KPMG

The new economics of power and compute

Why coordination is becoming more valuable than ownership

kpmg.com

Executive summary

A few weeks before this paper went to print, events caught up with its central argument. As a brutal heat wave pushed heat indexes past 110°F from Washington to Chicago, the US Department of Energy issued an emergency order authorizing PJM, the nation's largest grid operator, to direct data centers and other large loads to switch onto their own backup generation, a last resort ahead of rolling blackouts.[1] It was the third federal order of its kind in six months.[2] Coordinating energy and compute has become so critical that keeping it working now requires a cabinet secretary's signature. That coordination gap is the focus of this paper.

The first three papers in this series made the case for why utilities need to become energy orchestrators, the roles and operating model that shift requires, and the digital foundation that makes it possible[3]. This one asks a harder question, and one the market is answering with or without utilities: Who controls the coordination layer as artificial intelligence (AI) compute and energy begin to converge?

Two developments are reshaping the grid. AI inference is migrating toward the edge, to substations, regional infrastructure, and unused grid capacity that has sat idle for years, because that is where the economics make sense. At the same time, operators of the largest loads on the system are looking for ways to reduce their dependence on the grid through behind-the-meter generation and small modular reactors.

Together, these shifts point to a simple conclusion: Coordinating energy and compute is becoming more valuable than owning either. Utilities already hold what this new layer requires: regulated interconnection and dispatch authority; substation and grid-edge real estate; a century of balancing expertise; and the only trusted, regulated relationship with every ratepayer they serve. Most of these advantages cannot be replicated, regardless of how much capital technology companies are willing to invest. Utilities still treat coordination as a byproduct of running the grid, rather than a strategic asset.

This paper lays out three paths for utilities to capture that value—host, market-maker, and platform—and examines the technical, commercial, and regulatory trade-offs each entails. Standing still means leaving the coordination layer for others to build and control. The window to choose deliberately is measured in planning cycles, not decades.

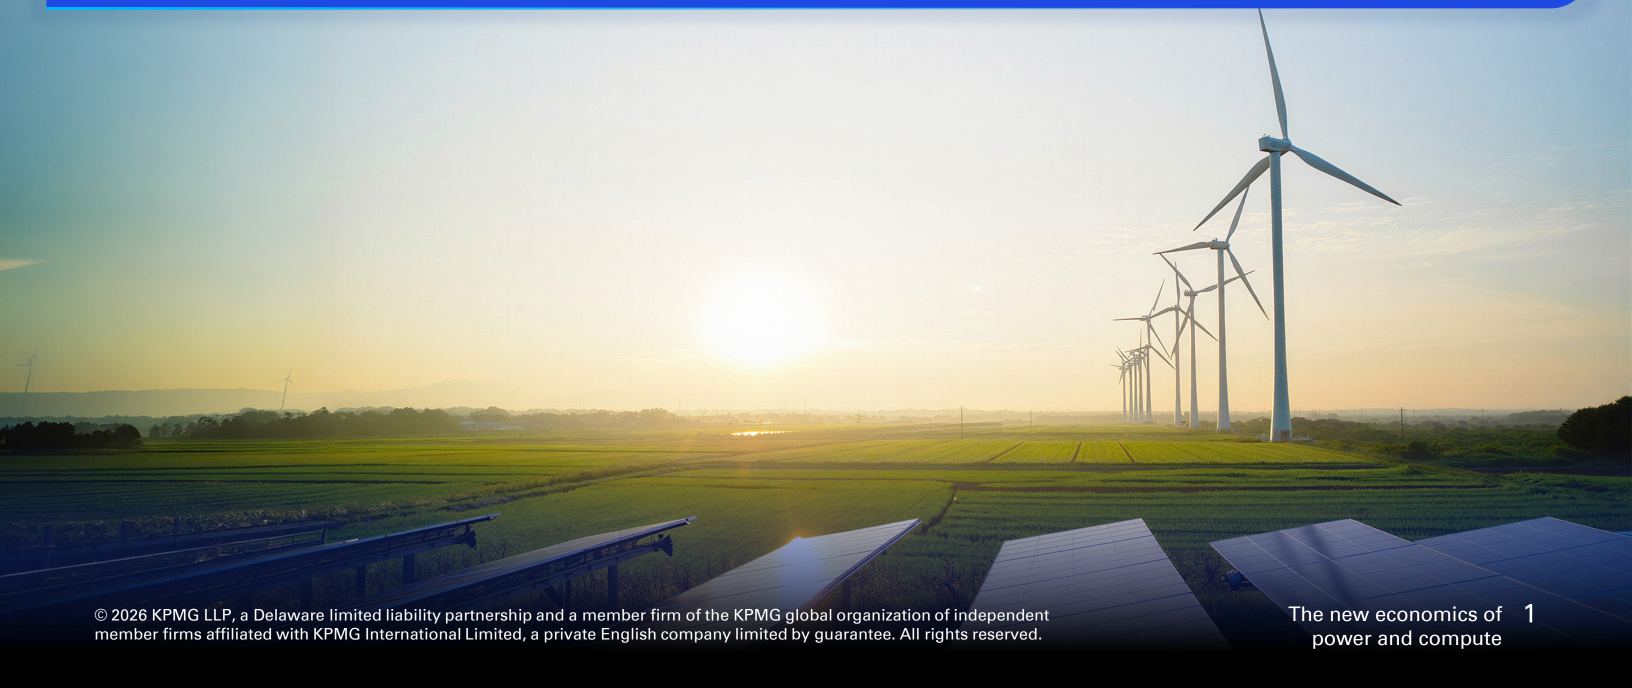

The coordination gap is already here

In late June 2026, a heat wave settled over the central and eastern United States and persisted through the July 4 weekend.

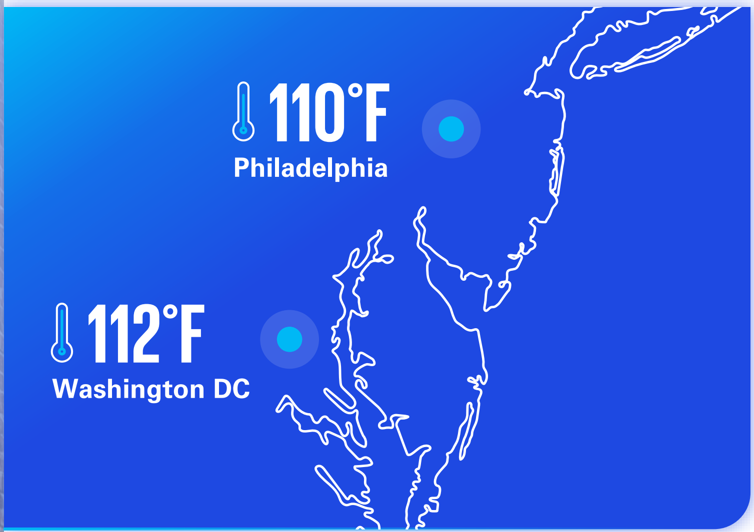

Heat indexes hit 110°F in Philadelphia and 112°F in Washington. Then the Department of Energy, acting on PJM's request under Section 202(c) of the Federal Power Act, authorized PJM to direct utilities to order data centers and other large customers of 50 megawatts or more to switch to their own backup generation within 15 minutes of an emergency signal. Hospitals, 911 centers, water treatment plants, air traffic control, and defense sites were exempt. Everyone else with a large load and a generator was on call. Never before in the AI era had backup generation been mobilized at this scale.

Leading up to that directive, demand on PJM, the grid operator serving 67 million people across 13 states and the District of Columbia, climbed to roughly 163 gigawatts at its peak, within about 2 percent of the all-time record of 165.6 gigawatts, a mark that had stood since 2006.[4] PJM recalled generators from maintenance, curtailed exports to neighboring grids, and activated emergency demand response.[5] Although the lights stayed on, no one was declaring victory. PJM's 2006 peak demand nearly fell, and the margin that saved it was assembled by hand, under federal emergency authority.[6]

PJM received similar emergency authorization during Winter Storm Fern in January and again during a May heat wave, when the grid operator expected reserves under 5.8 gigawatts against a peak load near 136 gigawatts.[7] Three federal emergency orders in six months made one thing clear: the backup generation and flexible load were available, but there was no established mechanism to coordinate them, short of federal emergency authority.

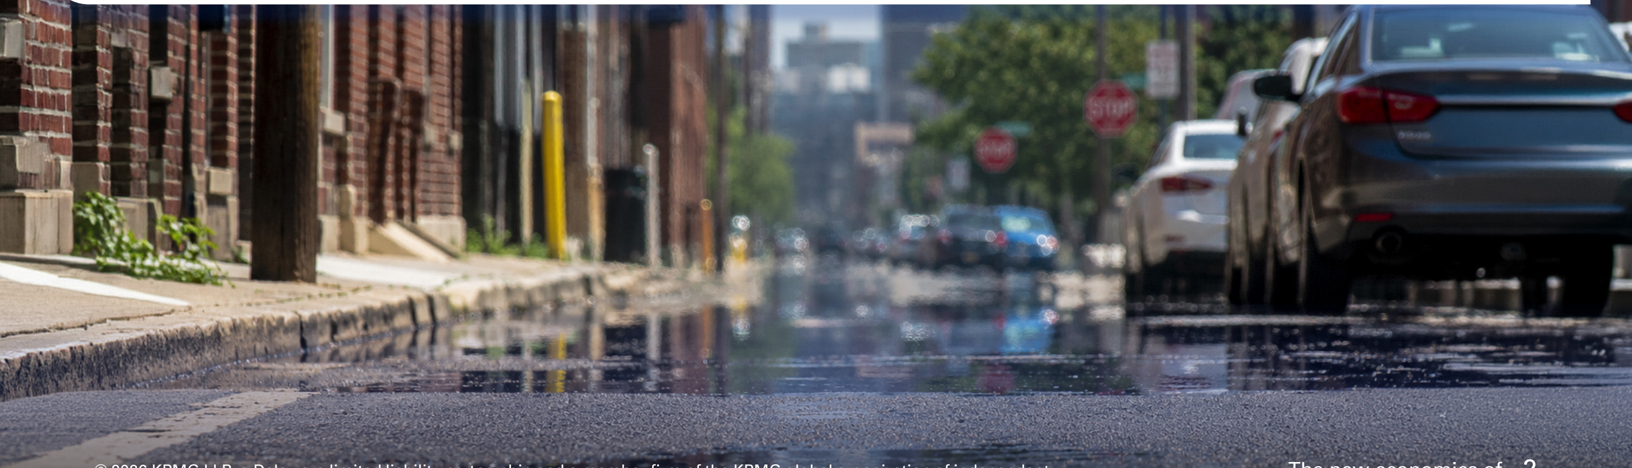

By the Department of Energy's own estimate:

>35 gigawatts

of backup generation **sit idle** across the country.

26 million homes

could be powered by this.

Almost none

of it had been tapped during a grid emergency, **until this year.**[8]

The need for that flexibility is only expected to grow. PJM forecasts point to roughly 32 gigawatts of peak demand growth by 2030, nearly all of it in data centers, while recent capacity auctions clear at prices roughly 11 times higher than two years ago.[9]

The solution used during the heat wave deserves scrutiny, too, because it was crude. Most backup capacity burns diesel. Virginia alone has permitted more than 8,000 diesel generators at data centers in recent years, and communities near those facilities spent the heat wave worrying about the impact on local air quality of firing them up.[10]

The reliance on emergency orders and diesel generation underscores the need for more durable coordination mechanisms.

> **Three emergency orders in six months isn't a coordination strategy. It's the consequence of not having one.**
>
> **Todd Durocher,**
> Principal, Power & Utilities Transformation Leader, KPMG US

Two shifts reshaping the grid

This summer's emergency orders exposed a problem that has been building for years. Two shifts in particular are reshaping the relationship between energy and compute.

Compute is moving closer to demand

For years, computing did not decentralize the way many expected it would.

Cloud platforms:

- **Scaled faster than expected**
- **Network costs stayed low**
- **Running a handful of massive hyperscale campuses turned out to be easier than managing dispersed infrastructure.**

The edge computing thesis showed up years before there was a workload that actually needed it.[11]

AI inference might finally be that workload, and part of the reason is the rise of smaller, inference-optimized AI models. The most advanced AI models are now more capable than many users and companies need, so competition has shifted from raw performance to efficiency, and in AI, efficiency is largely measured in watts. Distillation, quantization, and small models tuned for specific jobs keep shrinking the hardware that a given task needs, and every turn of that crank moves another class of workload closer to the user.

At NVIDIA's GTC 2026 conference, six major telecom and distributed cloud operators announced what NVIDIA calls "AI grids," distributed AI infrastructure built on existing network assets and even homes using proprietary Xfra Nodes. Telecom and distributed cloud providers run roughly:

100,000
distributed network data centers worldwide

>100 gigawatts
of **spare power** that could go toward new AI capacity[12]

Spectrum's deployment alone spans more than 1,000 edge data centers and hundreds of megawatts, within 10 milliseconds of 500 million devices.[13] Akamai has pushed its inference platform across more than 4,400 edge locations.[14] One application running on this infrastructure, Personal AI, gets sub-500-millisecond latency and cuts cost per token by more than 50 percent simply by moving language models off centralized infrastructure.[15] Another, Linker Vision, says it detects traffic accidents 10 times faster and coordinates disaster response 15 times faster than centralized processing allowed.[16]

Little of that capacity is housed inside a hyperscale campus. Much of it is closer to where utilities already operate. EPRI, the utility industry's own research institute, has noticed. In February 2026, EPRI announced a collaboration with NVIDIA, Prologis, and InfraPartners to pilot micro data centers, 5 to 20 megawatts, located at or near utility substations specifically to soak up capacity that is already available and otherwise going to waste.[17]

At least five US pilot sites are targeted by the end of 2026, built deliberately as a model other utilities can copy.[18] And the migration will not stop at the substation fence. For more on how AI and energy converge at the distribution edge, see "Energy Orchestrator, Part V: The grid's client-server moment."[19]

This paper focuses on the layer where the value of coordinating energy and compute is being created.

> **For the foreseeable future, the US Electric system needs every incremental megawatt of power that can be provided. Utilities have the opportunity to be commercial and innovative in how they help bring these megawatts online.**
>
> **Brad Stansberry,**
> Partner, Energy & Chemicals Advisory Leader, KPMG US

Large energy users are diversifying their power supply

At the same time, operators of the largest loads on the system are pursuing power strategies that reduce, though rarely eliminate, reliance on the grid. Small modular reactors are drawing serious interest from major technology companies partly because their design can support running independent of the broader grid.

Google's $10 billion investment in nuclear developer Kairos Power, aiming for 500 megawatts of carbon-free capacity by 2035, is one example.[20] NuScale Power, the only company with SMR designs certified by the Nuclear Regulatory Commission (NRC) (its original 50-megawatt design was certified in 2023; an uprated 77-megawatt design received Standard Design Approval in May 2025 after the fastest review in NRC history), describes its passive safety systems as enabling operation independent of grid power.[21] NuScale's chief executive has said the company is in talks with five "tier-one hyperscalers" he cannot name because of nondisclosure agreements.[22] Separately, at least one AI company brought large-scale on-site gas turbines online before it had secured a utility interconnection agreement, underscoring the extent to which some operators are prioritizing speed to power.[23]

Efforts to bypass the grid are real. Complete independence is not.

One of the largest cloud computing providers originally arranged to draw power directly from a Pennsylvania nuclear plant behind the meter, a structure that would have skipped grid oversight entirely for that load.[24] Two utilities, AEP Ohio and Exelon, formally objected at the Federal Energy Regulatory Commission (FERC), arguing the arrangement would push costs onto other ratepayers.[25] It worked. By mid-2025, that provider and the plant's owner signed a 17-year, roughly $18 billion agreement that explicitly moved the deal to a front-of-the-meter, grid-connected structure, ramping toward 1,920 megawatts of nuclear capacity by 2032 and running through 2042.[26] The deal ultimately moved in the opposite direction from where it began.[27]

Both shifts point to the same conclusion. The value will accrue to whoever can coordinate energy and compute across organizations, ownership structures, and physical locations. No one owns that role yet.

Why utilities still hold the advantage

The Susquehanna nuclear plant in Pennsylvania is a reminder that utilities retain more influence than many observers assume. A utility challenge to the original arrangement helped drive a fundamental restructuring of the deal. This summer sharpened the point. Under the emergency framework now being used repeatedly, large loads with their own generation stand first in line for curtailment, ahead of any interruption to homes and businesses.[28] As power constraints become a larger factor in siting and investment decisions, the value of an energy strategy extends beyond cost to include access, reliability, and resilience during periods of system stress.

The structural case goes well beyond one season. Many of the assets that make this emerging coordination layer valuable already sit with utilities, and few can be replicated easily by other market participants:

Regulated dispatch and balancing authority.

Utilities have spent a century coordinating electrons in real time across thousands of devices. The same capability now extends to compute.

Physical interconnection and substation real estate.

EPRI, NVIDIA, Prologis, and InfraPartners did not pick substations at random. They targeted sites where spare capacity and physical access already exist.[29]

Trusted, regulated relationships with every ratepayer they serve.

In March 2026, a coalition of major technology companies signed a voluntary "Ratepayer Protection Pledge" at the White House, agreeing to secure their own power and cover related grid upgrade costs.[30] The pledge reflects a broader recognition that public acceptance and regulatory support matter. Utilities remain central to both.

Statutory backing that keeps getting stronger.

In Texas, Senate Bill 6 now requires any load of 75 megawatts or larger to install equipment allowing it to be remotely disconnected during firm load shed events.[31] No customer negotiates around state law, and this one puts curtailment authority for the largest loads squarely in the grid operator's hands.

A demonstrated ability to set terms across utilities and jurisdictions.

At least 36 utilities nationally, including Dominion Energy, Arizona Public Service, Florida Power and Light, and AEP Ohio, have adopted large-load tariffs governing how data centers connect to and pay for grid capacity.[32] A review of 11 of those tariffs filed with FERC found utilities converging on similar terms: minimum demand commitments averaging 80 percent of contracted capacity, contract lengths averaging 14 years, exit fees averaging 5 years of minimum charges, and security deposits averaging 7 years of minimum charges.

Hyperscale technology companies need grid capacity, interconnection, and regulatory standing that only utilities provide. Utilities benefit from the capital, technology, and demand growth those companies bring. As AI-driven demand reshapes the power system, the relationship remains deeply interdependent.

Three paths to capture the value

No single strategy fits every utility. Differences in scale, regulatory environment, and risk appetite will shape which path makes the most sense. The options that follow are not sequential stages but distinct approaches, each with its own advantages and trade-offs.

Host	Market-maker	Platform

Host: Monetize grid-adjacent assets

The Host strategy builds on assets utilities already own and capabilities they already possess. By leasing grid-adjacent real estate and available capacity to AI inference operators—an approach being piloted by EPRI, NVIDIA, Prologis, and InfraPartners at utility substations [33] — utilities can participate in the growth of distributed AI infrastructure without fundamentally changing their operating model. It is the lowest risk of the three paths because it extends existing strengths in infrastructure, real estate, and interconnection rather than requiring new commercial or technical capabilities. The revenue stream can be durable, but it is largely passive: the utility is compensated for location, access, and connectivity rather than for coordinating or operating the energy-compute ecosystem itself.

Market-maker: Create and price new markets

The Market-maker path builds on a broader shift in the utility role: from delivering power to designing the markets, products, and commercial frameworks that surround it. Rather than simply connecting new demand, utilities help define the terms under which it enters and operates within the system.

Utilities are already pursuing this approach in different ways. AEP Ohio has focused on managing demand growth to protect system reliability.[34] Its PUCO-approved tariff requires new data center customers above 25 megawatts to commit to minimum demand charges equal to 85 percent of contracted capacity for a 12-year term.[35] The tariff was adopted after data center load in AEP Ohio's territory grew from 100 megawatts in 2020 to 600 megawatts in 2024, with more than 30,000 megawatts now representing inquiries and preliminary requests for service.[36]

Georgia Power has taken a different approach. In December 2025, the Georgia Public Service Commission approved 9,885 megawatts of new generation, approximately 80 percent of it tied to data center demand.[37] Under the arrangement, Georgia Power agreed to financially backstop the buildout through 2031 if the contracted load does not materialize, while large customers assume most of the associated costs.[38]

The mechanisms differ, but the underlying role is the same. In both cases, the utility is shaping how new demand enters the system, how risk is allocated, and how value is captured.

The emergency orders also exposed an opportunity. They demonstrated that significant flexibility already exists on both sides of the equation: more than 35 gigawatts of backup generation and a growing class of AI workloads capable of reducing power consumption when conditions require it.[39] What does not yet exist is a commercial mechanism for coordinating that flexibility at scale. That gap helps explain why accessing it required emergency federal authority under Section 202(c), three times in six months.

The research puts a large number on the prize. A Duke University Nicholas Institute study found that:

Trimming data center peak demand by just 1% to 2%

can cut electricity rates 0.5% to 2.8%

while protecting reliability,

and earlier Duke work found the existing US system could absorb about 76 gigawatts of new load if those loads were willing to reduce consumption for just a few hours each year.[40]

The market components are beginning to emerge. EPRI's DCFlex initiative includes a pilot in which Silicon Valley Power, a municipal utility serving 58 data centers within a 20-square-mile territory, can see an NVIDIA data center's load in real time, send a curtailment signal, and observe the response.[41] The effort is also informing new infrastructure designs. A 96-megawatt AI facility in Manassas, Virginia, designed for flexibility from the ground up, is planned for late 2026.[42]

The Market-maker approach turns emergency measures into market structures: interruptible-compute tariffs, bring-your-own-capacity interconnection, and priority access in exchange for binding curtailment rights.

Platform:
Coordinate energy and compute at scale

The Platform path is the most ambitious of the three. It extends the underlying digital foundation into a full coordination platform spanning grid operations, compute placement, and partner ecosystems.[43] Rather than participating in the coordination layer, the utility becomes the focal point for energy-compute coordination across its territory. Hyperscale technology companies, colocation providers, and distributed resource aggregators connect and coordinate through the utility's platform. This path offers the greatest potential value, but it also carries the highest level of complexity and the greatest exposure to the regulatory exposure.

	Host	Market-maker	Platform
Execution risk	Low	Moderate	High
Value ceiling	Modest, durable	Substantial	Highest
Core requirement	Real estate and interconnection	Commercial product design	Full digital and partner architecture

Each path involves a different balance of opportunity, complexity, and risk. Host builds on existing strengths and offers the most straightforward path to participation. Platform offers the broadest strategic role, but also demands the greatest investment and regulatory engagement. Many utilities will fall somewhere between those two poles. The appropriate choice depends on a utility's scale, capabilities, and objectives.

What the transition requires of utilities

The three paths described above differ in ambition, complexity, and risk, but none is simply a strategic choice. Success depends on a combination of technology, commercial capabilities, and regulatory alignment.

Technical

All three paths build on the SAP S/4HANA and Joule foundation: real-time intelligence, integrated process management, and ecosystem connectivity.[44] The platform path applies those capabilities beyond internal operations, extending coordination across energy resources, compute infrastructure, and external partners.

There is an important symmetry here: The technology driving this demand can also help coordinate it. Market-maker products depend on the ability to forecast large-load behavior, verify curtailment compliance through telemetry, and settle large volumes of flexibility transactions efficiently.

That same foundation can support market-facing functions as well. AI-enabled agents could verify curtailment events against meter data, evaluate flexible-interconnection requests against available hosting capacity, and help manage large-load queues at a scale that would be difficult to support manually.

For more on how AI can be applied to coordinate energy resources at the distribution edge and behind the meter, see the fifth paper in this series, The Energy Orchestrator series, part V: The grid's client-server moment.[45]

Commercial

All three paths need new commercial structures. Host relies on capacity-for-compute leasing arrangements. Market-maker depends on products such as flexible interconnection and new capacity offerings. Platform requires participation agreements that govern how ecosystem partners connect, share information, and transact.

None of these exist today as a standard utility tariff, making commercial innovation as important as the underlying technology. Large-load tariffs are one early example of this evolution, but they govern access to power, not yet access to coordination.

Regulatory

The opportunity is significant, but it is constrained. Utilities cannot simply choose to become unregulated technology platforms. Rate-of-return regulation, affiliate-transaction rules, and ring-fencing requirements limit the extent to which regulated utilities can use ratepayer-backed assets to support unregulated, competitive ventures.

These constraints will shape how utilities pursue each path. In many cases, the host and market-maker strategies align more naturally with traditional utility roles, while platform models may require clearer separation between regulated and competitive activities.

The regulatory landscape is evolving as policymakers respond to the growth of large loads and data center demand. In June 2026, a House Energy and Commerce subpanel began advancing the Ratepayer Protection Act, which would encourage states to adopt a “large load standard” requiring data center developers to fund the grid upgrades their demand creates.[46]

The regulatory framework for flexible load is still evolving. FERC has directed PJM to modify its rules to support more flexible data center interconnections.[47] At the same time, PJM’s independent market monitor has questioned whether flexible load arrangements provide sufficient certainty without binding curtailment commitments and has urged FERC to require large loads to bring matching generation.[48]

The outcome of that debate will depend on whether flexible load can be demonstrated, measured, and enforced in practice. Binding tariffs, telemetry, and verification will be central to that effort. Whatever the outcome of any single proceeding, the broader direction is clear: the regulatory framework for large loads and flexible demand is being shaped at both the federal and state levels, not just in utility boardrooms.

Regulatory sandboxes may provide a useful starting point for utilities pursuing platform ambitions, allowing new integration models to be tested before broader regulatory frameworks are established.[49]

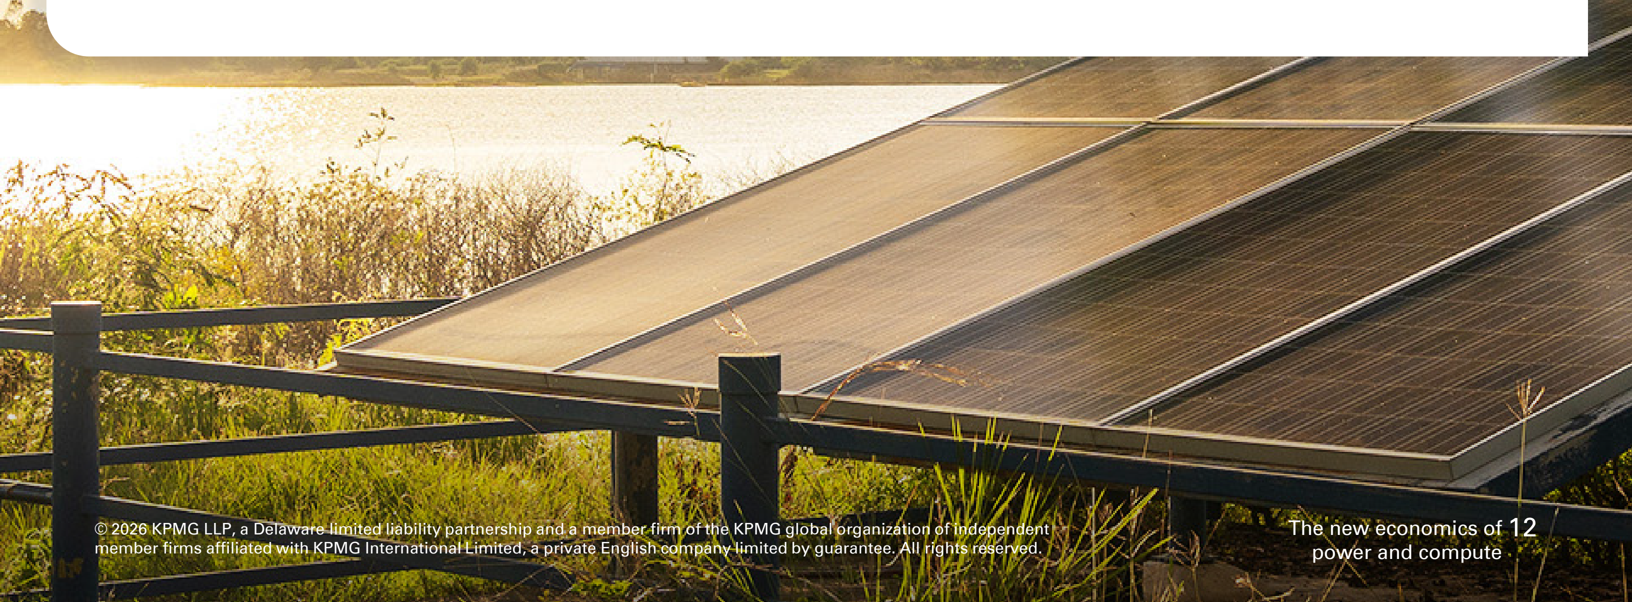

The cost of standing still

None of the three paths call for moving recklessly. All require deliberate action.

The alternative is drift and drift. In that scenario, the utility becomes the physical infrastructure beneath a coordination layer designed and controlled by others. Nothing makes that ending inevitable, but delays have a way of making decisions for you. This summer already showed what fills the gap in the meantime: emergency federal orders, diesel generators, a grid operator improvising under Section 202(c), and a demand record that survived by 2 percent. The country cannot manage the growth of AI through emergency measures indefinitely. Eventually, those resources will be coordinated through markets, platforms, or operating structures designed for that purpose.

The only question is who will build them.

How KPMG can help

As utilities navigate the complex transformation from Network Integrators to Energy Orchestrators, KPMG offers comprehensive support across the entire journey. Our deep industry expertise, combined with specialized capabilities in strategy, business transformation, AI, and data, positions us as the ideal partner to guide utilities through this critical evolution.

Strategic transformation support

KPMG's strategic advisors help utilities develop a clear vision and roadmap for their Energy Orchestrator transformation. We work with executive leadership to:

- Assess current capabilities against the Energy Orchestrator model requirements
- Define a strategic vision that balances innovation with reliability obligations
- Develop phased implementation plans that manage risk while accelerating transformation
- Create business cases and value realization frameworks for new platform business models
- Design ecosystem partnership strategies, particularly for technology company collaboration

Business transformation expertise

Becoming an Energy Orchestrator requires fundamental changes to operating models, organizational structures, and core processes. KPMG's transformation specialists help utilities:

- Design and implement new organizational structures aligned with Energy Orchestrator roles
- Develop change management strategies that build organizational capability and cultural readiness
- Transform regulatory engagement approaches to support new business models
- Redesign core processes to enable ecosystem orchestration rather than linear operations
- Implement performance metrics that measure platform success and ecosystem value creation

Data and technology enablement

The Energy Orchestrator model depends on robust data foundations and advanced technology architecture. KPMG's technology professionals deliver:

- Data strategy development that establishes data as a strategic utility asset
- Technology architecture design for platform integration, advanced analytics, and digital twins
- Implementation support for marketplace technologies and transactional platforms
- Cybersecurity frameworks to protect critical infrastructure while enabling ecosystem collaboration
- Cloud strategy development to support the scalability requirements of the Energy Orchestrator

Transformation journey support

KPMG recognizes that each utility's transformation journey is unique. We tailor our support to meet utilities wherever they are on their path to becoming Energy Orchestrators.

Phase 01 Foundation building	Phase 02 Capability scaling	Phase 03 Ecosystem orchestration
Capability assessment and gap analysis	Platform business model implementation	Full-scale platform deployment
Strategic roadmap development	Ecosystem partnership development	Sophisticated ecosystem value optimization
Initial use case implementation	Advanced analytics capability expansion	Advanced AI implementation
Technology architecture design	Regulatory framework evolution	Performance-based regulatory models
Data governance framework establishment	Organizational redesign and change management	Continuous innovation frameworks

By partnering with KPMG, utilities gain access to a powerful combination of industry knowledge, transformation experience, and technical expertise designed specifically to support the journey to becoming successful Energy Orchestrators in a rapidly evolving energy landscape.

Contact us

Todd Durocher
Principal, Power & Utilities Transformation Leader
KPMG US
tdurocher@kpmg.com

Todd Durocher is a leader in the Power & Utilities Advisory practice at KPMG with more than 25 years of experience in strategy, business transformation, and technology-enabled change. He advises utilities on large-scale transformation initiatives, including AI operating model modernization, cost optimization, grid modernization, acquisitions and integrations, and emerging technologies. Todd also leads KPMG's Energy SAP delivery in the US helping clients from strategizing and developing their program roadmaps, governance, business cases, and target state operating models through the implementation to post-go-live support.

Brad Stansberry
Partner, Energy & Chemicals Advisory Leader
KPMG US
bstansberry@kpmg.com

Brad Stansberry brings more than 25 years of experience delivering results on complex projects for finance function leaders in the energy and utility industry. He also leads the Energy & Chemicals Advisory business at KPMG. His focus is on helping chief financial officers and finance executives in the energy and chemical industries to run the "business of finance" better. This includes a focus on defining finance organization strategy and objectives, improving finance processes, deploying enabling technologies, enhancing finance talent and skills, and being a more valued service provider to their business constituents.

The authors would like to thank Leah Lockwood and Karen Henrie for their support and contributions to this paper.

Related thought leadership

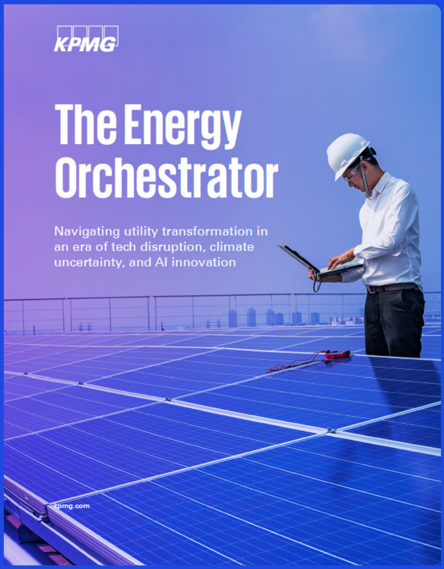

[The Energy Orchestrator: Redefining the Future of Utilities]

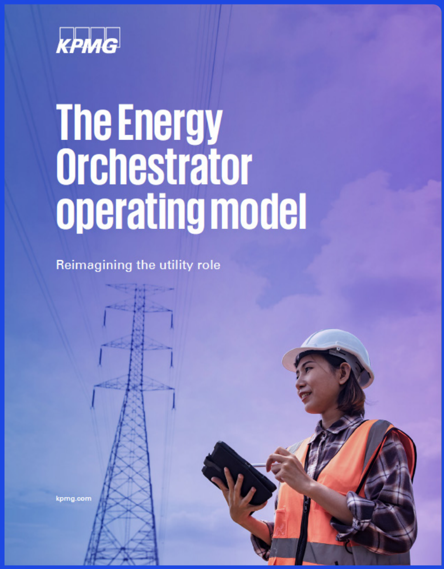

[The Energy Orchestrator: A New Operating Model]

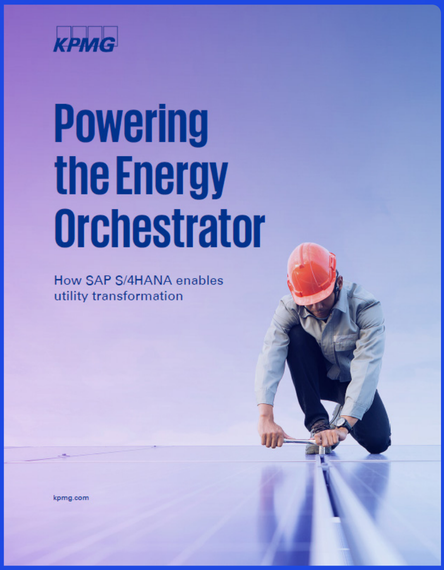

[Powering the Energy Orchestrator: How SAP S/4HANA Enables Utility Transformation]

Energy Orchestrator Part V: The grid's client-server moment

References

[1] US Department of Energy, Order Nos. 202-26-32 and 202-26-33 under Section 202(c) of the Federal Power Act, June 30, 2026; Ella Nilsen, "Energy Dept. directs data centers to use backup generators during heat wave, freeing up power for AC," CNN, July 2, 2026; Christine Condon, "PJM gets green light to push data centers onto back-up power during heat wave," Maryland Matters, June 30, 2026.

[2] Ethan Howland, "PJM gets emergency approval to curtail data centers, large loads during hot weather," Utility Dive, May 19, 2026; Ethan Howland, "PJM prepares to call on data center, large load backup generation to avoid blackouts," Utility Dive, January 30, 2026; "2026 DOE 202(c) Orders," US Department of Energy, accessed July 22, 2026.

[3] "The Energy Orchestrator: Redefining the Future of Utilities," KPMG LLP, 2025; "Inside the Energy Orchestrator: A New Operating Model," KPMG LLP, 2025; "Powering the Energy Orchestrator: How SAP S/4HANA Enables Utility Transformation," KPMG LLP, 2025.

[4] Health Department declares heath health emergency," City of Philadelphia, June 30, 2026; "A scorching stretch for D.C. begins" How hot will it get?," The Washington Post, July 2, 2026; PJM Interconnection, "Application for Presidential Permit of PJM Interconnection, L.L.C. (PP-33)," U.S. Department of Energy, accessed October 26, 2023.

[5] PJM Inside Lines, "PJM hot weather operations update," July 3, 2026, and June 30, 2026.

[6] Ibid.

[7] PJM Interconnection market data and load forecasts as reported by Reuters and industry press, July 1–3, 2026 (peak demand of approximately 163 gigawatts against the 2006 record of 165,563 megawatts; operating reserves; emergency demand response activation; projected peak demand growth of approximately 32 gigawatts by 2030; capacity auction price levels).

[8] Ethan Howland, "PJM gets emergency approval to curtail data centers, large loads during hot weather," Utility Dive, May 19, 2026; Ethan Howland, "PJM prepares to call on data center, large load backup generation to avoid blackouts," Utility Dive, January 30, 2026.

[9] US Department of Energy, Order Nos. 202-26-32 and 202-26-33 under Section 202(c) of the Federal Power Act, June 30, 2026; Bruce Gil, Gizmodo, "Energy department wants data centers to stop draining the grid during brutal heat wave," July 2, 2026.

[10] PJM Interconnection market data and load forecasts as reported by Reuters and industry press, July 1–3, 2026 (peak demand of approximately 163 gigawatts against the 2006 record of 165,563 megawatts; operating reserves; emergency demand response activation; projected peak demand growth of approximately 32 gigawatts by 2030; capacity auction price levels).

[11] Ariel Wittenberg and Benjamin Storrow, "'We are screwed': People near data centers dread heat wave pollution," E&E News by Politico, July 1, 2026.

[12] How AI Inference is Reshaping Edge Data Center Strategy," ValorC3, May 26, 2026.

[13] Kanika Atri, "NVIDIA, Telecom Leaders Build AI Grids to Optimize Inference on Distributed Networks," NVIDIA Blog, March 17, 2026.

[14] Ibid.

[15] Ibid.

[16] Ibid.

[17] Ibid.

[18] "EPRI, NVIDIA, Prologis, and InfraPartners Announce Collaboration to Produce Smaller Data Centers for the AI Era." EPRI press release, February 3, 2026.

[19] Ibid.

[20] 'The Energy Orchestrator series, part V: The grid's client-server moment," KPMG, 2026.

[21] "Google and Kairos power to partner to deploy 500 MW of clean electricity generation," Kairos Power, October 14, 2024.

[22] "Uprated NuScale SMR design gets US approval," World Nuclear News, May 30, 2025; NuScale Power press release, June 2025; "How SMRs gain design approval in the US," NuScale, accessed July 22, 2026.

[23] "US approves NuScale's bigger nuclear reactor design," Reuters, May 29, 2025.

[24] Chris Newman, "Data centers going off-grid with natural gas to 'find any way to get power'," NGI, June 16, 2025.

[25] Ethan Howland, "Talen to sell Amazon 1.9 GW from Susquehanna nuclear plant," Utility Dive, June 11, 2025.

[26] Ethan Howland, "Talen urges FERC to reject AEP-Exelon protest over interconnection pact for Amazon data center," Utility Dive, July 8, 2024.

[27] Talen Energy Corporation, Form 8-K, U.S. Securities and Exchange Commission, June 11, 2025.

[28] Talen Energy Corporation, Form 8-K, US Securities and Exchange Commission, June 11, 2025 (power purchase agreement restructuring, Susquehanna nuclear plant); Utility Dive and Data Center Dynamics, industry coverage, June 11–12, 2025.

[29] Zachary Skidmore, "PJM granted emergency approval to curtail data centers due to hot weather concerns," DCD, May 20, 2026.

[30] "EPRI, NVIDIA, Prologis, and InfraPartners Announce Collaboration to Produce Smaller Data Centers for the AI Era," EPRI press release, February 3, 2026.

[31] Garrett Downs, "Tech companies would have to pay AI data center energy costs under bill moving in Congress," CNBC, June 24, 2026; Nicole Greenfield, "AI Data Centers: Big Tech's Impact on Electric Bills, Water, and More," Consumer Reports, March 20, 2026.

[32] Maeve Allsup, "A reality check on flexible data centers," Latitude Media, January 20, 2026.

[33] "Utilities adopt large load tariffs to cope with the costs and power demands of data centers," EUCI, December 16, 2025.

[34] "EPRI, NVIDIA, Prologis, and InfraPartners Announce Collaboration to Produce Smaller Data Centers for the AI Era," EPRI press release, February 3, 2026.

[35] "AEP Ohio Proposal on Data Centers to Protect Ohio Consumers Adopted by PUCO," AEP Ohio press release, July 9, 2025; Sonal C. Patel, "Regulator Approves AEP Ohio's Landmark Data Center Tariff," POWER Magazine, July 10, 2025.

[36] Ibid.

[37] Ibid.

[38] "Data Center Fact Sheet," Georgia Public Service Commission, March 2026; Alander Rocha, "Georgia regulators approve massive power grid expansion to serve data centers," Georgia Recorder, December 19, 2025.

[39] Ibid.

[40] Chris Williams, et al., "Power-Flexible AI Data Centers: A New Paradigm for Grid-Responsive Compute,"arXiv preprint 2606.25098, June 23, 2026; Ella Nilsen, "Energy Department directs data centers to use backup generators during heat wave, freeing up power for AC," CNN Weather, July 2, 2026; 2026 DOE 202(c) Orders, US Department of Energy, July 20, 2026.

[41] Herman K. Trabish, “Data centers are ready to negotiate flexibility for speed,” Utility Dive, June 26, 2026; Zachary Skidmore, “Data centers could unlock 76GW of US grid capacity through optional curtailment,” DCD, February 18, 2025.

[42] Ibid.

[43] Ibid.

[44] “Powering the Energy Orchestrator,” KPMG LLP, 2025.

[45] Ibid.

[46] 'The Energy Orchestrator series, part V: The grid’s client-server moment," KPMG, 2026.

[47] Garrett Downs, “Tech companies would have to pay AI data center energy costs under bill moving in Congress,” CNBC, June 24, 2026; Nicole Greenfield, “AI Data Centers: Big Tech’s Impact on Electric Bills, Water, and More,” Consumer Reports, March 20, 2026.

[48] Maeve Allsup, “A reality check on flexible data centers,” Latitude Media, January 20, 2026.

[49] Ibid.

[50] “The Energy Orchestrator operating model,” KPMG LLP, 2025.

Some or all of the services described herein may not be permissible for KPMG audit clients and their affiliates or related entities.

Learn about us: | **kpmg.com**