profile, a phased ramp schedule, a curtailment commitment and a proposal to co-invest in dedicated transmission. Alongside its power request, the second brings a clear community engagement plan that shows how it will account for local interests, address concerns directly and make credible commitments on community investment, job growth, responsible resource use and utility rate protection.

**Both want the same outcome.
Only one secures power on schedule.**

$1B+

estimated revenue deferred from a one-year delay on a 100 megawatt AI data center

For years, the data center industry has treated power as a transactional problem to solve after selecting a site. That approach no longer works. The projects moving forward today are operating in partnerships across developers, off-takers, utilities and communities to secure constrained power resources and earn community permission to build.
The cost of taking either power availability or local support for granted is significant. For an AI model provider, a single year of delay on a 100 MW facility defers more than US$1 billion in revenue. At the current market speed, deferred revenue quickly becomes lost revenue and lost market share.[1] As power and community constraints tighten, the US market is becoming an early signal of how data center development will need to change globally. The data center industry worldwide should pay attention.

Unlike chips, where capital and relationships can move buyers up the queue, power cannot be secured simply by spending more. It depends on physical infrastructure and close coordination across a complex system.

Now, with AI pushing demand to unprecedented levels, the race to secure as much capacity as possible is compounding the problem. Utilities and infrastructure providers simply cannot meet large capacity demands on the timelines data centers need. To secure power in the next two to three years, data centers need to change their approach.

Our analysis and client experience point to four moves developers can make before breaking ground to improve their path to energization. Winning developers will select potential sites based on power readiness, earn community permission to energize through community engagement, engineer flexibility into design and co-plan capacity with their utility partners.

With interconnection requests outpacing what utilities can realistically deliver in the next three to five years, and behind-the-meter (BTM) infrastructure also in short supply, the first data centers to come online will be those that are easiest to serve and deliver clear benefits to their communities. And the ones that will win in the market will maximize what they can deliver with the power they secure.

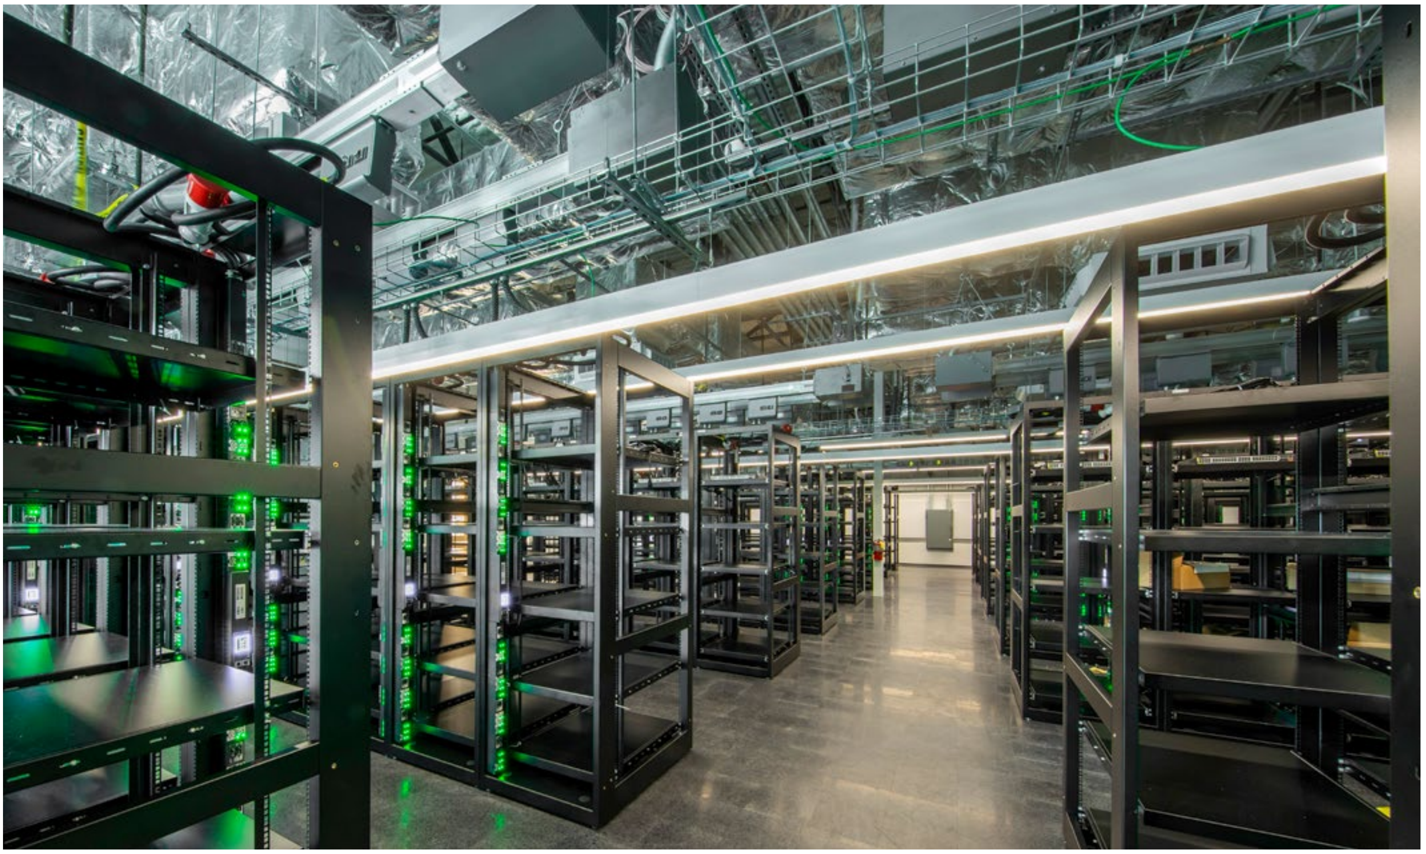

The cost of falling behind

Global CapEx investments for AI infrastructure are expected to reach US$1 trillion in 2027 and grow to nearly US$1.6 trillion by 2031.[2] At current build costs, a 100 MW facility runs between US$3.4 billion and US$5.5 billion.[3] Without power, this is a very, very expensive building.

> "Speed equals revenue. Whoever wins the speed race wins the revenue."[4]
>
> Sophie Smith,
> Vice President of the Americas at Soben (part of Accenture)

Speed requires both planning a path to power and getting all the key stakeholders aligned at the start.

For an AI model provider, long delays cede room for rivals to train and launch the next generation of models and accumulate the usage data that improves model quality over time, turning delay into permanent lost market share.

> For colocation operators, a one-year delay on a 100 MW facility defers approximately US$235 million in annualized lease revenue, and lease rates continue to climb at double-digit rates.[5]

Delay risk is not limited to monthly rent revenue. Prolonged delays can trigger missed financing milestones, broader financing strain, loss of project control through the exercise of step-in rights and, ultimately, the loss of committed tenants. The colocation tenants most likely to walk away are often those with the highest opportunity cost of delay, which can also mean the largest and most valuable leases.

The industry's instinct has been to move faster: file more interconnection requests, lock up more capacity and bring any BTM generation online as quickly as possible. But without a strategic approach, that urgency can create challenges of its own.

Community approval can make or break data center growth

The rush to development is turning communities into opponents rather than partners. Americans are increasingly well informed and vocal about their opposition to data center development in their local communities. In a 2026 poll, a resounding 71% opposed local construction of AI data centers.[6] In the past three years, US$85 billion of planned data center investments (53 projects) were canceled because of local opposition.[7] Several states have even proposed legislative restrictions on data center development.[8]

78%
of Americans said they worry that new data centers will raise their energy bills

Many are concerned that data centers will drive electricity price hikes, strain their water resources, generate noise, increase air pollution and cause associated health impacts. Communities have seen this first-hand. In the early waves of data center growth, some operators installed gas turbines equivalent to a small power plant, bringing in concentrated nitrogen oxides (NOx), particulate matter and jet-engine-level noise.

Those concerns are playing out in an already strained affordability environment. According to a nationally representative Consumer Reports survey, 78% of Americans said they worry that new data centers will raise their energy bills.[9] After roughly a decade of tracking inflation, residential electricity prices rose 10.5% from January to August 2025. Meanwhile, the share of households at risk of utility shutoffs because of unpaid bills rose 14% from 2024. [10 11]

These increases are driven in part by major grid infrastructure investment to support rising data center demand; a meaningful share of those costs ultimately flows through to ratepayers. At the same time, the promised economic benefits from local tax revenue and job growth have often fallen short or failed to materialize at all.

These challenges, considered together, show that community concerns are understandable. The full picture, however, is more complex. The same data center load that drives new infrastructure needs can also benefit the grid and the surrounding community.

Interconnection agreements that increase grid utilization spread fixed grid costs across more kilowatt-hours (kWh) and lower average rates. A 10% increase in grid utilization could reduce average residential rates by more than 3%.[12] Similarly, while many promises have fallen short, there is potential to grow local employment from data center development. While the quantity and quality of job growth and investment in workforce development and education vary significantly between projects, data centers providing long-term, local jobs. Six years after data centers are developed, local communities are seeing an average estimated 4-5% increase in jobs.[13]

50%

increase in public support for nuclear expansion in nine weeks

For developers looking to build a data center, securing power at scale and on time hinges on community buy-in. Communities are far more likely to support data centers when they are engaged early rather than hearing about a 500 MW facility the same day the interconnection request is filed. Engaging communities as stakeholders early, and backing plans with credible commitments, helps ensure facilities are seen as local assets and prevents long delays or canceled projects.

Trust in the AI and data center industry is strained. Perceptions of today's data center developments have been shaped by recent projects that overpromised benefits, underexplained resource impacts or appeared to shift costs onto local residents. The nuclear sector has faced similar trust and development concerns. Generation Atomic and Constellation Energy showed the power of targeted, fact-based community engagement through a public education campaign around nuclear energy, developed with Accenture. The campaign helped increase public support for nuclear expansion by 50% in nine weeks.[14] The same dynamic applies to data centers when operators credibly address concerns and demonstrate local value.

> Data center developers can use similar approaches to build local support, but only when public education is backed by credible commitments that address community concerns and demonstrate tangible local value. Those commitments require developers to change how they approach powering new data centers.

Today's data center demand is outpacing the grid

The compute required to train AI models is scaling at roughly four times per year. Even with hardware efficiency improving 26–40% annually, the net effective training power demand has been roughly doubling annually and is expected to continue to rise.[15] The data center industry is scrambling to capture as much power as it can to drive these models and the revenue-generating inference that follows.

To date, the industry has focused more on finding new sources of power than on rethinking demand itself.

Some of the largest players have pursued direct investments in generation via solutions such as BTM on-site, natural gas turbines to close the gap between what the grid can deliver today and what they need.
Early examples provided quicker workarounds. But direct generation no longer offers a faster route to power. Supply constraints and long build times have eroded that advantage, and these investments often still depend on some grid connection.

Neither the grid infrastructure nor planning processes were built for what's happening today. For nearly two decades, US electricity consumption growth was almost entirely offset by efficiency gains, growing just 0.2% per year and broadly distributed across the grid.[16] Now, national electricity use is expected to grow by 1% this year and by 3% in 2027.[17]

Further, AI infrastructure is concentrating enormous load in specific markets, hitting individual utility service territories with requests that dwarf their entire existing system peak. Many of these interconnection requests are real, however a large volume of speculative filings have also flooded grid systems with demand that never materializes, making utilities skeptical of every forecast, including the credible ones.

In Virginia, Dominion's interconnection queue has reached roughly 70 gigawatts (GW), nearly three times its all-time system peak.[18] The Electric Reliability Council of Texas (ERCOT) queue jumped from 63 GW to 226 GW in under a year, with 73% tied to data centers.[19] Utilities cannot respond at the speed the industry needs.

> The projected incremental US generation required through 2030 is expected to require CapEx investments of $58-89 billion, and even if capital is mobilized, money alone cannot compress a five-year interconnection queue, remove grid congestion or build a substation in under 18 months.[20]

Building against a moving power target

Data center power demand is shaped by chip technology, cooling design, workload mix and underlying software architecture. Step changes across these domains are outpacing data center build timelines. While these changes are driving efficiency, growth in industry power demand from increased intensity of workload and AI adoption is far exceeding any efficiency gains. At the facility level, these technology advancements and design changes can meaningfully shift power needs, making baseline power assumptions informing facility build outdated before the first rack is installed.

As the AI landscape matures, frontier model companies are adjusting their model strategies and investment priorities. As a result, while the aggregate AI power demand continues to climb, the power and power intensity for any individual data center is becoming harder to predict. Compute used to train top frontier models has grown roughly 400% per year since 2018, and – traditionally lower intensity – inference workloads can demand 15 times the per-query power of earlier large language models (LLM) generations.[21 22] OpenAI's Sora offers a recent example of how quickly AI workload economics can shift, with the video-generation app reportedly discontinued as high compute demands and operating costs pressured its business case.[23] That volatility makes point-in-time capacity planning for individual projects especially difficult.

Rack-level and model-level computing efficiency are improving fast enough to raise a fair question about whether power demand projections are overstated. For example, **Nvidia's upcoming Vera Rubin NVL72 configurations will provide roughly five times more computing performance per watt over its Blackwell counterpart.** This will drastically reduce the power (and cost) per token generated. Meanwhile, software training and deployment techniques such as those incorporated by DeepSeek showed that frontier-level performance may be achievable with less compute than many expected.

Together, these hardware and software efficiency gains may lower the ceiling of AI power demand from today's most aggressive projections, but they do not eliminate the demand for new facilities and the need to plan for rapid changes in power and thermal designs.

> Growth in AI model complexity, adoption and query volume continues to outpace what efficiency alone can absorb.

Four no-regret moves to start today

With more credible forecasting, greater design flexibility and earlier coordination across utilities, communities and operators, the industry can unlock capacity more effectively while building a stronger foundation for long-term growth.

The following no-regret recommendations apply before breaking any ground, and they remain relevant no matter where your business is building or how constrained the local grid is today.

01

Site for power readiness

For decades, the industry has prioritized factors like fiber connectivity or tax incentives, assuming that power could be procured for any chosen location. That assumption is now the single biggest point of failure in development. In today's resource-constrained environment, developers must define what they need and then follow the power. Speed to power starts with identifying where power capacity is readily available on a credible timeline. That means prioritizing sites with available grid headroom, reusable interconnection rights or infrastructure, manageable permitting risk and community dynamics that offer a realistic path to local support.

30-70% overbuild is required to achieve grid-level reliability from on-site gas

A proper power strategy is rarely a single-source answer. Most large AI data center developments will need to rely on a mix of power sources to match the planned ramp schedule and address various types of power needs, including near-term bridge power, back-up power, reliable firm power and dispatchable surge power. That mix typically spans available grid interconnection, existing generation, behind-the-meter resources (or microgrids), battery energy storage systems (BESS), thermal energy storage and longer-term power purchase agreements (PPAs) from new or expanded generation.

The fastest path to energization often runs through existing infrastructure, but the unique site context points to the right site-specific starting point. Adjacency to a nuclear facility may point toward a long-term PPA, proximity to a gas pipeline may support gas-fired BTM generation, and excess utility generation capacity may point toward colocation with an existing power generation facility. The best plans will look at the full menu of existing infrastructure pathways, including secondary-market acquisition of legacy interconnections via brownfields, colocation near existing generation, access to existing gas pipelines, as well as incorporating some level of grid interconnection and match availability to ramp schedule (Figure 1).

Regardless of power source, successful siting plans must have a long-term, reliable power solution in mind. The BTM gas solutions that recent data center developers have assembled under pressure are often insufficient over the long-term and can sometimes fail even as short-term solutions. While gas solutions offer firm power, limitations in speed to dispatch can conflict with quick surges in power demands for AI compute. Achieving grid-level reliability from on-site gas alone can require 30–70% overcapacity, pushing costs well above prevailing grid prices. The resulting localized emissions can also increase community opposition to future power demand or development. Without coordination with long-term interconnection plans or replacement generation, these bridge solutions can become permanent by default.

Recent work reinforces the need to start with available infrastructure. A large US energy company worked with Accenture to evaluate BTM generation opportunities for brownfield and greenfield sites by examining existing infrastructure, gas lines and capacity to expand. Similarly, DLB's work has increasingly required evaluation of brownfield retrofits to access available infrastructure, including a recent assessment of carrier hub sites across multiple US metros where existing power and permitting were the primary levers for speed to energization.

Permitting should be evaluated in parallel. Air permits for backup and on-site generation are the most common cause of late-stage delay. Brownfield sites often eliminate or reduce this risk by allowing use of existing permits or making it easier to secure them through established regulatory precedent and community familiarity.

The fastest projects will be those that align site selection, power sourcing and ramp strategy from the start.

Figure 1: Speed to power across procurement options

Power sources should be evaluated against a range of factors

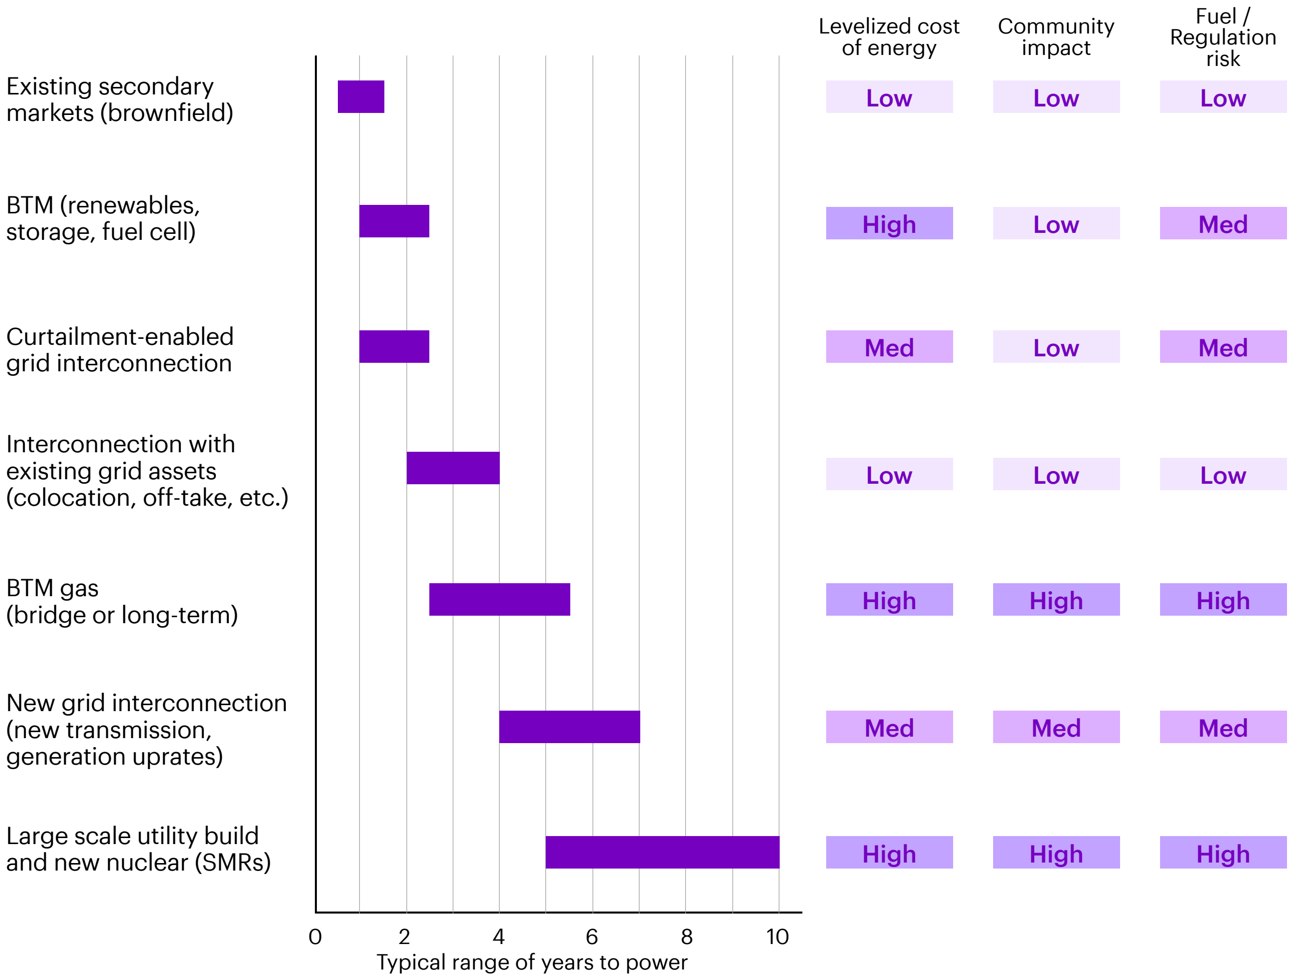

Source: Accenture research

02

Earn community permission

A data center is not truly power ready until the community is ready for it too. Data centers are large physical infrastructure projects with visible local impacts, including electricity demand, water use, noise, air emissions, land use and ratepayer concerns. In communities already facing affordability pressure, developers cannot assume that promised tax revenue or job creation will be enough to secure support.

Developers should build community engagement into the earliest stages of site selection, especially for sites identified as power ready. That means mapping local stakeholders before public announcements, engaging community organizations, residents, local officials and economic development leaders early, and creating a clear process to listen, respond and update plans as the project evolves. In a low-trust environment, this can be a critical differentiator between developers competing for similar available power.

Engagement must be backed by credible commitments such as those established through a community benefit agreement (CBA). Communities are more likely to support projects when developers can point to clear facility benefits and measurable actions to mitigate local impacts.

A CBA should spell out commitments on utility rate protection, water use, air emissions, noise, light pollution, traffic, land use and public health considerations. It should also define the local value the project will create, including realistic projections for tax revenue, construction and operating jobs, workforce training, education partnerships, infrastructure improvements, community investment and digital access. These commitments should be specific enough to be evaluated, tracked and incorporated into permitting, utility agreements and operating plans.

Community support must be earned before it is needed. The most effective developers will treat community engagement as part of the power strategy, not as a communications activity. A transparent engagement plan and credible CBA can help utilities, regulators and local leaders see that the project is easier to serve, easier to permit and more likely to deliver local value. It can also reduce the risk that opposition delays or cancels projects after major capital has already been committed.

03

Engineer flexibility into design

50-60%

of secured data center power capacity is stranded in traditional designs

Optimizing power use effectiveness is now table stakes. In a power-constrained market, developers need to focus on right-sizing what they request and optimizing how much useful compute they can produce from each available megawatt. Conventional designs unnecessarily strand 50-60% of secured utility capacity in headroom for rare peaks, redundancy and rigid system boundaries.[25]

Flexible design can adapt to a range of demand outcomes and speed power access by reducing the megawatts required to deliver the same compute. This can shrink bridge power needs, limit the scale of grid upgrades and shorten the path to energization. In parallel, flexible design can improve the data center's ability to adapt to final chip selection and the best-fit cooling technology for those chips.

Flexibility has two complementary pathways: facility-level design that makes power easier to share, shift and adapt to final chip and cooling technologies; and workload-level operations and software that determines which workloads can pause, move or run when power is available.

Facility-level flexibility starts with treating power, cooling and network as one integrated system. Operators should align rack density, cooling architecture, cabling and network topology across the facility, row, rack and server, while reducing avoidable conversion and distribution losses through shorter cable runs and, where appropriate, DC power distribution.

Pooled power infrastructure is a core design choice. By breaking down rigid electrical silos between IT and cooling, a shared power pool can capture non-coincident peaks, reduce stranded capacity and lower total provisioned load. Paired with thermal and battery storage, it can also shift demand across time, running chillers off-peak and freeing more power for IT when the grid is constrained. Verrus' work with the National Renewable Energy Laboratory (NREL) demonstrates this physical design approach, using power pooling, storage, controls and grid-interactive architecture to respond to utility signals while maintaining mission-critical operations.[26]

Designing facilities to be adaptable for a range of chip and cooling architectures is also key to right-sizing equipment load and power demand. Initial designs are often developed years before a data center comes online. Allowing for adjustment based on changes in AI strategy or available technology will avoid oversizing equipment and incremental stranded capacity from late-stage redesign.

> DLB's recent work with data center clients to develop flexible design for chips and cooling incorporated these principles by designing for both direct-to-chip liquid-cooling and multi-building compute commissioning spanning different rack densities and cooling architectures.

Workload-level flexibility adds the software layer. Training and batch jobs can often shift across time or geography, while inference and real-time workloads face tighter service-level limits. Operators should classify workloads by their ability to pause, defer, throttle or move, then use orchestration software to reallocate non-critical compute based on power availability. Emerald AI demonstrates this software-led model by using workload orchestration alongside NVIDIA technology to throttle non-critical loads, shift eligible compute and run secondary "filler" workloads when power is available.[27][28]

04

Co-plan capacity with utilities

The gap between nameplate capacity and actual operational load is preventing data centers from accessing grid power. Urgent requests for nameplate-based capacity overstate the likely operational load and require utilities to plan for firm power at any hour, which expands peak strain. These inflexible, large loads are nearly impossible for utilities to accommodate in fewer than three years. Financial investment alone cannot shorten the timeline required to build and expand physical grid infrastructure. Nameplate-driven loads can trigger a cascade of grid upgrades, new generation and coordination requirements that a right-sized, flexible request for the same facility could largely avoid or shrink.

Traditional firm interconnection structures force utilities to plan around peak demand that will rarely occur, if ever, as data centers routinely operate well below design limits.[29] A more flexible structure can materially change the path to energization. A commitment to curtail just 0.5% of annual load can make a facility much easier for the grid to serve and defer, minimize or avoid costly transmission upgrades.[30]

3-5

years faster to power with a flexible grid interconnection

Utility co-planning speeds power access by turning a broad capacity request into a staged, credible service plan, giving the utility the load certainty, flexibility parameters and upgrade sequence needed to energize the site sooner. Instead of treating every contracted megawatt as firm from day one, developers should separate firm, interruptible, curtailable and shiftable capacity in their utility agreements. That gives utilities a clearer view of when load will arrive, how much must be served at peak, what flexibility the site can provide and which upgrades are truly required.

A minimum curtailment agreement is already a requirement in some markets, such as ERCOT. These types of requirements are expected to be adopted in other markets to expand protections for residential load.[31] One US data center analysis found that pairing flexible capacity with bring-your-own generation or storage could enable full operation three to five years earlier than a traditional firm interconnection, while limiting curtailment to 40 to 70 hours per year.[32] Google's recent agreements with Indiana Michigan Power and Tennessee Valley Authority show this structure moving into utility contracts, building on successful demonstration with Omaha Public Power District where it reduced load during three grid events without undermining operations.[33]

Developers should also look beyond peak shaving. Utilities face different challenges across markets: some need peak-hour curtailment, others need backup capacity, frequency support, renewable integration or alternatives to transmission upgrades. Alternative co-investments and collaborative agreements should be assessed against expected impact on local electricity rates and resource use.

Aurora AI Factory, a collaboration between NVIDIA, Emerald AI, EPRI, Digital Realty and PJM, will be a first-of-its-kind facility integrating power flexibility into the core of its operations to speed interconnection ahead of network upgrades.[34]

Co-planning can also widen the infrastructure conversation. Utilities and adjacent infrastructure providers control or influence rights-of-way, easements, substations, fiber corridors, and water and gas service. Bringing these parties into one planning process can replace sequential negotiations with a coordinated buildout plan, reducing timing risk across the infrastructure needed to energize a data center.

The developers winning today's interconnection conversations are the ones who approach utilities as planning partners, looking for where the project and the local grid can benefit each other. That reframes the negotiation from a transaction over megawatts into a search for mutually beneficial outcomes. The ultimate collaborative agreements serve to demonstrate credible commitments and expected impacts in community engagement.

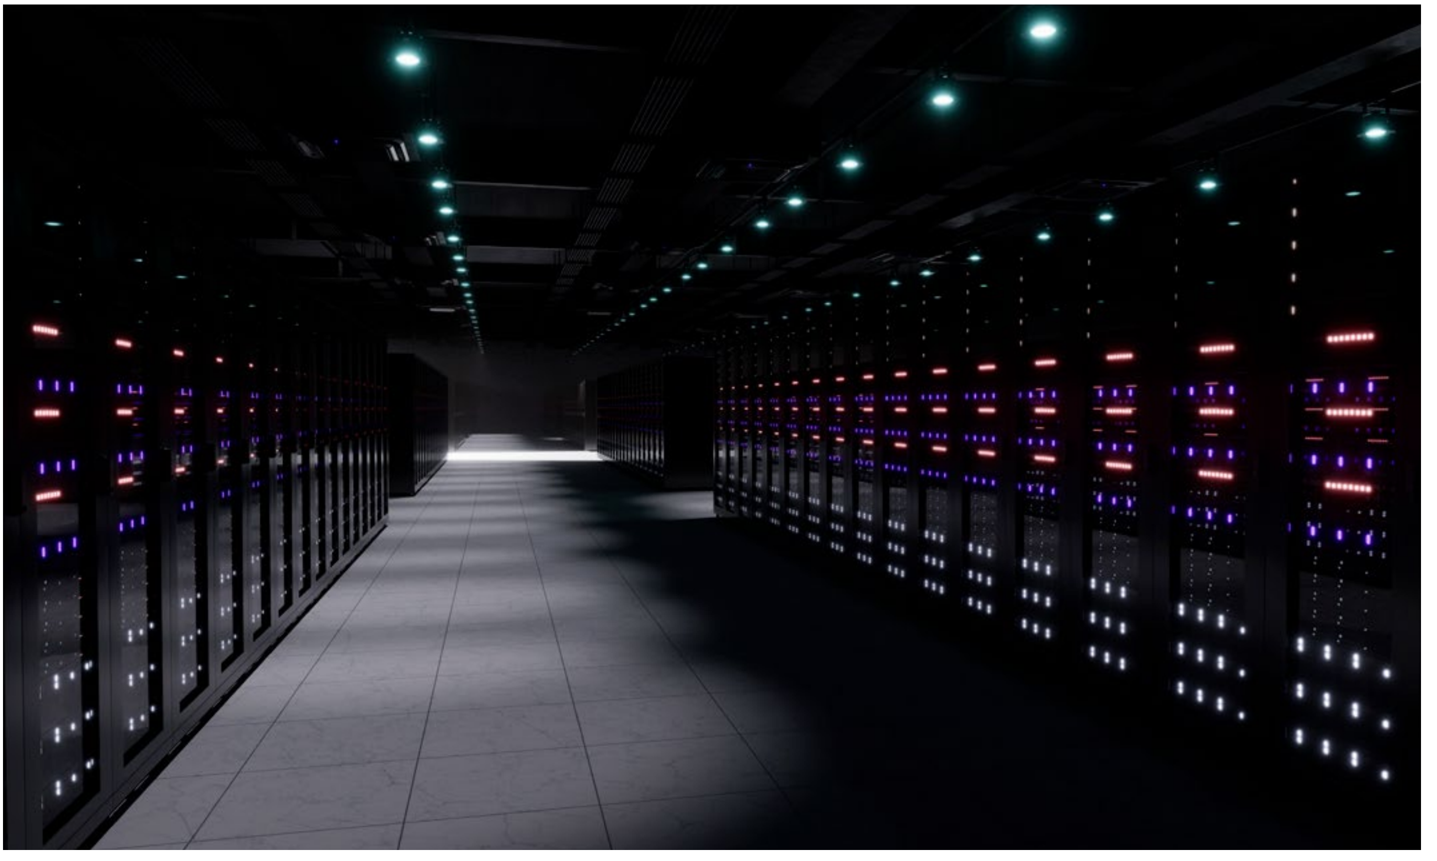

Putting strategy into action

Powering data centers today requires disciplined execution across the full development lifecycle, from design and planning to operations.

The four recommendations guide concrete actions from the earliest planning stages all the way to live service (Figure 2).

Figure 2: Immediate actions to drive recommendations across the build lifecycle

Actions are organized by phase so each no-regret recommendation becomes a set of tangible next steps.

Recommendation	Phase 1: Design & Planning	Phase 2: Procurement & Fit-Out	Phase 3: Service & Ops
Site for power readiness	• Audit potential sites against power availability • Select best-fit bridge power solutions • Develop long-term power supply strategy	• Lock long-lead power equipment • Phase fit-out and commissioning to power delivery	• Dispatch BTM/on-site generation within guardrails • Review capacity and gate the next ramp
Earn community permission to build and operate	• Identify key stakeholder groups to engage for each identified power-ready site • Assess community needs to understand core community concerns and priorities • Define preliminary benefit and mitigation commitments	• Formalize commitments through community benefit agreements • Define metrics for jobs, resource use and local investment • Establish feedback channels	• Report progress against commitments • Maintain local engagement and response channels • Update commitments as the facility ramps
Engineer flexibility into design	• Co-design power, cooling and network • Standardize cluster blueprint for workload portability • Design pooled power architecture to reduce losses and stranded capacity • Define the flexibility budget	• Stand up workload orchestration capability	• Tag and classify workloads and run orchestration plan to manage loads • Adjust cooling set points using real-time signals • Run the peak-event response operations plan • Track power productivity and flexibility performance
Co-plan capacity with utilities	• Begin collaborative utility conversations and long-range capacity planning • Define the flexibility budget • Lock long-term power supply strategy	• Finalize collaborative utility agreement and contract • Embed flexibility economics into commercial and operating agreements	• Run power partner coordination for peak-events and response execution • Enable partners to validate power productivity and flexibility performance • Review capacity and gate the next ramp

Source: Accenture research

The data center developers who get it right

The next phase of data center development will favor operators that can pair power access with a credible, flexible and executable plan for using constrained power productively.

As grid timelines lengthen and AI workloads continue to evolve, developers will need to move beyond broad capacity requests toward a more disciplined, power-first process. Developers will need to follow power from the start, earn community permission to energize, design facilities that can adapt as workloads and chip architectures evolve, co-plan capacity with utilities and match the power mix to the site and ramp schedule. These moves are mutually reinforcing. A more flexible design makes a utility agreement more credible. A better utility agreement gives communities more confidence that local impacts will be managed. A site selected for power readiness gives the whole project a faster and more realistic path to energization.

In a constrained power market, following this blueprint is what will separate projects that wait in the queue from projects that get built.

How Accenture can help

Accenture Data Center Services helps hyperscalers, enterprises and infrastructure developers build AI data centers faster and with greater confidence. The four recommendations in this report—following the power, earning community trust, designing for flexibility and co-planning with utilities—each require decisions that cut across engineering, community relations, power strategy and operations simultaneously.

What Accenture brings is integration. Data center development has traditionally been organized around functional silos: power, civil, mechanical, IT infrastructure and operations that are each optimized independently. Decisions made in one domain routinely create constraints in another. Accenture Data Center Services, working alongside DLB Associates and its network of affiliates, is structured to work across those boundaries, connecting power strategy, facility design, community engagement and operational planning into a coordinated effort optimized for the whole program.

Accenture brings a software-driven delivery model centered on a live digital twin that connects decisions and tracks program status across the full lifecycle, and DLB brings deep mission-critical engineering and delivery expertise across every phase, from concept through commissioning. Together we act as a single integrated partner from site selection and power origination through engineering, procurement and construction management (EPCm) to live operations and help clients move from capital to live capacity with greater speed, visibility and confidence.

About the research

We conducted in-depth interviews with senior Accenture leaders and practitioners working directly in data center development, power and grid advisory, facility engineering, utility negotiations, and AI infrastructure. These conversations were designed to test the core arguments in this report, pressure-check findings against client experience, and identify where the evidence was strongest.

We also analyzed publicly available data on utility interconnection queues, grid investment trends, regulatory filings and data center development programs across North America, Europe and Asia-Pacific. That analysis was further informed by Accenture's work with energy, infrastructure and technology clients and by the engineering experience of DLB Associates across mission-critical facility programs. All insights have been validated through triangulation across multiple data sources and expert inputs. Where we cite specific figures or external sources, these are noted in the references.

We use generative AI in our research production process. Our research experts review and validate the generative AI outputs with traditional research methods where possible, applying Accenture's Responsible AI standards.

Contributors

Kevan Yalowitz
Senior Managing Director
Software & Platforms, Global Lead

Bill Zhao
Managing Director
AI Lead for the Anthropic Business Group

Neil Chauhan
Chief Disruption & Differentiation Officer
DLB Associates

Bryan Colopy
Managing Director
Software & Platforms, Americas Lead

Andy Bhattacharyya
Senior Manager
High Tech

Michael Higgins
Principal Director
Utilities

Jason dePreaux
Senior Principal
Industry & Enterprise

References

1. MLQ.ai, Neocloud infrastructure: The race to power AI's compute needs, November 2025; Accenture research analysis.
2. Goldman Sachs, Tracking Trillions: The Assumptions Shaping the Scale of the AI Build-Out, May 1.
3. Alpha Matica, Deconstructing the data center: A look at the cost structure igniting the AI boom, November 13, 2025.
4. Accenture (Soben), Data centre trends 2026: Shifting up a gear, 2026.
5. JLLs 2026 Global Data Center Outlook, CBRE North America data center trends H1, 2025.
6. New American Journal, Americans oppose A.I. data centers in their area, May 13, 2026.
7. Heatmap News, Local opposition to data centers explodes in 2026, May 6, 2026.
8. The Wall Street Journal, These cities and states are taking aim at data centers, April 2026.
9. Consumer Reports, AI data centers: Big Tech's impact on electric bills, water, and more, March 20, 2026.
10. National Energy Assistance Directors Association (NEADA), Residential energy prices update, November 19, 2025.
11. National Energy Assistance Directors Association (NEADA), The cost of power: How soaring electric rates are deepening energy poverty in America, August 2025.
12. The Brattle Group, The untapped grid: How better utilization of the power system can improve energy affordability, March 2026.
13. Brookings, New evidence on data center employment effects, May 2026.
14. Accenture, Powered for Change 2025: Industrial decarbonization in the age of gen AI, 2025.
15. Electric Power Research Institute (EPRI), Scaling intelligence: The exponential growth of AI's power needs, August 2025.
16. US Energy Information Administration (EIA), Electricity end use and electric vehicle use data, 2005–2021, May 2026.
17. US Energy Information Administration (EIA), EIA forecasts strongest four-year growth in US electricity demand since 2000, fueled by data centers, January 13, 2026.
18. Virginia State Corporation Commission, Application of Virginia Electric and Power Company for approval of its large-load connection queue process standards, February 2, 2026.
19. Electric Reliability Council of Texas (ERCOT), System planning and weatherization update, December 2025.
20. Accenture, Powering the future of US data centers, 2025.
21. EPRI, Scaling intelligence: The exponential growth of data center power demand, 2025.
22. Oviedo, Felipe, et al., Energy use of AI inference: Efficiency pathways and test-time compute, September 24, 2025.
23. Jin, Berber and Jessica Toonkel, The sudden fall of OpenAI's most hyped product since ChatGPT, The Wall Street Journal, March 30, 2026.
24. Accenture, Powering Sustainable AI Balancing growth with environmental responsibility, 2025
25. Verrus, Beyond PUE: Rethinking data center efficiency for the era of power scarcity, January 15, 2026.
26. NREL, Vulcan test platform: Demonstrating the data center as a flexible grid asset, June 2025.
27. ENTSO-E, Data centres and the power system: Expected trends, challenges, and opportunities, May 2026.
28. Latitude Media, Nvidia and Oracle tapped this startup to flex a Phoenix data center, July 2025.
29. Shehabi, Arman; Newkirk, Alex; Smith, Sarah J.; Hubbard, Alex; Lei, Nuoa; Siddik, Md Abu Bakar; Holecek, Billie; Koomey, Jonathan; Masanet, Eric; Sartor, Dale, 2024 United States data center energy usage report, Lawrence Berkeley National Laboratory, December 2024.
30. Norris, T. H., Profeta, T., Patino-Echeverri, D., Cowie-Haskell, A. Rethinking load growth: Assessing the potential for integration of large flexible loads in US power systems, Duke University, 2025.
31. Norris et al., Rethinking load growth: Assessing the potential for integration of large flexible loads in US power systems, Duke University, 2025.
32. Brancucci, Carlo; Cutler, Dylan; Jenkins, Jesse. Flexible data centers: A faster, more affordable path to power. Camus, encoord, and Princeton ZERO Lab, December 2025.
33. Google, How we're making data centers more flexible to benefit power grids, 2025.
34. Latitude Media, How the world's first flexible AI factory will work in tandem with the grid, 2025.

About Accenture

Accenture helps the world's leading enterprises reinvent by building their digital core and unleashing the power of AI to create value at speed for organizations across industries. Our strategy is to be the reinvention partner of choice for our clients and lead in the safe, widespread adoption of AI, and to be the most client-focused, AI-enabled, great place to work in the world. We bring together the talent of our approximately 799,000 people with proprietary assets and platforms, deep process and industry expertise, and leading ecosystem relationships to deliver end-to-end solutions and measurable outcomes at scale. Through our Reinvention Services, we offer broad expertise across Cybersecurity, Digital Core, Finance, Industry and Enterprise, Song, Supply Chain and Engineering, and Talent, with advanced capabilities in AI and Data, Industry and Process, and Technology. We serve approximately 9,000 clients and generated approximately $70 billion in FY25 revenue.

Visit us at accenture.com

Accenture Research

Accenture Research creates thought leadership about the most pressing business issues organizations face. Combining innovative research techniques, such as data-science-led analysis, with a deep understanding of industry and technology, our team of 300+ researchers in 20 countries publishes hundreds of reports, articles and points of view every year. Our thought-provoking research developed with world-leading organizations helps our clients embrace change, create value and deliver on the power of technology and human ingenuity.

Visit us at accenture.com/research

About DLB Associates

DLB Associates is a global engineering and consulting firm with more than 40 years of expertise in mission-critical infrastructure—specializing in data centers, advanced technology environments and high-performance facilities for hyperscale, colocation, enterprise and telecom clients. The firm delivers integrated mechanical, electrical, controls, commissioning and technology solutions across the full project lifecycle. In January 2026, Accenture acquired a 65% majority stake in DLB and its affiliated companies, integrating the firm into Accenture to expand end-to-end AI data center capabilities.