SUSTAINABLE DATA CENTERS ROADMAP

CHAPTER 4

Accelerating Low-Carbon Power with AI Data Centers

Ayse Coskun, Varun Sivaram and Swasti Jain

October 2025

Innovation for Cool Earth Forum

4 Accelerating Low-Carbon Power with AI Data Centers

Ayse Coskun, Varun Sivaram and Swasti Jain

A. Introduction: The Double-Edged Sword of AI's Energy Consumption 2
B. Mechanism 1: Advanced Market Commitments as a Catalyst for Clean-Firm Power 8
C. Mechanism 2: The Flexible Data Center: Transforming Power Grids 10
D. Mechanism 3: Siting for Sustainability 17
E. Mechanism 4: AI as the Architect of a Clean Energy Future 20
F. Recommendations 22
G. References 24

A. Introduction: The Double-Edged Sword of AI's Energy Consumption

i. Where AI meets the grid

Over the past decade, the proliferation of digital technologies has quietly reshaped global electricity demand. Now, with the exponential growth of artificial intelligence (AI), the digital sector—specifically data centers powering AI workloads—is rapidly emerging as one of the most transformative forces in global energy systems. As deployment accelerates across sectors, from finance and healthcare to manufacturing, hyperscale AI models are triggering surging demand for computational power and, by extension, electricity. Behind every model training run and wave of new applications sits a physical footprint of racks, wires, cooling loops and a rising, highly concentrated draw on the grid. Globally, data center electricity use is now about 1.5% of demand and, in the International Energy Agency's (IEA) Base Case Scenario, doubles to ~3% by 2030. Much of this load is clustered in areas where the scale of data center electricity demand strains local power adequacy and complicates long-term grid planning.[1] (See Chapter 1(C) of this Roadmap.) According to the IEA's Energy and AI Observatory, notable clusters include Loudoun County in Virginia, Dublin in Ireland, Singapore, Tokyo and the Amsterdam–Frankfurt–London corridor (see Figure 4-1).[2] While many of these clusters are urban or metropolitan hubs, large AI campuses are rising in more rural or remote regions where the relative strain on smaller grids can be even more pronounced. In places like rural Grant County in Washington (home to a cluster of hyperscale facilities around Quincy), the public utility has capped data center load growth because demand is outpacing local transmission capacity, underscoring how concentrated build-outs can strain smaller grids.[3]

Figure 4-1. *International Energy Agency (IEA) analysis of data center locations, powerlines and submarine fiber optic network data. The data center hubs shown below represent the capacity-weighted centroid of a cluster of data centers within 100 km of each other and total over 500 MW of installed capacity.*

What fills this near-term gap is not pre-ordained. Some planners lean on familiar dispatchable options, such as natural gas or coal-fired power plants, to guarantee reliability. Yet recent experience tells a more nuanced story. In the United States, solar and utility-scale batteries led new capacity additions in 2024. This shift signals multiple viable pathways to serve new loads and create workforce and institutional capacity to deploy these systems more quickly and at lower cost.[2,4] While capacity additions in 2023 were dominated by wind and solar, forward-looking integrated resource plans (IRPs) in high-growth states now include substantial natural gas capacity. This strategy is driven in part by differing views on the ability of firmed renewables, including wind and solar paired with batteries, to serve load with similar reliability to fully dispatchable generation. Crucially, evidence from real systems shows that a clean, flexible portfolio performs well under stress, successfully contributing to meeting reliability requirements. In this context, flexible portfolios are combinations of energy resources that adapt output or consumption to the grid's needs. They pair wind and solar with battery storage, and demand response as variable renewables grow. For example, during the Electric Reliability Council of Texas' (ERCOT's) record-hot summer of 2024, mid-day solar met a large share of load, and evening battery discharge bridged the ramp as the sun set. This maintained reliability and avoided conservation appeals or load shedding even as new demand records were set. This example serves as an operational proof point that firm service can emerge from the combination of renewables, storage and market design rather than from one resource alone.[5]

Globally, the IEA forecasts that renewables (primarily wind, solar and hydro) will be used to meet nearly 50% of the additional electricity demand from data centers through 2030, while natural gas and coal together will supply over 40%.[6] Specifically, in the United States, much of the near-term capacity for announced mega-campuses will be powered by fossil fuels. For example, Meta's $10 billion Louisiana campus is slated to be served by multiple new gas-fired plants, while the planned Stargate AI campus in Abilene, Texas includes a 360 MW on-site natural gas facility.[7] In contrast, markets like Canada, Japan and Brazil are responding to compute demand with renewable-forward infrastructure planning.[8]

With transparent forecasting and grid modernization investments, jurisdictions can meet rising compute loads while bending the supply mix toward low-carbon and flexible portfolios, rather than locking in outcomes by default.[9]

ii. The opportunity

On the ground, developers are bringing data centers online wherever firm power is available—even as they aim to run on renewables or nuclear later—because clean additions and transmission take time, and near-term supply remains largely fossil-based as a result.[6] In the United States, interconnection timelines have stretched from under two years in the 2000s to over four years recently, with a median of five years for projects completed in 2023, creating a real pace gap between compute build-out and a fully clean and reliable grid.[10]

This chapter advances a central thesis: the immense and concentrated power demand of AI data centers could be a primary driver to accelerate the clean energy transition. By proactively aligning AI computational use and infrastructure growth with clean energy deployment, stakeholders can unlock scale, financing and innovation in four key ways:

1. Advanced Market Commitments (AMCs) for Clean-Firm Power: AMCs are purchase agreements intended to de-risk an innovative technology. AMCs differ from power purchase agreements (PPAs), which are long-term commitments to buy power regardless of whether the generation source is innovative. This distinction matters as compute demand is outpacing clean capacity in many regions where availability-based offtake can de-risk and scale a broader portfolio. In this context, clean firm power refers to low-carbon resources that can generate electricity, independent of weather, such as enhanced geothermal, advanced nuclear, carbon capture and storage (CCS). Clean firm power also includes green hydrogen to complement the large renewable PPAs that hyperscalers (i.e., ≥50 MW at a site or ≥100 MW portfolio) already sign.[11] Even so, near-term reliability gaps remain: a planned 1.8 GW AI campus near Cheyenne, Wyoming will use dedicated natural gas with proposed carbon capture to meet immediate needs—underscoring why AMC-style deals that pull forward new clean-firm projects are essential to avoid long-term fossil lock-in.[12]

2. Demand Flexibility as a Grid Asset: Demand flexibility refers to the demand side, such as a data center, adjusting its power use in response to grid needs while maintaining quality of service (QoS). Many data centers today behave as always-on loads (e.g., Google and Meta fleet data show diurnal power swings of only approximately 4% on average).[13] Workloads like AI training and other batch analytics have hours-to-days completion windows. This enables temporal flexibility, shifting non-urgent workloads to another time, such as overnight or during demand-response events (see Figure 4-2). It also enables spatial flexibility, shifting workloads across regions or sites, to follow abundant low-carbon power and avoid local grid stress.[14] For example, Google's carbon-intelligent platform already shifts non-urgent compute to lower-carbon hours and has been used for targeted demand response during grid stress.[15]

Figure 4-2. *Demand flexibility means shifting delay-tolerant workloads in time, shifting power consumption to on-site resources like batteries, and shifting workloads or siting across regions to follow abundant, low-cost clean energy. (Source: Rocky Mountain Institute (RMI))[14]*

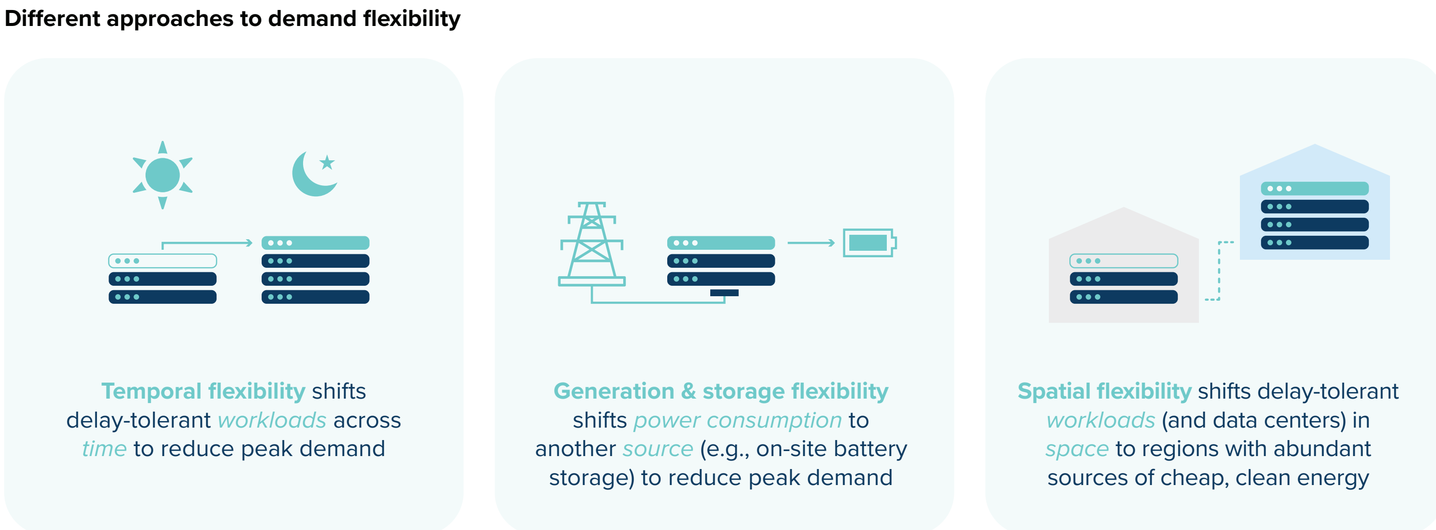

3. Siting for Sustainability: Siting refers to where new capacity is built; this is distinct from real-time workload shifting within the current data center footprint. In principle, latency-tolerant workloads (e.g., AI training, batch analytics) can be sited in renewable-rich regions, but in practice, many mega-campuses go where power is available fastest. Thus, near-term supply often tends to be fossil-based. Given multi-year interconnection timelines and constraints on moving data, near-term siting should prioritize grids with headroom and curtailment, pair builds with storage and use availability-based clean offtake to pull forward new clean-firm capacity.

4. AI as the Architect of a Clean Energy Future: Machine learning can optimize grid operations, forecast renewable output and accelerate materials discovery for next-generation energy technologies. Focusing AI applications on opportunities and challenges can reduce costs, environmental stresses and speed to market.[16]

This transformation of aligning the AI compute build-out with clean, reliable power will not occur automatically. It will require knowledge-sharing across firms and utilities by publishing anonymized load-shape datasets and sharing procurement playbooks for flexible, clean capacity. Yet the window of opportunity is now: as trillions of dollars flow into AI infrastructure, strategic intervention can shape the energy future at a global scale.

iii. Global lens

AI data center growth and its energy implications occur in a global context, creating new opportunities. However, most capacity remains centered in the United States, given that approximately 54% of global hyperscale data center capacity and North American build has clustered in Northern Virginia, Phoenix, Dallas–Fort Worth and Atlanta.[17] Northern Virginia alone has approximately 4-5 GW of current data center load, with growth outpacing generation and transmission additions.[18] In parallel, near-term supply for several marquee US campuses is being met with fossil dispatch or planned new gas capacity, such as Project Stargate's on-site ~360 MW gas plant in Abilene, Texas.[19]

Against this backdrop, material projects and policies are advancing outside the United States. In Japan, the Ministry of Economy, Trade and Industry (METI) has launched public-private partnerships to co-develop AI data centers with integrated offshore wind and floating solar infrastructure.[20] In Brazil, clean-power strategies for data centers are being led by renewable PPAs, signed by companies such as Scala, while true hydroelectric co-location exists at Itaipu's utility-campus data center.[21] Malaysia's state of Johor houses the Green Data Center Park, powered by 500 MW of on-site solar, positioning Southeast Asia as a competitive region for clean information technology (IT) growth.[22] In Canada, hydro-rich provinces—especially Québec and British Columbia—offer low-carbon power for AI-class data centers and have recently tightened and streamlined how large new loads connect. Québec's 2025 energy reform created a ministerial authorization regime for large connections, while British Columbia enacted a streamlined permitting law, opened consecutive clean-power procurements, and kept a pause on new crypto-mining interconnections to reserve capacity for strategic industrial customers.[23,24]

Even in power-constrained regions like sub-Saharan Africa, the possibility of AI data center hubs is driving renewed interest in distributed solar and grid modernization. However, challenges related to reliability, financing and workforce capacity remain significant. For example, Microsoft and G42 announced a $1 billion investment in data centers in Kenya designed to run on geothermal: an example of pairing latency-tolerant AI with a firm, low-carbon supply at the source.[25] In Namibia, the European Union's Global Gateway is backing giga-scale renewables for green hydrogen and even green iron—an "energy-at-the-resource" model that data center developers could emulate.[26] These global projects underscore that data center planning and investment are starting to influence national energy planning well beyond the traditional technology hub geographies.

B. Mechanism 1: Advanced Market Commitments as a Catalyst for Clean-Firm Power

i. The power of the purse: how hyperscalers move markets

Large hyperscalers have used PPAs for many years to expand deployment of solar and wind by agreeing to buy power for 10-20 years at predictable terms (long-tenor offtakes) to enable project financing. These hyperscalers are now the largest corporate buyers of renewables.[27] PPA prices generally clear near market rates, so pricing is relatively well established. Sometimes the PPAs also incorporate a green premium, which is the additional cost of a choosing cleaner technology over one that emits more greenhouse gases. More recently, hyperscalers have begun using AMCs, usually tied to technology maturation or performance demonstrations to signal demand for electricity from novel technologies. However, AMCs are more speculative, with limited precedent on price points and timelines. Thus, they function as pioneer instruments to support technology development, rather than commoditized instruments simply intended to secure power.

A visible example is Google's multi-year offtake with Fervo Energy, announced in May 2021, which has allowed Fervo Energy to develop a next-generation geothermal project. Fervo began delivering carbon-free electricity to NV Energy's system in late 2023.[28] Further, after Nevada regulators approved a Clean Transition Tariff in May 2025, Fervo is slated to deliver 115 MW of round-the-clock enhanced geothermal to NV Energy's grid, which NV Energy will sell to Google at a set rate. With deals such as this one, 2024 set a global record for corporate clean-energy deals.[11,29]

In parallel, for advanced geothermal, federal analyses and new resource assessments show a large upside if progress continues. The Department of Energy's (DOE's) Enhanced Geothermal Shot targets $45/MWh by 2035, and the US Geological Survey (USGS) estimates the potential of Enhanced Geothermal Systems (EGS) is approximately 135 GW in the Great Basin upper crust to 6 km.[30] Independent modeling finds the national potential is more than 60-100 GW by mid-century with technology advances.[31] Even with this potential, significant hurdles remain: costs remain high, long-term performance is unproven at scale, seismic risks must be managed and permitting risk still hinders financing. As a result, AMCs like the Google-Fervo deal demonstrate greater bankability.

ii. Case study: nuclear renaissance

Small modular reactors (SMRs) are fission technologies that follow nuclear-plant licensing. They aim to gain cost and schedule control through standardization and factory fabrication. DOE has an active SMR support and solicitation pathway to progress near-term designs. Fusion, by contrast, operates under a different US regulatory path. For example, the Nuclear Regulatory Commission (NRC) placed fusion under the byproduct-materials framework rather than power-reactor rules, potentially enabling faster first-plant timelines.[32] Recent corporate offtakes illustrate both lanes. In the case of fission, Microsoft-Constellation's 20-year PPA to restart Three Mile Island Unit 1 was announced September 20, 2024. Constellation is targeting a 2028 return to service, pending approvals and execution (see Figure 4-3).[33-35] In fusion, Microsoft and Helion announced their agreement on May 10, 2023, aiming for an initial ~50 MW of power generation in 2028. Google and Commonwealth Fusion Systems (CFS) announced their 200 MW offtake on June 30, 2025. This offtake targets early-2030s power from the CFS plant in Virginia.[36,37] Of note, these fusion and restart deals are forward-dated commitments rather than active deliveries, whereas the Google-Fervo project is already online and scaling under the approved tariff.

Figure 4-3. *Three Mile Island Nuclear Power Plant. (Source: American Nuclear Society)*[35]

iii. Scaling pathways and spillover effects beyond data centers

AMCs can launch first-of-a-kind (FOAK) projects. However, the doubling rate of installed capacity is what primarily drives cost declines (via learning-by-doing, supply-chain maturation and finance de-risking). Thus, spillover benefits only materialize when deployments repeat and cumulative capacity grows. In modular clean technologies, meta-analyses consistently find experience-curve cost declines.

Specifically, as synthesized by Rubin et al. (2015) across 11 electricity-supply technologies, these cost declines are typically 10-30% per doubling of installed capacity.[38]

DOE's Liftoff analysis suggests that moving from FOAK to subsequent builds of the same design could reduce overnight capital costs (the up-front cost of construction excluding financing) by roughly 30-40%.[39,40] This is contingent on delivering FOAK projects competently and then achieving approximately 10-20 serial deployments with reactor-to-reactor learning rates of roughly 12-15%. (Note that this is a modeled potential, as opposed to being evidence from standardized SMR fleets to date.)[41] Likewise, major assessments note that SMR economic competitiveness is still unproven at a commercial scale. By contrast, where standardization and serial build have been achieved in large reactor programs, costs and schedules have improved, offering a directional precedent but not a guarantee for SMRs.[42] In this context, structuring hyperscaler AMCs as both multi-unit and multi-year offtakes is best viewed as an enabling mechanism to increase the odds that cumulative deployment doublings occur. For example, pairing AMCs with policy support, finance backstops, liability frameworks and workforce pipelines further improves the chances that cost trajectories bend downward rather than stall.

There is also a clear policy analogue: the United Kingdom's Contracts-for-Difference (CfD), a public AMC, has delivered multi-GW auction rounds at declining strike prices over time, demonstrating how bankable revenue certainty mobilizes private capital at scale. The latest round awarded 9.6 GW across 128 projects, illustrating how demand-certainty mechanisms can push technologies down their cost curves.[43] For AI buyers, the lesson is strategic: catalyzing spillovers today helps secure scarce, reliable clean power near key load clusters and drive long-run costs down as technologies learn. Many hyperscalers even have explicit carbon-free energy (CFE) commitments, such as Google which aims to run on 24/7 carbon-free energy by 2030.[44] Unlike most large loads, hyperscalers are already helping co-design utility structures, such as Nevada's Clean Transition Tariff that enables the Google and Fervo geothermal scale-up (as discussed above).[45,46]

C. Mechanism 2: The Flexible Data Center: Transforming Power Grids

i. From rigid consumer to grid stabilizer

The next wave of owner-operated hyperscale clouds already modulates load in practice. Google's carbon-aware computing shifts movable batch work across data centers based on hourly carbon-free energy forecasts, and its "elastic training" and checkpointing tools show how some training jobs can pause and resume without losing progress.[47] Facilities can also flex on the non-IT side. For example, Microsoft's Dublin campus is using data center uninterruptible power supply (UPS) batteries for grid frequency services, demonstrating a pathway to make sites grid-responsive.[48] By contrast, most co-location operators—companies that lease space, power, and cooling to multiple tenants—face a split-incentive. The operators run the facility, but tenants control their own servers and workloads, which limits operator-led orchestration beyond facility-side measures.[49]

Duke University's Nicholas Institute estimates that US balancing authorities could integrate nearly 100 GW of new flexible large loads using existing capacity if those loads curtailed briefly during the most stressed hours. Specifically, these loads would have to curtail an average of 0.5% of annual uptime and roughly 2 hours on average during events.[50] This finding speaks to available headroom in today's system if new loads are flexible. Rather than claiming that more renewables must be built to serve these large flexible loads, the study shows modest flexibility can reduce the need for near-term capacity additions. Further, the same flexibility also improves renewable utilization when wind and solar are added by cutting curtailment when renewable resources are highly available and still trimming load during tight hours.[50]

Beyond Google and Microsoft, operators such as Digital Realty and Equinix in Dublin, Aeven in Denmark, and Basefarm in Norway have already deployed grid-interactive UPS or demand-response systems. These systems provide frequency regulation, which means quickly adjusting power to keep the grid stable, while maintaining quality of service (QoS) guarantees. [51-54]

On the modeling side, the National Renewable Energy Laboratory's (NREL's) Electrification Futures work finds demand-side flexibility cuts operating costs and eases the need for variable renewable energy (VRE) integration in highly electrified systems. However, its flexible-load assumptions are drawn mainly from electric vehicle (EV) charging and flexible industrial end uses.[55] This insight is still relevant to data centers because they can also match this profile when owners time-shift batch training or non-urgent inference and when sites use UPS batteries for grid services.[56] Of note, a portion of modern data centers now reach the 100+ MW class, comparable to classic single-site giants such as aluminum smelters (hundreds of MW), electric arc furnaces (EAF) and electrified steam crackers (hundreds of MW). However, data centers are increasingly software-defined, fast-ramping, and already equipped with battery and UPS assets that can provide services.[57-59] These features make owner-operated clouds uniquely positioned to catalyze clean energy and flexibility at scale, even though multi-tenant co-location sites will remain less controllable at the workload level.

ii. Innovations in action

First movers are already showing the playbook. For example, Google's carbon-aware computing shifts non-urgent workloads across hours and, where possible, across locations based on hourly carbon-free energy forecasts.[60] However, only a subset of organizations can do this in practice: owner-operated hyperscale clouds with global fleets and batchable jobs can orchestrate load, whereas most multi-tenant co-location sites cannot directly control tenant workloads.[61] In parallel, new platforms are emerging to translate that flexibility into operational grid support. For example, in May 2025, Emerald AI[a] (backed by NVIDIA's venture arm) conducted a field experiment with Salt River Project (SRP) in Phoenix and sustained a 25% power reduction for three hours on a live AI graphics processing unit (GPU) cluster (see Figure 4-4). In the Phoenix demonstration, AI jobs were tagged by how much performance they can temporarily give up during grid stress: Flex 1 (up to 10% performance reduction), Flex 2 (up to 25%) and Flex 3 (up to 50%). Given coincident-peak (CP) and capacity charges hinge on only a few stressed hours, a 25% curtailment capability typically yields about $40,000 to $70,000 per MW per year in ERCOT's 4 Coincident Peak (4CP), and about $17,000 to $160,000 per MW per year in PJM's 5 Coincident Peak (5CP). Therefore, a 100 MW site that can flex 25 MW often sees low seven figure annual bill reductions.[62-65]

This three-hour test translates to 24/7/365 operations because the peaks are brief and forecastable, so software can pre-stage checkpoints and apply power caps only during those intervals while running normally the rest of the year.[66] Emerald AI's Phoenix demonstration serves as a significant proof point for grid-resilience and cost reduction by shaving critical peaks without violating compute performance targets.[67] Together, these approaches map to what IEA expects the parts of the sector that are capable of shifting load to deliver this decade: flexible demand that helps integrate more renewables while easing immediate supply-side pressures in hot-spot regions.[6]

a Co-author Varun Sivaram is the founder and CEO of Emerald AI. Co-authors Ayse Coskun and Swasti Jain are affiliated with Emerald AI, as well.

Figure 4-4. *Emerald AI's Phoenix demonstration shows the power reduction curve and applications with different flexible service-level agreements (SLAs) meeting their performance requirements. (Source: ArXiv)[68]*

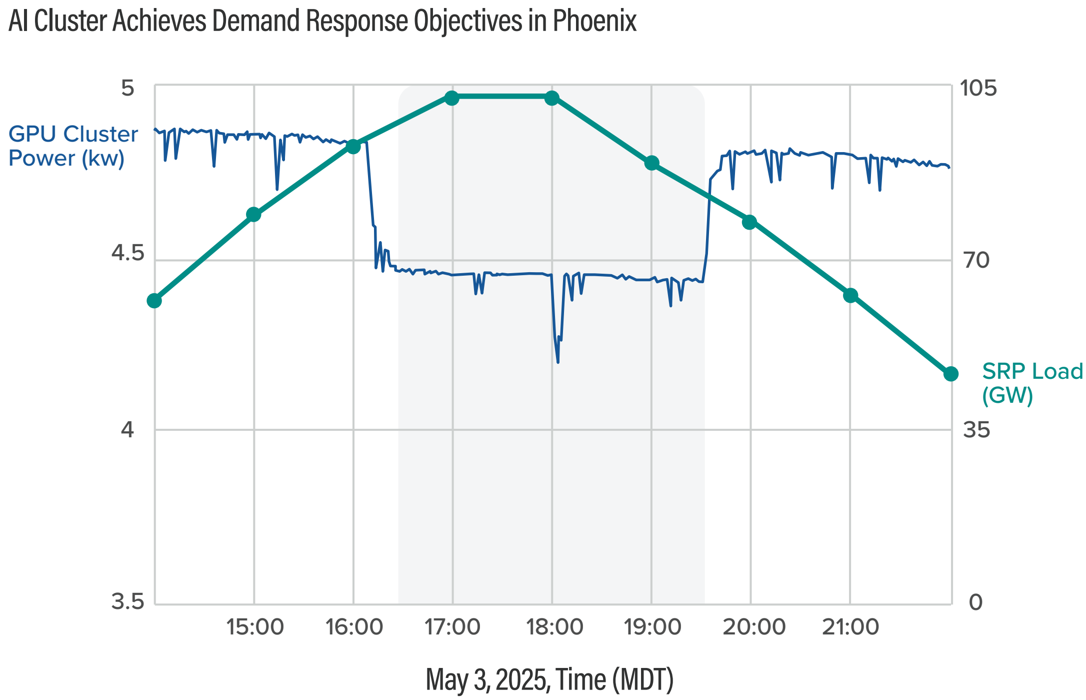

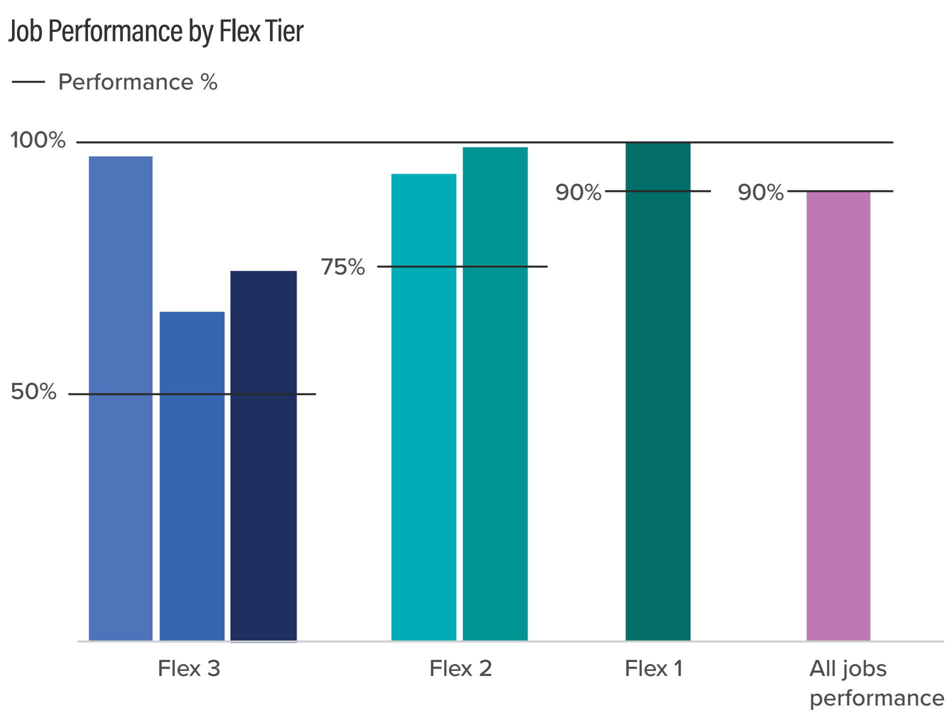

iii. A paradigm shift for grid management

As shown in Figure 4-5 below, the IEA expects flexible demand to absorb essentially all (~100%) projected data center additions through 2035 if operators are flexible for just ~0.1-1% of hours, complementing conventional flexibility from natural gas where it remains system-critical.[69] System context matters: in Japan, for example, tight reserve margins (a small amount of extra generation capacity), the division of the electric grid into two regions with different frequencies (50 Hz in the east and 60 Hz in the west) and policy design mean large, routine industrial curtailment is more constrained. In this case, Japanese market operators emphasize securing decarbonized supply, while demand response is being used primarily as a resource during tight hours in capacity markets.[70] In Japan, demand response participation is concentrated among large commercial and industrial (C&I) customers, such as high-voltage factories, campuses and big buildings, typically enrolled through aggregators certified by Japan's Energy Ministry. These aggregators also bundle behind-the-meter (BTM) batteries, combined heat and power, EVs and building controls.[71] This approach succeeds because there are clear, paying routes to market in the capacity market and the supply-demand balancing market. In 2022, for example, about 2.3 GW of demand response won roughly 60% of the "Power Source I" auction.[72]

Figure 4-5. *Data center capacity additions through 2035 and feasible integration into the current electricity system under different flexibility cases. (Source: IEA)*[69]

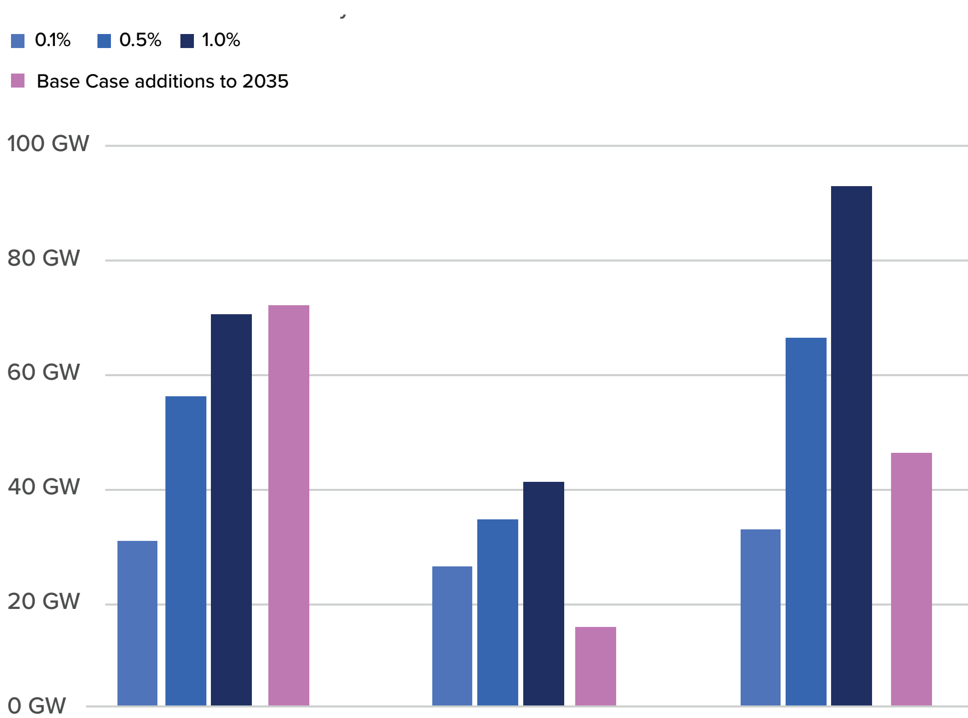

Current electricity systems can already integrate all data centre additions to 2035 if a mix of back up activation and workload management reduces grid demand 1% of the time.

iv. Enhanced renewable integration

Flexible demand enhances renewable integration by absorbing surplus solar and wind when it is plentiful and standing down during constrained hours. In doing so, flexible demand reduces the need for curtailment, enables deeper renewable penetration, and avoids overbuilding supply for rare peaks.[73] In California ISO's (Independent System Operator) grid, renewable curtailment reached the Terawatt-hour (TWh) scale in 2024, concentrated in high solar shoulder months: energy that shiftable loads (including data centers) could take up, especially when paired with short-duration storage and carbon-aware scheduling.[74] The IEA quantified the system impact with the US example, showing only ~0.3% of annual grid electricity needing active management during short 3-5 hour events.[69]

Recent research from the Association for Computing Machinery (ACM) Journal and ETH Zurich (Federal Institute of Technology Zurich) shows that QoS-aware scheduling and power-capping can deliver demand response and even regulation services without breaching service levels and that carbon-aware orchestration across sites is practical at scale.[75-77] Field evidence now complements this literature: Emerald AI's May 2025 demonstration on a 256-GPU cluster in a hyperscale cloud reduced power by 25% for three hours using a software-only approach while maintaining QoS.[68] On the value side, IEA modeling indicates that modest, time-bound flexibility reduces the need for curtailment and defers some peak-driven supply build, consistent with system-level findings.[69] The forward question is therefore not whether natural gas can provide flexibility, since it does, but what blend of flexible demand, storage and dispatchable supply delivers the best reliability and cost profile in each region.[78]

v. Grid resilience and cost reduction

When dispatched at the right times and places, flexible load can downsize networks and peaking investments, improving resilience and affordability. As shown in Figure 4-6 below, DOE defines virtual power plants (VPPs) as aggregations of distributed energy resources (DERs), such as rooftop solar with customer-sited batteries, EVs, and smart buildings that can balance electricity supply and demand and deliver utility-grade services like a power plant.[79] This chapter treats data center flexibility as virtual capacity that is compatible with virtual power plants, so DOE's VPP analysis is a useful proxy for system value. DOE estimates that tripling VPP capacity to approximately 80-160 GW by 2030 would reduce overall grid costs by roughly $10 billion per year by avoiding use of peakers (power-generating facilities that run during highest-demand periods), deferring grid investments and lowering operating costs.[79] In the United Kingdom, the National Energy System Operator's (NESO's) Demand Flexibility Service delivered 3.3 GWh of verified peak-time reductions in winter 2022 and 3917.7 MWh across 44 events in winter 2024 under its merit-based design.[80,81] As a result, NESO demonstrates a practical route to reliability at lower balancing cost.[82] At the distribution edge, National Grid—the investor-owned utility that owns the high-voltage transmission network in England and Wales—reports £80 million of investment deferral via 17 GWh of procured flexibility and greater than 70,000 registered flexible assets.[83] This is a different operational model than data centers. Aggregating thousands of devices is complex, while one or two large data centers can deliver the same order of magnitude of response with fewer parties and direct telemetry. This serves as evidence that market-based flexibility can substitute for some wires-and-steel upgrades. Although results will vary by system design and siting, the direction is consistent: well-designed flexibility reduces the need for both curtailment and high-cost build-outs.

Figure 4-6. *Virtual power plants (VPPs) aggregate distributed energy resources, such as rooftop solar with customer-sited batteries, electric vehicles (EVs) and chargers, smart buildings and equipment and their controls, and flexible commercial and industrial (C&I) loads. (Source: Department of Energy Loan Programs Office (DOE LPO))*[79]

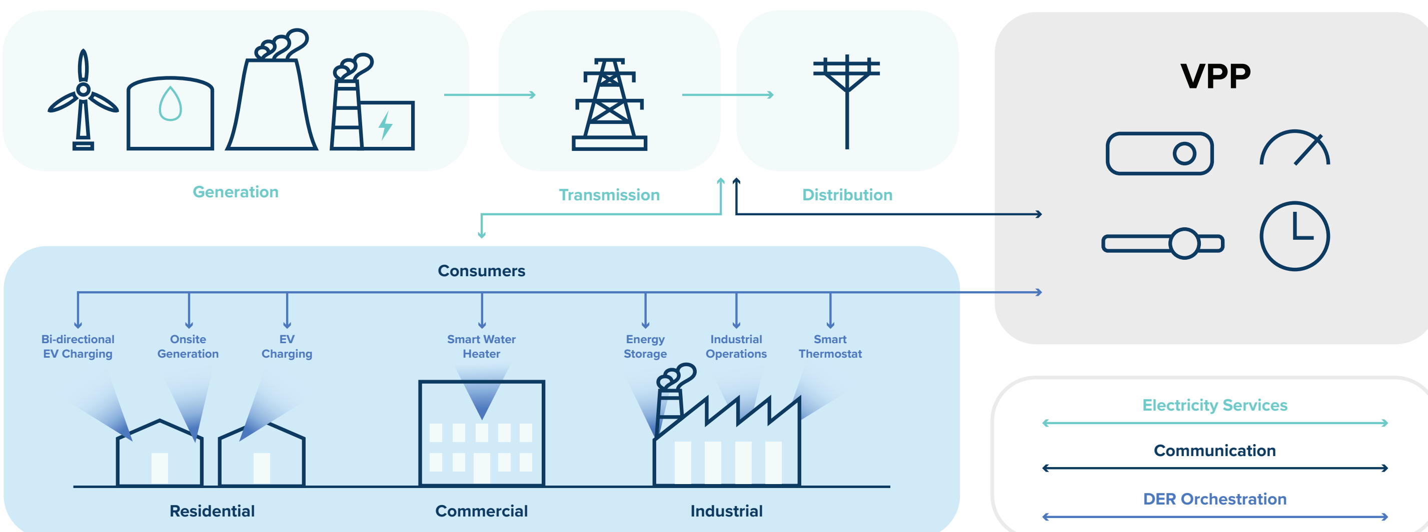

vi. Accelerated AI innovation (why flexibility can be a win-win)

From the grid's perspective, flexible data centers shave critical peaks, reduce the need for curtailment and defer costly reinforcements—improving reliability and affordability. Likewise, from the operator's perspective, flexibility can be a win-win where accelerated interconnection is a binding constraint. In other words, transmission operators may advance interconnection dates or offer ramped and conditional interconnections when projects commit to flexible operating modes. Not every market faces the same pressure, but the challenge is particularly acute in the United States, where the median time from interconnection request to operation reached about five years for projects built in 2023.[10] NESO has reached a point where it can bring forward 20 GW of connections in approximately 4 years via queue reforms and flexible arrangements.[84] In the United States, Federal Energy Regulatory Commission (FERC) Order 2023 is streamlining how new projects connect to the grid by studying many requests together in “clusters” and imposing withdrawal penalties to shorten queues. In parallel, some regions are trying “flexible interconnection,” through which very large new loads can connect sooner if they agree to operate within limits and share real-time data. For example, in Texas, ERCOT’s large-load process requires telemetry and allows voluntary curtailment during grid stress.[63] On the customer side, making this work usually takes three simple building blocks:

1. Power caps: Software limits that quickly lower server or site power for short periods.

2. Thermal buffers: Modest on-site chilled water or ice storage lets facilities store cooling when supply is plentiful and draw it down when the grid is constrained.[85]

3. Real-time telemetry tied to Service-Level Agreements (SLAs): Continuous, machine-readable telemetry linked to clear SLAs verifies the response for the grid while ensuring applications still meet their performance targets.

D. Mechanism 3: Siting for Sustainability

i. Siting for clean, reliable compute

Training and many batch inference jobs are tolerant to network delay, so operators can place new capacity in regions with strong renewable resources and then schedule work to follow local peaks. In the United States, the Great Plains offers exceptional wind potential and the desert in the Southwest offers high solar irradiance, which makes these regions natural candidates for “compute-next-to-resource” strategies.[86] Geothermal is also expanding beyond traditional fields as enhanced geothermal systems mature, widening the map for firm clean supply.[87] Remote siting often runs into limited grid headroom, which is why some projects consider BTM generation at the campus. This option remains the exception because most hyperscalers prefer utility or third-party power under bankable tariffs or PPAs rather than owning plants at scale. Recent experience in Pennsylvania shows how co-located or BTM nuclear supply raised regulatory and cost-allocation questions and ultimately shifted toward a standard grid-connected structure.[88,89] This preference aligns with the broader corporate market’s use of utility green tariffs and long-term contracts to add new clean capacity.[89]

When the goal is to accelerate new clean energy, siting choices work best when paired with enabling mechanisms that bring additional projects onto the grid. One pathway is utility tariffs designed for hourly matched portfolios, such as Nevada's approved structure that enables Google to source new geothermal and storage through NV Energy.[87,90] A second pathway is risk-reduction finance for firm and long-duration storage so remote renewable hubs can deliver around the clock. Federal loan guarantees have begun backing multi-hour storage that firms renewables and lowers system costs. Finally, grid modernization investments, such as advanced conductors, dynamic line ratings and new high-capacity corridors, expand the ability of resource-rich zones to host large flexible loads.[91-93]

ii. Catalyzing new clean energy development

Large AI data center projects, if not planned carefully, can drive up costs for other ratepayers and strain grid stability. However, relocatable data center loads offer a unique opportunity to accelerate low-carbon power deployment when aligned with renewable energy growth. By being flexible on location, these energy-hungry data centers can be placed in resource-rich regions, such as deserts or windy plains, where new solar and wind farms can be built to supply them.[86] This strategic siting turns data centers into demand catalysts for additional clean generation capacity. In practice, a relocatable AI campus can serve as an anchor offtaker that makes a new renewable project financially viable, bringing extra wind or solar online faster than traditional grid planning might.[14] One concrete example is Microsoft's New Zealand cloud region. Microsoft signed a 10-year agreement tied to Contact Energy's new 51.4 MW Te Huka III geothermal plant, and Contact explicitly says the contract supported its investment decision to build the project.[104] The plant began delivering to the grid in October 2024, illustrating how a relocatable cloud load can underwrite firm clean capacity that serves both the data center and the broader system. In short, when AI data centers are free to move to the power (instead of always bringing power to the load), they can help drive an accelerated build-out of wind, solar and other low-carbon resources.

Critically, realizing these benefits depends on project design and grid integration. Reaping community-wide gains requires that the clean electricity procured for a data center also strengthen the public grid. For example, using utility green tariffs or CFDs ensures that new renewable capacity is not just behind-the-fence—on the customer side of the meter, serving only the facility—but also contributes energy to everyone over time.[94,95] As Figure 4-7 shows, relocatable data centers should pair any new renewables with energy storage and embrace flexible operations (e.g., drawing less power at grid peak times or using on-site batteries), so their presence reduces the need for curtailment and eases grid congestion instead of exacerbating it.[14] This strategy defers the need for certain network upgrades and helps ensure reliability is maintained with renewables. Only when power procurement, interconnection terms and community agreements are structured for system-wide value will a relocatable load accelerate the clean energy transition.[96] In practice, this means the data center's clean power investments should lower overall system costs over time (by adding new capacity and reducing fuel use), and its operators should coordinate with utilities to avoid saddling other customers with infrastructure expenses. When done right, a flexible, renewables-driven data center can act as a springboard for more clean energy on the grid to deliver not just low-carbon electricity for its own operations, but cheaper and more reliable power for the surrounding region as well.[14]

Figure 4-7. *Solar and onshore wind (especially when paired with batteries) are the fastest, lowest-cost new capacity, while coal/CCUS (carbon capture, utilization and storage), hydroelectric, geothermal and offshore wind are slower and more expensive. (Source: Rocky Mountain Institute (RMI))*[14]

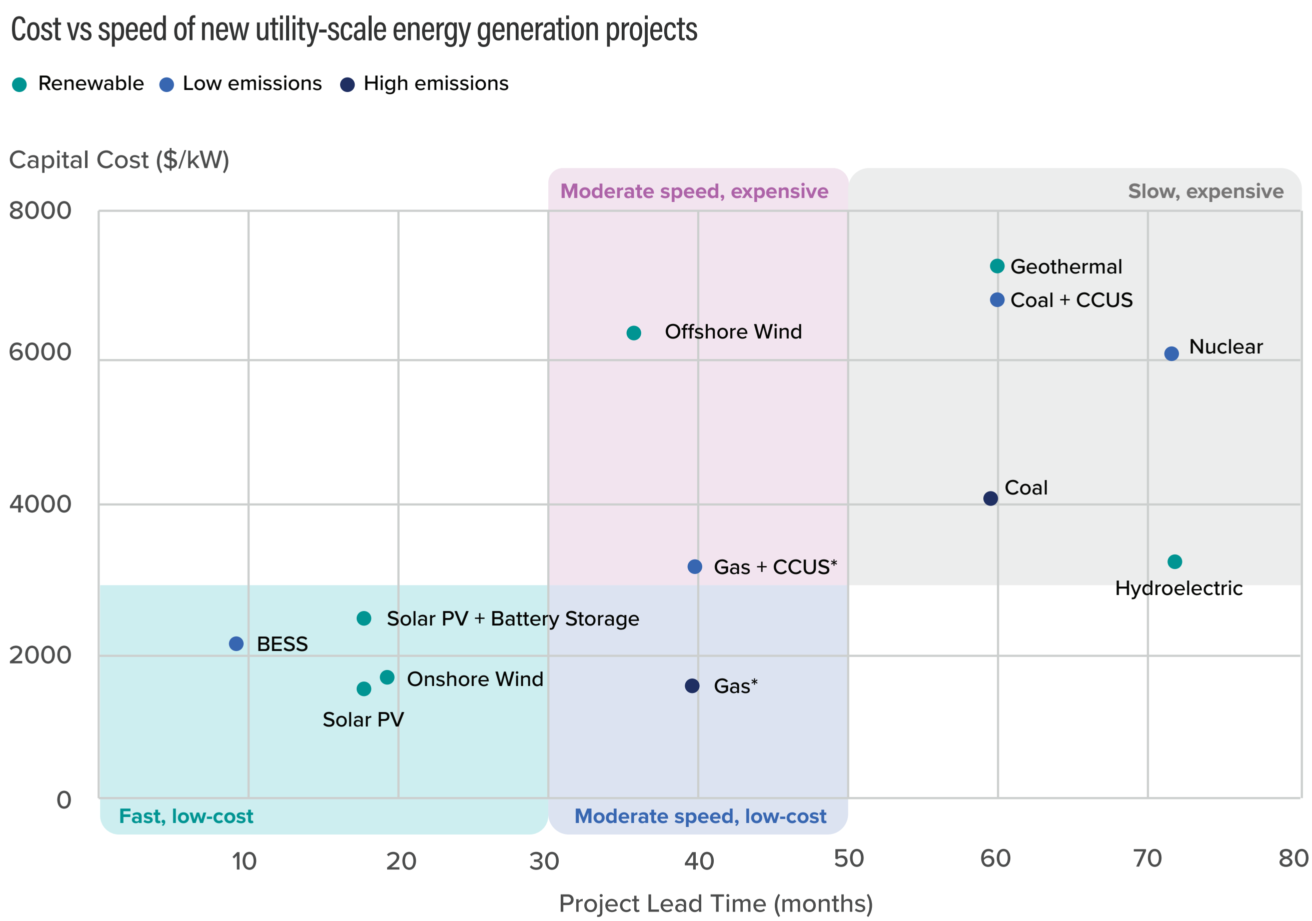

E. Mechanism 4: AI as the Architect of a Clean Energy Future

i. Grid optimization

One near-term, high-impact role for AI is speeding interconnection and transmission studies: an "offline" use that directly unlocks renewable projects. As Figure 4-8 shows, Southwest Power Pool's (SPP's) generator-interconnection (GI) queue is facing significant backlogs. Therefore, SPP is partnering with Hitachi Energy to apply AI to generator-interconnection queue triage and reliability assessments so planners can screen power-flow cases, prioritize upgrades and shorten study timelines.[97] Hitachi and SPP will use an integrated NVIDIA-based compute and AI platform to automate processes, add predictive analytics and pilot AI-augmented simulation, with systems acceleration and data-management optimization due by December 2025. These tools complement both FERC Order 2023 queue reforms and DOE's Grid Modernization Initiative by automating contingency ranking, hosting-capacity mapping, and topology and stability screening. All these partnerships help connect renewables and storage sooner, while reducing the need for curtailment and fuel burn.[97]

Figure 4-8. *Southwest Power Pool's (SPP's) generator-interconnection (GI) queue lists 659 projects totaling ~157.7 GW, with ~380 currently in study and clusters still open from 2018. Source: RTO Insider[97]*

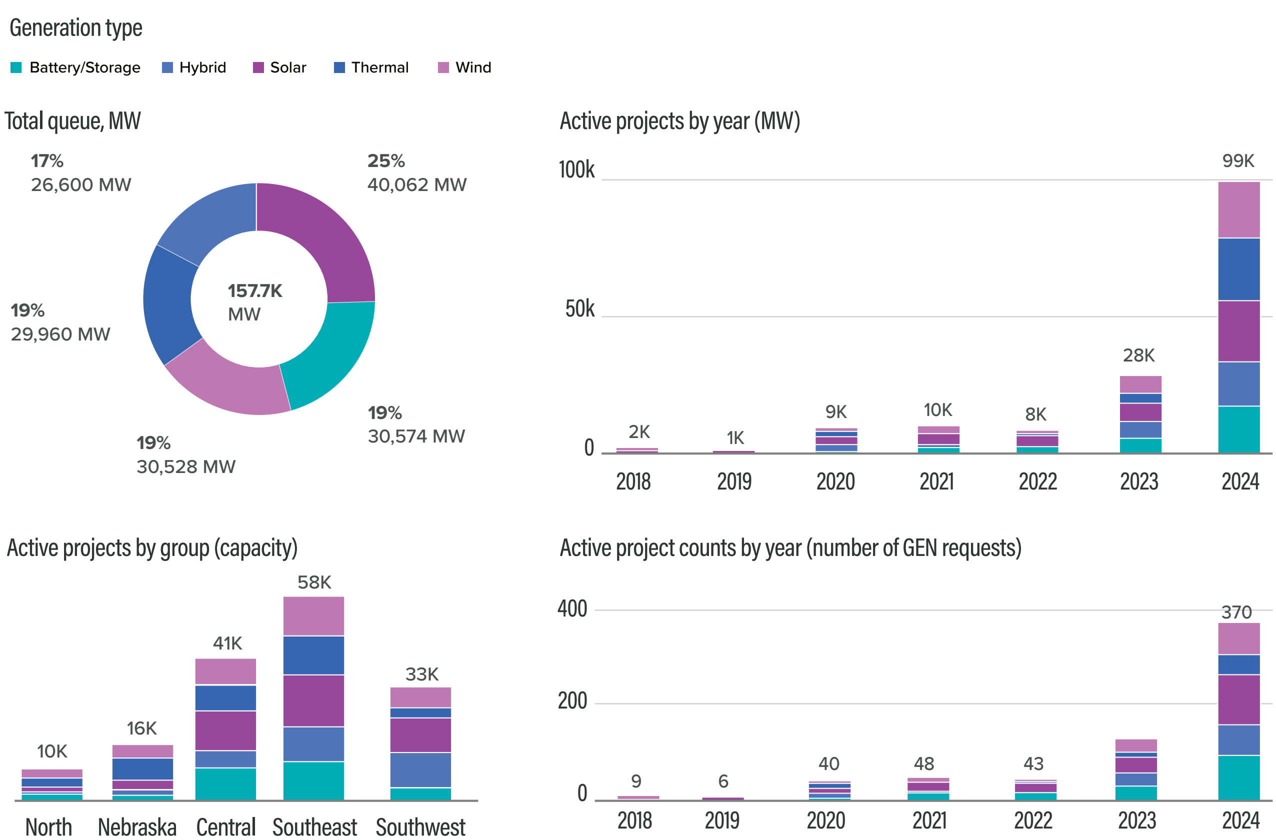

ii. Materials discovery and design

AI is compressing the discovery-to-device cycle for clean energy materials. For example, instead of testing materials one by one, AI learns from big materials databases and predicts which compositions and crystal structures are most promising. It can sift through tens of millions of options and surface a few thousand good candidates for battery electrodes, solid electrolytes, photovoltaic absorbers and catalysts, among other materials.[98] These models are coupled with high-throughput first-principles calculations and open data infrastructures, such as The Materials Project, to validate stability and properties and to hand off the best leads to synthesis.[99] On the lab side, AI tools mine the literature for synthesis recipes and processing conditions, feeding autonomous platforms that can make and test candidates rapidly: shortening iteration times for solar, battery and CO_2-conversion technologies while lowering cost-to-proof.[100]

iii. Improved forecasting

AI is also closing the gap between weather hazards and power-system impacts by improving both the speed and the skill of forecasts that feed unit commitment, dispatch, storage scheduling and outage preparedness. Thus, current versions of AI are already building on the "hazard-to-impact" framing highlighted in last year's ICEF Roadmap chapter on extreme-weather response, which emphasized more actionable signals for heat waves, wind ramps and atmospheric rivers.[101] Today's AI models deliver state-of-the-art medium-range skill at a fraction of the compute and energy cost of traditional physics-based forecasts on supercomputers. The European Centre for Medium-Range Weather Forecast's (ECMRWF) now operates an Artificial Intelligence Forecasting System (AIFS) in parallel with its traditional physics-based Integrated Forecasting System (IFS). Peer-reviewed AI models, such as GraphCast demonstrate strong 10-day global forecast performance, giving grid operators more frequent, higher-skill inputs for renewable scheduling and contingency analysis.[102] The AI outputs are trained and calibrated against authoritative observational datasets from the National Oceanic and Atmospheric Administration (NOAA) and other government weather services, which remain the backbone for assimilation and verification.[103]

Recent American Meteorological Society (AMS) work shows that machine-learning corrections to models can dramatically reduce wind-forecast error, which directly lowers reserve needs, and the need for curtailment and balancing costs for renewable-heavy grids. Taken together, these three AI applications are beginning to reinforce one another. AI for interconnection and other offline studies helps planners clear queues and move clean projects to "shovel-ready" status faster. AI-enhanced operational forecasts improve day-ahead and intraday scheduling of wind, solar and storage. Lastly, major weather centers are now deploying probabilistic AI ensembles. As a result, this combination creates a virtuous cycle that speeds grid build-out, strengthens extreme event readiness, and enables deeper renewable penetration and more reliable service.[103]

F. Recommendations

1. *Utilities and independent power producers (IPP) should:*

    - ***Deploy advanced control tools*** *to accelerate interconnection and grid studies and to operate flexible portfolios. These tools include model-predictive control, enhanced forecasting and, where appropriate, AI.*

    - ***Adopt staged or ramped interconnections for large loads*** *(within standard planning cycles) and require telemetry and fast power-capping from data centers and consider on-site storage to provide demand response and regulation while maintaining service-level objectives.*

    - *Use these tools to* ***prioritize non-wires alternatives*** *and to reduce curtailment in renewable-rich zones.*

2. *Electricity regulators should:*

    - ***Establish clear 24/7 carbon-free energy procurement pathways*** *that treat storage and clean-firm resources as first-class options alongside renewables.*

    - *Enable* ***advanced market commitments (AMCs)*** *that allow multi-buyer participation, recognize hourly matching, and credit verifiable flexible-load performance.*

3. *National governments, regulators, and utilities should:*

    - ***Expand targeted public-private risk-sharing*** *to lower the cost of* *firm, low-carbon supply while keeping rates affordable amid rising public concern about electricity bill impacts from data center–driven capacity additions.*

- ***Pair corporate offtake with loan guarantees**, liability and fuel frameworks (where relevant), and long-duration storage demonstrations.*
- ***Adapt CfD-style mechanisms to clean-firm resources and storage** so FOAK projects are followed by repeat builds of the same design.*

4. *Large data center operators with load flexibility, hyperscalers and procurement authorities should:*
    - ***Commit to portfolio-based, 24/7 carbon-free procurement** that include renewables, storage and clean-firm resources where available.*
    - ***Publish transparent hourly performance** and adopt grid-supportive operating modes, such as fast power caps and brief curtail on-signal, to **unlock faster interconnection and lower system costs**.*
    - ***Prioritize deliverable power** to the public grid (when siting in resource-rich regions), rather than exclusively behind-the-fence supply.*
5. *National and local governments should:*
    - ***Link siting incentives for new AI campuses to system and community value.***
    - *Require additional deliverable **clean capacity, storage co-procurement and community benefit plans** that include workforce pipelines, water stewardship and shared transmission upgrades.*
6. *National governments and utilities, including public power and transmission owners, should **invest in grid modernization**, including advanced transmission, system visibility and congestion management, so resource-rich zones can host large and flexible loads without unnecessary overbuild. Regulators should authorize these investments, set incentives and ensure timely cost recovery.*
7. *Academic experts and system operators should **advance operations-ready forecasting for clean grids**. Priorities include post-processing and downscaling of weather models (including via AI) for wind and solar, probabilistic products that feed unit commitment and storage scheduling, and open benchmarks that **connect forecast improvements to avoided reserves, reduced curtailment, and emissions reductions**.*

G. References

1. International Energy Agency (IEA). Chapter 3: Energy Demand from AI in Energy and AI (IEA, Paris, France, 2025, https://www.iea.org/reports/energy-and-ai/energy-demand-from-ai).

2. International Energy Agency (IEA). Energy and AI Observatory; IEA, Paris, France, https://www.iea.org/data-and-statistics/data-tools/energy-and-ai-observatory?tab=Energy+for+AI (Accessed August 2025 (Last updated June 2025).

3. Cheryl Schweizer. Grant PUD places limits on electrical demand from data centers; Columbia Basin Herald, Moses Lake, Washington, https://columbiabasinherald.com/news/2025/mar/31/grant-pud-places-limits-on-electrical-demand-from-data-centers/ (2025).

4. Office of Energy Statistics Staff. Solar, battery storage to lead new U.S. generating capacity additions in 2025; US Energy Information Administration (EIA), Washington, DC, https://www.eia.gov/todayinenergy/detail.php?id=64586 (2025).

5. Garrett Golding. Solar, battery capacity saved the Texas grid last summer; an uncertain future awaits; Federal Reserve Bank of Dallas, Dallas, Texas, https://www.dallasfed.org/research/economics/2025/0114 (2025).

6. International Energy Agency (IEA). Chapter 4: Energy Supply for AI in Energy and AI (IEA, Paris, France, 2025, https://www.iea.org/reports/energy-and-ai/energy-supply-for-ai).

7. Kristin Toussaint. Meta is building a new data center in Louisiana—and this Senate committee wants to know why it's being powered by gas (exclusive); Fast Company, New York, New York, https://www.fastcompany.com/91334846/meta-is-building-a-new-data-center-in-louisiana-and-this-senate-committee-wants-to-know-why-its-being-powered-by-gas-exclusive (2025).

8. Julian Cruz. Brazil's Data Center Boom: A Renewable-Fueled Infrastructure Play in a Strategic Market; AInvest, New York, New York, https://www.ainvest.com/news/brazil-data-center-boom-renewable-fueled-infrastructure-play-strategic-market-2505/ (2025).

9. Arman Shehabi, Alex Newkirk, Sarah J Smith, Alex Hubbard, Nuoa Lei, Md Abu Bakar Siddik, Billie Holecek, Jonathan Koomey, Eric Masanet & Dale Sartor. 2024 United States Data Center Energy Usage Report (Report #: LBNL-2001637); Lawrence Berkeley National Laboratory, https://doi.org/10.71468/P1WC7Q; Retrieved from https://escholarship.org/uc/item/32d6m0d1 (2024).

10. Joseph Rand, Nick Manderlink, Will Gorman, Ryan H Wiser, Joachim Seel, Julie Mulvaney Kemp, Seongeun Jeong & Fritz Kahrl. Queued Up: 2024 Edition, Characteristics of Power Plants Seeking Transmission Interconnection As of the End of 2023; Lawrence Berkeley National Laboratory Energy Analysis and Environmental Impacts Division, Energy Markets and Policy Department Berkeley, California, https://emp.lbl.gov/sites/default/files/2024-04/Queued%20Up%202024%20Edition_1.pdf (2024).

11. Tony Lenoir. Corporate PPA leaderboard – Microsoft leap cuts into Amazon lead; S&P Global, New York, New York, https://www.spglobal.com/market-intelligence/en/news-insights/research/corporate-ppa-leaderboard-microsoft-leap-cuts-into-amazon-lead (2025).

12. Mead Gruver & Matt O'Brien. Cheyenne to host massive AI data center using more electricity than all Wyoming homes combined; The Associated Press (AP), New York, New York, https://apnews.com/article/ai-artificial-intelligence-data-center-electricity-wyoming-cheyenne-44da7974e2d942acd8bf003ebe2e855a (2025).

13. Bilge Acun, Benjamin Lee, Fiodar Kazhamiaka, Kiwan Maeng, Manoj Chakkaravarthy, Udit Gupta, David Brooks & Carole-Jean Wu. Carbon Explorer: A Holistic Approach for Designing Carbon Aware Datacenters. arXiv:2201.10036 (2022). https://doi.org/10.48550/arXiv.2201.10036.

14. Yuki Numata, Alexandra Gorin, Laurens Speelman, Lauren Shwisberg & Chiara Gulli. Fast, Flexible Solutions for Data Centers; Rocky Mountain Institute (RMI), Basalt, Colorado, https://rmi.org/fast-flexible-solutions-for-data-centers/ (2025).

15. Emma Penrod. Google taps 'carbon-intelligent' computing platform to help maintain grid reliability in power crises; Utility Dive (Informa TechTarget), Newton, Massachusetts, https://www.utilitydive.com/news/google-carbon-intelligent-computing-platform-system-reliability-demand-response-grid-emergency/698958/ (2023).

16. David Sandalow, Colin McCormick, Alp Kucukelbir, Julio Friedmann, Michal Nachmany, Hoesung Lee, Alice Hill, Daniel Loehr, Matthew Wald, Antoine Halff, Ruben Glatt, Philippe Benoit, Kevin Karl et al. Artificial Intelligence for Climate Change Mitigation Roadmap (Second Edition) (ICEF Innovation Roadmap Project, November 2024); https://doi.org/10.7916/2j4p-nw61 (2024).

17. Synergy Research Group. Hyperscale Data Center Count Hits 1,136; Average Size Increases; US Accounts for 54% of Total Capacity; Reno, Nevada, https://www.srgresearch.com/articles/hyperscale-data-center-count-hits-1136-average-size-increases-us-accounts-for-54-of-total-capacity (2025).

18. Matt Busse. Data centers driving 'immense increase' in Virginia's energy demand, report says; Cardinal News, Roanoke, Virginia, https://cardinalnews.org/2024/12/10/data-centers-driving-immense-increase-in-virginia-energy-demand-report-says/ (2024).

19. Lancium. Lancium and the Stargate Project in Abilene, TX: Bringing Hyperscale Campuses to Texas; Energy Systems Integration Group (ESIG), Reston, Virginia, https://www.esig.energy/wp-content/uploads/2025/05/ESIG_LLTF_PresentationLancium.pdf (2025).

20. Trade and Industry (METI) and the Ministry of Internal Affairs and Communications (MIC) Ministry of Economy. Public-Private Advisory Council on Watt-Bit Collaboration to Launch; METI, Tokyo, Japan, https://www.meti.go.jp/english/press/2025/0318_001.html (2025).

21. Scala Data Centers. Scala Data Centers and Serena Announce the Largest Renewable Energy Partnership in Latin America's Data Center Industry; São Paulo, Brazil, https://scaladatacenters.com/en/scala-data-centers-and-serena-announce-the-largest-renewable-energy-partnership-in-latin-americas-data-center-industry/

22. Azeem Azhar. Big Tech Will Scour the Globe in Its Search for Cheap Energy; WIRED Magazine, San Francisco, California, https://www.wired.com/story/big-tech-data-centers-cheap-energy/ (2024).

23. Hydro-Québec Stratégies et Finances. Action Plan 2035: Towards a Decarbonized and Prosperous Québec (English Version); Hydro-Québec, Montreal, Québec, Canada, https://www.hydroquebec.com/data/a-propos/pdf/action-plan-2035.pdf (2023).

24. Dominique Rolland, Patricia Larouche, Shannon Corbeil Bathurst & Mariane Gagné. Québec Government Adopts Bill 69 to Modernize Energy Legislation; Stikeman Elliott LLP Toronto, Ontario, Canada, https://stikeman.com/en-ca/kh/canadian-energy-law/quebec-government-adopts-bill-69-to-modernize-energy-legislation (2025).

25. Darryl Willis. Microsoft partners celebrate AI innovation; Microsoft, Redmond, Washington, https://www.microsoft.com/en-us/industry/blog/energy-and-resources/2024/06/27/microsoft-partners-celebrate-ai-innovation (2024).

26. Directorate-General for International Partnerships (DG-INTPA). Global Gateway: Namibia becomes a pioneer for Africa's green transition; European Commission DG-INTPA, Brussels, Belgium, https://international-partnerships.ec.europa.eu/news-and-events/news/global-gateway-namibia-becomes-pioneer-africas-green-transition-2025-04-11_en (2025).

27. Lorenzo Moavero Milanesi, Tjark Freundt, Yuito Yamada, Fridolin Pflugmann & Marc Ludwig. How hyperscalers are fueling the race for 24/7 clean power; McKinsey & Company, New York, New York, https://www.mckinsey.com/industries/electric-power-and-natural-gas/our-insights/how-hyperscalers-are-fueling-the-race-for-24-7-clean-power (2024).

28. Michael Terrell. With new geothermal project, it's full steam ahead for 24/7 carbon-free energy; Google Cloud, Mountain View, California, https://cloud.google.com/blog/products/infrastructure/google-fervo-geothermal-project-creates-carbon-free-energy (2021).

29. Emma Penrod. NV Energy seeks new tariff to supply Google with 24/7 power from Fervo geothermal plant; Utility Dive (Informa TechTarget), Newton, Massachusetts, https://www.utilitydive.com/news/google-fervo-nv-energy-nevada-puc-clean-energy-tariff/719472/ (2024).

30. Erick R. Burns, Colin Williams & Jacob DeAngelo. Enhanced geothermal systems electric-resource assessment for the Great Basin, southwestern United States. Report No. 2025-3027, 4 (US Geological Survey (USGS); Reston, VA, 2025), https://doi.org/10.3133/fs20253027.

31. Chen Chen, Daniel Merino-Garcia, Timothy D. G. H. Lines & Daniel S. Cohan. Geothermal power generation potential in the United States by 2050. Environmental Research: Energy 1, 025003 (2024). https://dx.doi.org/10.1088/2753-3751/ad3fbb.

32. US Nuclear Regulatory Commission (NRC). Backgrounder on Nuclear Insurance and Disaster Relief; US NRC, Washington, DC, https://www.nrc.gov/reading-rm/doc-collections/fact-sheets/nuclear-insurance.html (Accessed August 2025 (Last updated August 2025).

33. World Nuclear News (WNN). Constellation to restart Three Mile Island unit, powering Microsoft; WNN, London, United Kingdom, https://www.world-nuclear-news.org/articles/constellation-to-restart-three-mile-island-unit-powering-microsoft (2024).

34. Brian Martucci. Constellation plans 2028 restart of Three Mile Island unit 1, spurred by Microsoft PPA; Utility Dive (Informa TechTarget), Newton, Massachusetts, https://www.utilitydive.com/news/constellation-three-mile-island-nuclear-power-plant-microsoft-data-center-ppa/727652/ (2024).

35. Matt Wald. Resurrecting Three Mile Island; Nuclear Newswire, American Nuclear Society (ANS) Westmont, Illinois, https://www.ans.org/news/2025-01-24/article-6714/resurrecting-three-mile-island/ (2025).

36. Commonwealth Fusion Systems (CFS) and Google. Google and Commonwealth Fusion Systems Sign Strategic Partnership; CFS and Google, Devens, Massachusetts and Mountain View, California, https://cfs.energy/news-and-media/google-and-commonwealth-fusion-systems-sign-strategic-partnershipGoogle (2025).

37. David Kirtley. Announcing Helion's fusion power purchase agreement with Microsoft; Helion Energy, Everett, Washington, https://www.helionenergy.com/articles/announcing-helion-fusion-ppa-with-microsoft-constellation/ (Accessed August 2025).

38. Edward S. Rubin, Inês M. L. Azevedo, Paulina Jaramillo & Sonia Yeh. A review of learning rates for electricity supply technologies. Energy Policy 86, 198-218 (2015). https://doi.org/10.1016/j.enpol.2015.06.011.

39. US Department of Energy (DOE) Office of Nuclear Energy. Commercializing Advanced Nuclear Reactors Explained in Five Charts; Washington, DC, https://www.energy.gov/ne/articles/commercializing-advanced-nuclear-reactors-explained-five-charts (2023).

40. Julie Kozeracki, Chris Vlahoplus, Katheryn Scott, Melissa Bates, Billy Valderrama, Erica Bickford, Tim Stuhldreher, Andrew Foss & Tom Fanning. Pathways to Commercial Liftoff: Advanced Nuclear; US Department of Energy (DOE), Washington, DC, https://www.energy.gov/sites/default/files/2025-07/LIFTOFF_DOE_Advanced-Nuclear.pdf (2023).

41. Kelly Cummins, Melissa Klembara, Alexa Thompson, Olivia Corriere, Emily Goldfield & William Dean. Portfolio Insights: Learning from Case Studies: Financing and Development Approaches from Recent First-of-a-Kind Projects; Office of Clean Energy Demonstration (OCED), Washington, DC, https://www.energy.gov/sites/default/files/2024-11/FOAK%20Financing%20and%20Development%20Approaches_112024_vf.pdf (2024).

42. Ryan Spangler, Sunming Qin, Levi M Larsen, Chandu Bolisetti, Abdalla Abou-Jaoude, Mehdi Asgari, Koroush Shirvan, W. Robb Stewart, James Krellenstein & Garrett Wilkinson. Potential Cost Reduction in New

Nuclear Deployments Based on Recent AP1000 Experience: Systems Analysis & Integration Campaign; US Department of Energy (DOE) Systems Analysis & Integration Campaign, Washington, DC, https://inldigitallibrary.inl.gov/sites/STI/STI/Sort_173162.pdf (2025).

43. Department for Energy Security and Net Zero. Guidance: Contracts for Difference (CfD) Allocation Round 6: results GOV.UK, London, United Kingdom, https://www.gov.uk/government/publications/contracts-for-difference-cfd-allocation-round-6-results (2024).

44. Google Sustainability. 24/7 by 2030: Realizing a Carbon-free Future; Mountain View, California, https://sustainability.google/reports/247-carbon-free-energy/ (2020).

45. Lisa Martine Jenkins. The 'clean transition tariff' won approval in Nevada. What's next for Fervo?; Latitude Media, Somerville, Massachusetts, https://www.latitudemedia.com/news/the-clean-transition-tariff-won-approval-in-nevada-whats-next-for-fervo/ (2025).

46. Vicki M. Baldwin (Parsons Behle & Latimer). Public Utilities Commission of Nevada Electronic Filing RE: Docket No. 23-04016 - Investigation regarding Sierra Pacific Power Company d/b/a NV Energy's Progress of Developing a Clean Transition Tariff and its Appropriateness for Development as a New Tariff Offering Pursuant to the Modified Final Order Issued in. Docket Nos. 22-06014, 22-06015 and 22-06016; Public Utilities Commission of Nevada, Carson City, Nevada, https://pucweb1.state.nv.us/PDF/AxImages/DOCKETS_2020_THRU_PRESENT/2023-4/26664.pdf (2023).

47. Deepak Patil. Train AI for less: Improve ML Goodput with elastic training and optimized checkpointing; Google Cloud, Mountain View, California, https://cloud.google.com/blog/products/ai-machine-learning/elastic-training-and-optimized-checkpointing-improve-ml-goodput (2025).

48. John Roach. Microsoft datacenter batteries to support growth of renewables on the power grid; Microsoft, Redmond, Washington, https://news.microsoft.com/source/features/sustainability/ireland-wind-farm-datacenter-ups (2022).

49. Alex C. Newkirk, Nichole Hanus & Christopher T. Payne. Expert and operator perspectives on barriers to energy efficiency in data centers. Energy Efficiency 17, 63 (2024). https://doi.org/10.1007/s12053-024-10244-7.

50. Tyler H. Norris, Tim Profeta, Dalia Patino-Echeverri & Adam Cowie-Haskel. Three Key Takeaways: Rethinking Load Growth in U.S. Power Systems; Nicholas Institute for Energy, Environment & Sustainability at Duke University, Durham, North Carolina, https://nicholasinstitute.duke.edu/articles/three-key-takeaways-rethinking-load-growth-us-power-systems (2025).

51. Will Reynolds & Emma Long. Digital Realty and Enel X Partner to Support Ireland's Renewable Energy Transition; Digital Realty, Dublin, Ireland, https://www.digitalrealty.com/about/newsroom/press-releases/123259/digital-realty-and-enel-x-partner-to-support-irelands-renewable-energy-transition (2024).

52. Georgia Butler. Schneider Electric & Aeven to deliver excess power to Danish grid; Data Center Dynamics, London, United Kingdom, https://www.datacenterdynamics.

com/en/news/schneider-electric-aeven-to-deliver-excess-power-to-danish-grid/ (2023).

53. Vera Grishchenko. Eaton UPS-as-a-Reserve Proven in Data Center Pilot; Eaton, Espoo, Finland, https://www.eaton.com/us/en-us/company/news-insights/news-releases/2018/UPSasR-pilot-Norway.html (2018).

54. Equinix. Response by Equinix to Cru Consultation: Large Energy Users Connection policy; Redwood City, California, https://consult.cru.ie/en/system/files/materials/200/CRU202504ac%20-%20Equinix.pdf (2024).

55. National Renewable Energy Laboratory (NREL). Flexible Loads and Renewable Energy Work Together in a Highly Electrified Future; NREL, Golden, Colorado, https://www.nrel.gov/manufacturing/news/program/2021/flexible-loads-and-renewable-energy-work-together-in-a-highly-electrified-future (2021).

56. Ella Zhou & Trieu Mai. Electrification Futures Study: Operational Analysis of U.S. Power Systems with Increased Electrification and Demand-Side Flexibility (NREL/TP-6A20-79094); National Renewable Energy Laboratory (NREL), Golden, Colorado, https://docs.nrel.gov/docs/fy21osti/79094.pdf (2021).

57. Andy Home. US aluminium smelters vie with Big Tech for scarce power; Reuters, London, United Kingdom, https://www.reuters.com/markets/commodities/us-aluminium-smelters-vie-with-big-tech-scarce-power-andy-home-2025-05-22/ (2025).

58. Ltd. (SRDL) Sanrui Electric Furnace Co. How Much Power Does an Electric Arc Furnace Require? ; SRDL, Xi'an, Shaanxi Province, China, https://www.srfurnace.com/how-much-power-does-an-electric-arc-furnace-require.html (2025).

59. Jim Middleton. Decarbonisation of steam crackers; Decarbonisation Technology (Crambeth Allen Publishing Ltd), Croydon, United Kingdom, https://www.ten.com/sites/energies/files/2023-05/decarbonisation_of_steam_crackers_technipenergies.pdf (2021).

60. Ana Radovanovic. Our data centers now work harder when the sun shines and wind blows (Blog); Google, Mountain View, California, https://blog.google/inside-google/infrastructure/data-centers-work-harder-sun-shines-wind-blows/ (2020).

61. Heather Klemick, Elizabeth Kopits & Ann Wolverton. Data Center Energy Efficiency Investments: Qualitative Evidence from Focus Groups and Interviews; US Environmental Protection Agency (EPA), Washington, DC, https://www.epa.gov/environmental-economics/data-center-energy-efficiency-investments-qualitative-evidence-focus-groups (2017).

62. CPower. 4CP Management System: A Dynamic Solution to Managing 4CP Charges; Baltimore, Maryland, https://cpowerenergy.com/wp-content/uploads/2016/12/ERCOT_4CP_Web_Download.pdf Accessed August 2025).

63. Electric Reliability Council of Texas (ERCOT). ERCOT Four Coincident Peak Calculations; ERCOT, Taylor, Texas, https://www.ercot.com/mktinfo/data_agg/4cp (2024).

64. Rodan Energy Solutions. Everything You Need to Know About 5CP In PJM (Blog); Princeton, New Jersey, https://rodanenergy.com/everything-you-need-to-know-about-5cp-in-pjm (2024).

65. Resource Adequacy Planning. PJM Manual 19: Load Forecasting and Analysis (Revision: 35); PJM (Pennsylvania-New Jersey-Maryland) Interconnection, Valley Forge, Pennsylvania, https://www.pjm.com/-/media/DotCom/documents/manuals/archive/m19/m19v35-load-forecasting-and-analysis-12-31-2021.pdf (2021).

66. Grid Status Exports. ERCOT's 4CP Summer Demand Roller Coaster Takes Off as Storage Flips Outcomes (Blog); Chicago, Illinois, https://blog.gridstatus.io/ercot-4cp-2025-june/ (2025).

67. Marc Spieler. How AI Factories Can Help Relieve Grid Stress; NVIDIA, Santa Clara, California, https://blogs.nvidia.com/blog/ai-factories-flexible-power-use/ (2025).

68. Philip Colangelo, Ayse K. Coskun, Jack Megrue, Ciaran Roberts, Shayan Sengupta, Varun Sivaram, Ethan Tiao, Aroon Vijaykar, Chris Williams, Daniel C. Wilson, Zack MacFarland, Daniel Dreiling, Nathan Morey et al. Turning AI Data Centers into Grid-Interactive Assets: Results from a Field Demonstration in Phoenix, Arizona. arXiv:2507.00909 (2025). https://doi.org/10.48550/arXiv.2507.00909.

69. International Energy Agency (IEA). Energy and AI; IEA, Paris, France, https://www.iea.org/reports/energy-and-ai (2025).

70. Enel X Japan. Demand Response; Tokyo, Japan, https://www.enelx.com/jp/en/demand-response (Accessed August 2025).

71. Agency for Natural Resources and Energy (ANRE). Guidelines for Energy Resource Aggregation Business; ANRE, Ministry of Economy, Trade and Industry (METI), Tokyo, Japan, https://www.enecho.meti.go.jp/en/category/vpp_dr/data/guidelines_for_energy_resource_aggregation_business.pdf (2015 (Last revised June 2020)).

72. Emi Bertoli, Vida Rozite, Brendan Reidenbach, Anthony Vautrin, Sungjin Oh & Matthieu Suire. Demand Response; International Energy Agency (IEA), Paris, France, https://www.iea.org/energy-system/energy-efficiency-and-demand/demand-response (Accessed August 2025 (Last updated July 2023)).

73. International Energy Agency (IEA). Executive Summary in Integrating Solar and Wind (IEA, Paris, France, 2024, https://www.iea.org/reports/integrating-solar-and-wind/executive-summary).

74. California Independent System Operator (ISO). Monthly Renewables Performance Report: July 2024; California ISO, Folsom, California, https://www.caiso.com/documents/monthly-renewables-performance-report-july-2024.html (2024).

75. Ali Jahanshahi, Nanpeng Yu & Daniel Wong. PowerMorph: QoS-Aware Server Power Reshaping for Data Center Regulation Service. ACM Trans. Archit. Code Optim. 19, Article 36 (2022). https://doi.org/10.1145/3524129.

76. Sophie Hall, Francesco Micheli, Giuseppe Belgioioso, Ana Radovanović & Florian Dörfler. Carbon-Aware Computing for Data Centers with Probabilistic Performance Guarantees. arXiv:2410.21510 (2024). https://ui.adsabs.harvard.edu/abs/2024arXiv241021510H.

77. C. A. Silva, R. Vilaça, A. Pereira & R. J. Bessa. A review on the decarbonization of high-performance computing centers. Renewable and Sustainable Energy Reviews 189, 114019 (2024). https://doi.org/10.1016/j.rser.2023.114019.

78. Lori Aniti. Solar and wind power curtailments are increasing in California; US Energy Information Administration (EIA), Washington, DC, https://www.eia.gov/todayinenergy/detail.php?id=65364 (2025).

79. Jigar Shah. DOE Releases New Report on Pathways to Commercial Liftoff for Virtual Power Plants (Blog); US Department of Energy (DOE) Loan Programs Office, Washington, DC, https://www.energy.gov/lpo/articles/doe-releases-new-report-pathways-commercial-liftoff-virtual-power-plants (2023).

80. National Energy System Operator (NESO). Demand Flexibility Service (DFS) (Report); NESO, Warwick, United Kingdom, https://www.neso.energy/document/363911/download (2025).

81. National Energy System Operator (NESO). Demand Flexibility Service delivers electricity to power 10 million households; NESO, Warwick, United Kingdom, https://www.neso.energy/news/demand-flexibility-service-delivers-electricity-power-10-million-households (2023).

82. National Energy System Operator (NESO). Demand Flexibility Service (DFS) (Website); NESO, Warwick, United Kingdom, https://www.neso.energy/industry-information/balancing-services/demand-flexibility-service-dfs (Accessed August 2025).

83. Solar Power Portal. National Grid DSO releases two-year expansion plan; Solar Power Portal. Informa PLC, London, United Kingdom, https://www.solarpowerportal.co.uk/energy-policy/national-grid-dso-releases-two-year-expansion-plan (2024).

84. National Energy System Operator (NESO). Homepage; NESO, Warwick, United Kingdom, https://www.neso.energy/ (Accessed August 2025).

85. A. Arteconi, N. J. Hewitt & F. Polonara. State of the art of thermal storage for demand-side management. Applied Energy 93, 371-389 (2012). https://doi.org/10.1016/j.apenergy.2011.12.045.

86. Manajit Sengupta, Yu Xie, Anthony Lopez, Aron Habte, Galen Maclaurin & James Shelby. Solar Resource Maps and Data (https://www.nrel.gov/gis/solar-resource-maps) from The National Solar Radiation Data Base (NSRDB). Renewable and Sustainable Energy Reviews 89, 51-60 (2018). https://doi.org/10.1016/j.rser.2018.03.003.

87. Bloom Energy. Survey: 27% of data centers are expected to run entirely on onsite power by 2030; Latitute Studios, Parkville, Maryland, https://www.latitudemedia.com/news/survey-27-of-data-centers-are-expected-to-run-entirely-on-onsite-power-by-2030/ (2025).

88. Ethan Howland. FERC rejects interconnection pact for Talen-Amazon data center deal at nuclear plant (Dive Brief); Utility Dive (Informa TechTarget), Newton, Massachusetts, https://www.utilitydive.com/news/ferc-interconnection-isa-talen-amazon-data-center-susquehanna-exelon/731841/ (2024).

89. Avi Salzman. How Amazon's Nuclear Deal Could Solve a Touchy Political Problem; Barron's (Dow Jones & Company), New York, New York, https://www.barrons.com/articles/amazon-nuclear-deal-talen-energy-42e8b78b (2025).

90. Zachary Skidmore. Diversity of power - the biggest data center energy stories of 2024; Data Center Dynamics, London, United Kingdom, https://www.datacenterdynamics.com/en/analysis/diversity-of-power-the-biggest-data-center-energy-stories-of-2024/ (2025).

91. Grid Deployment Office. Building a Better Grid Awards | January 2025; US Department of Energy (DOE), Washington, DC, https://www.energy.gov/gdo/articles/building-better-grid-awards-january-2025 (2025).

92. Loans Program Office. LPO Financing Programs; US Department of Energy (DOE), Washington, DC, https://www.energy.gov/lpo/lpo-financing-programs (Accessed August 2025).

93. Mack Ramsden & Joshua Rosen. Publication: Order No. 1920: A Guide to FERC's Landmark Transmission Planning Order; Foley Hoag LLP, Boston, Massachusetts, https://foleyhoag.com/news-and-insights/publications/alerts-and-updates/2024/may/order-no-1920-a-guide-to-fercs-landmark-transmission-planning-order/ (2024).

94. GOV.UK. Collection: Contracts for Difference London, United Kingdom, https://www.gov.uk/government/collections/contracts-for-difference (Accessed August 2025).

95. US Environmental Protection Agency (EPA). Utility Green Tariffs; EPA, Washington, DC, https://www.epa.gov/green-power-markets/utility-green-tariffs (Accessed August 2025).

96. Cindy Sabato. State Corporation Commission Must Set Fair Rates and Protect Virginia Residents from Subsidizing Data Center Infrastructure; Piedmont Environmental Council, Warrenton, Virginia, https://www.pecva.org/work/energy-work/state-corporation-commission-must-set-fair-rates-and-protect-virginia-residents-from-subsidizing-data-center-infrastructure/ (2025).

97. Tom Kleckner. SPP, Hitachi Partner to Use AI in Clearing GI Queue; RTO Insider LLC (Yes Energy), Boulder, Colorado, https://www.rtoinsider.com/107367-spp-hitachi-partner-ai-clearing-gi-queue/ (2025).

98. Amil Merchant, Simon Batzner, Samuel S. Schoenholz, Muratahan Aykol, Gowoon Cheon & Ekin Dogus Cubuk. Scaling deep learning for materials discovery. Nature 624, 80-85 (2023). https://doi.org/10.1038/s41586-023-06735-9.

99. The Materials Project. Materials Explorer: MgV_2O_5; Lawrence Berkeley National Laboratory (LBNL), Berkeley, California, https://next-gen.materialsproject.org/materials/mp-19003 (Accessed August 2025).

100. Colin McCormick. Chapter 13: Materials Innovation in Artificial Intelligence for Climate Change Mitigation Roadmap (Second Edition) (ICEF Innovation Roadmap Project; https://doi.org/10.7916/2j4p-nw61) (2024, Chapter 3 found at https://www.icef.go.jp/wp-content/themes/icef_new/pdf/roadmap/2024/13_ICEF2.0%20Materials%20Innovation_stand%20alone.pdf).

101. Alice C. Hill and Colin McCormick. Chapter 14: Extreme Weather Response in Artificial Intelligence for Climate Change Mitigation Roadmap (Second Edition) (ICEF Innovation Roadmap Project; https://doi.org/10.7916/2j4p-nw61) (2024, Chapter 14 found at https://www.icef.go.jp/wp-content/themes/icef_new/pdf/roadmap/2024/14_ICEF2.0%20Extreme%20Weather%20Response_stand%20alone.pdf).

102. Remi Lam, Alvaro Sanchez-Gonzalez, Matthew Willson, Peter Wirnsberger, Meire Fortunato, Ferran Alet, Suman Ravuri, Timo Ewalds, Zach Eaton-Rosen, Weihua Hu, Alexander Merose, Stephan Hoyer, George Holland et al. Learning skillful medium-range global weather forecasting. Science 382, 1416-1421 (2023). https://doi.org/10.1126/science.adi2336.

103. National Oceanic and Atmospheric Administration (NOAA) National Centers for Environmental Prediction (NCEP). Global Forecast System (GFS); NOAA NCEP, College Park, Maryland, https://www.ncei.noaa.gov/products/weather-climate-models/global-forecast (Accessed August 2025).

104. Contact Energy. Contact signs ten-year renewable energy agreement with Microsoft on new geothermal power station at Te Huka; Contact Energy Ltd, New Zealand, https://contact.co.nz/investor-centre/news/2023/contact-signs-ten-year-renewable-energy-agreement-with-microsoft (2023).