Powering the Digital Era: State of the Art and Challenges

1st Nazmul Haque Anik
Electrical Engineering Department
University of New Mexico
nazmul757@unm.edu

2nd Masoud Davoudi
masoudd@ieee.org

3rd Ali Bidram
Electrical Engineering Department
University of New Mexico
bidram@unm.edu

***Abstract*—The rapid escalation of US data center electricity consumption has raised concerns among electric utilities, regulators, and the Artificial Intelligence (AI) industry regarding the grid's ability to support the emerging digital era. National demand from data centers has tripled over the past decade and is projected to double or triple again by 2028, driven by compute-intensive AI workloads and parallel growth in other large, fast-changing loads. This paper examines the implications of this unprecedented demand surge across three key pillars of the digital-era energy ecosystem: utilities, regulators and standards organizations, and AI data-center developers. For each pillar, we assess the current state of the art, identify major technical and regulatory challenges, and highlight critical gaps in visibility, coordination, interconnection processes, operational readiness, and resource planning. The paper provides a discussion of emerging issues and potential improvements needed to ensure that future large-load growth can be integrated reliably and sustainably as AI-driven electricity demand continues to expand.**

***Index Terms*—Artificial Intelligence, Data Centers, Energy Demand, Large Loads**

I. Introduction

Recent analyses highlight the rapid acceleration of US data-center electricity consumption. According to the 2024 US Data Center Energy Use Report, published by the US Department of Energy (DOE) and Lawrence Berkeley National Laboratory (LBNL), US data centers consumed 176 TWh in 2023, representing approximately 4.4% of total national electricity use [1]. Projections indicate that this demand could reach 6.7% to 12.0% of all US electricity consumption by 2028 [1]. The report also notes that data-center electricity use increased from 58 TWh in 2014 to 176 TWh in 2023 and is expected to reach 325–580 TWh by 2028 [1]. At the component level, power requirements are rising as well: CPUs now draw 150–350 W and advanced GPUs demand 350–700 W [2], reinforcing expectations of continued rapid growth in total data-center energy use. DOE–LBNL projections assume rapid AI growth, whereas Congressional Research Service(CRS) assessments highlight the dependence of future demand on efficiency, policy, and infrastructure constraints.

A. Beyond Data Centers

Several other emerging sectors are positioned to become major electricity consumers in the coming years. The (North American Electric Reliability Corporation) NERC Large Load Task Force (LLTF) [3] defines a large load as a commercial or industrial facility, or a collection of facilities, behind one or more points of interconnection, whose size or operating characteristics may introduce reliability challenges for the Bulk Power System (BPS). Examples include cryptocurrency-mining facilities, EV-charging infrastructure, hydrogen-production electrolyzers, electric-arc-furnace (EAF) steel production, advanced manufacturing, and (liquefied natural gas) LNG export terminals. These diverse loads require updated interconnection processes and coordinated planning to ensure system reliability as they scale.

US cryptocurrency-mining electricity consumption has increased substantially in recent years, and the US Energy Information Administration estimates that it now accounts for roughly 0.6% to 2.3% of annual US electricity use [4]. Large, fast-changing loads of this type can affect frequency, voltage stability, and system dynamics, requiring enhanced interconnection requirements, more rigorous planning analyses, and improved real-time coordination.

Electrification trends further amplify expected load growth. The National Renewable Energy Laboratory projects that by 2030, approximately 33 million electric vehicles (EVs) will be on US roads, supported by an estimated 28 million charging ports [5]. Hydrogen electrolyzers could expand domestic capacity nearly forty-fold [6], and industrial electricity consumption grew 6% from 2018 to 2022 [7]. These rapid increases necessitate more advanced planning tools and cross-sector coordination, so agencies should develop integrated load forecasts that incorporate transportation, industry, and digital infrastructure.

B. A Holistic Necessity

Meeting these growing demands requires a holistic understanding of cross-resource interactions, not solely electricity consumption. For instance, the International Energy Agency estimates that a 100 MW US data center uses roughly as much water as 2,600 households [2], underscoring the substantial cooling-related water requirements associated with large-scale digital infrastructure. Given these multidimensional resource demands, planners should integrate land, water, fiber, and grid constraints into siting and permitting processes to minimize conflicts and accelerate development.

979-8-3315-8138-1

C. This Paper

The scale and pace of these emerging loads highlight the need for improved approaches to planning, interconnection, operational coordination, and cross-resource management. These challenges lie at the core of the 2026 IEEE PES General Meeting theme, *Powering the Digital Era*, which emphasizes the technical and institutional capabilities required to support a rapidly expanding digital and industrial ecosystem. To support this vision, this paper examines the implications of accelerating AI-era electricity demand across technical, regulatory, and stakeholder dimensions. The concluding section summarizes cross-sector needs and highlights areas requiring further research to maintain reliability as large loads continue to scale.

II. Role and Challenges of Utilities

Utilities are in charge of maintaining grid reliability and stability while allowing for the rapid growth of data centers. They face difficulties due to rapidly shifting demand profiles and the absence of reliable data center load models for operations and planning [8]. Large data centers impose unprecedented demands on transmission and distribution systems. Their rapid development cycles often outpace utility planning, as major upgrades typically require five to ten years to complete, while data centers can be constructed in one to two years [9]. This mismatch can lead to localized congestion and insufficient hosting capacity, so utilities should establish expedited but rigorous assessment processes for large-load interconnection requests. This section provides an overview of the main challenges utilities face as they adapt to rapidly expanding, fast-changing digital-era loads.

A. Interconnection Requirements and Performance Standards

To uphold reliability requirements and mitigate risks from large loads, NERC recommends that Transmission Owners strengthen their interconnection practices. This includes establishing clear facility design and performance criteria and requiring installation of high-speed disturbance monitoring devices such as Phasor Measurement Units (PMUs) to capture and assess load behavior during disturbances. These efforts can be enhanced by requiring real-time operational telemetry and dynamic models as conditions of interconnection. Utilities could also establish standardized large-load interconnection templates to streamline reviews and reduce uncertainty for developers.

B. Operational Uncertainty and Forecasting Challenges

AI data centers introduce significant uncertainty in both short- and long-term forecasting, as their workloads fluctuate rapidly with transitions between training, inference, and idle modes. These patterns can force grid operators to maintain higher spinning and non-spinning reserves, increasing operational costs and requiring flexible backup resources. Utilities can improve reserve planning by incorporating data-center-specific forecasting tools and workload visibility into operational models.

C. Planning Constraints and Infrastructure Expansion Needs

Rapid AI-driven load growth is also outpacing power generation and grid expansion. In the US, NERC warns that half of the national grid faces reliability risks as demand rises. For example, 69 proposed data centers in Georgia could increase peak load from 16.5 GW to 31 GW while capacity additions lag behind. Utilities can address this by partnering with regulators to pursue proactive generation and transmission expansion plans aligned with projected digital-era load growth.In addition, technical constraints—such as transformer limits, substation capacity, and protection system requirements—often delay interconnection. These issues can require costly upgrades or network reconfiguration, making them a major driver of utility planning and permitting timelines.

D. Stability and Protection Considerations

Large data centers can exhibit rapid load ramping on the order of tens of megawatts within seconds to minutes, particularly during transitions between AI training, inference, and idle [10]. Such large, abrupt AI load ramps can stress grid stability through steady-state voltage fluctuations, frequency dips in low-inertia systems, and excitation of oscillatory or resonance modes. These stability challenges require more sophisticated dynamic modeling, so utilities should develop EMT-based analysis capabilities and enforce performance standards that limit rapid load transitions.

E. Harmonics and Protection Considerations

Power-electronic interfaces associated with GPUs, rectifiers, and uninterruptible power supplies (UPS) may also introduce harmonic distortion and fast transient behavior, which can affect voltage quality, as well as protection coordination. These characteristics underscore the importance of considering representative ramp rates, harmonic profiles, and converter dynamics when assessing grid stability impacts of large, fast-changing data center loads.

F. Flexibility Mechanisms and Demand-Response Options

Utilities are increasingly investigating contract-based and incentive-based demand response arrangements in which data centers provide predetermined flexibility in exchange for compensation or priority service. These programs can reduce operational stress and improve grid balancing, so utilities should expand flexible interconnection agreements that define minimum ramping capabilities or temporary curtailment provisions.

III. Regulatory, Standards, and Policy Landscape for AI and Large Emerging Loads

The regulatory framework for AI data centers and other emerging large loads is expanding quickly but remains inconsistent across federal, state, and regional jurisdictions. No comprehensive U.S. legislation currently governs AI infrastructure directly, leaving integration challenges to be addressed through executive actions, agency guidance, and regional

979-8-3315-8138-1

reliability standards. Policymakers can strengthen national consistency by developing baseline requirements for large-load transparency, interconnection processes, and performance expectations.

A. Federal Actions Shaping Early AI Infrastructure Deployment

The evolving federal landscape is shaping how large data centers and AI facilities are permitted, financed, and integrated. Executive Order (EO) 14141 [11] emphasized U.S. leadership in AI infrastructure before being superseded by EO 14318 [12], which accelerates federal permitting for major data-center projects. Projects exceeding $500 million or 100 MW of additional load receive streamlined environmental review, financial support, and expedited access to federal lands. These steps can accelerate development, but federal agencies should pair them with energy-impact reviews to ensure the electric grid is prepared for the associated load growth.

B. State-Level Policies and Regional Differences

States are rapidly advancing their own policies to manage large-load growth, grid impacts, and local economic priorities. Texas created a Large Flexible Load Task Force and passed SB 6 [13], granting ERCOT authority to curtail very large loads during emergencies [14]. Virginia's rapid data-center expansion faced setbacks after zoning litigation halted development [16], while Georgia reconsidered its tax incentives in 2024 [17]. Colorado's HB 1177 authorizes utilities to offer discounted economic-development electricity rates to data centers, provided the discounts cover marginal costs and avoid shifting expenses to other customers [18]. Maryland's Critical Infrastructure Streamlining Act (2024) classifies data centers as critical infrastructure and offers faster pathways for permitting backup generation following earlier denials by the Maryland PSC [18]. These cases highlight growing tension between economic development goals and grid constraints, emphasizing the need for states to harmonize incentive programs with regional grid capacity assessments.

C. Congressional and Federal Regulatory Engagement

Congressional activity is increasing as large loads become a national concern. The Bipartisan AI Task Force's 2024 report [19] identified infrastructure and energy impacts as priority issues. FERC opened a 2025 proceeding to evaluate co-location of major loads with generation in PJM [20] and issued guidance encouraging all RTOs/ ISOs to adopt enhanced forecasting methods for large loads [21]. Establishing standardized forecasting and reporting requirements across all balancing authorities would significantly improve oversight and planning.

D. Industry Reliability Efforts and Technical Guidance

Industry organizations are strengthening technical guidance to support reliable integration of AI-driven loads. NERC's Large Load Task Force (LLTF) is developing recommendations for integrating AI, cryptocurrency, hydrogen production, and industrial electrification [22]. Although Rule 810 [3] requires registered entities to respond to NERC Alerts, many data centers are not registered entities, limiting situational awareness. NERC could mitigate this gap by creating a category of "observed large loads" requiring minimum telemetry and disturbance reporting standards.

E. Interconnection Constraints

Interconnection and permitting delays are among the largest challenges nationwide, as construction timelines for data centers are far shorter than the timelines for grid upgrades. In several cases, FERC has rejected interconnection requests due to reliability concerns. Regulators can ease pressure on interconnection queues by adopting fast-track processes for large loads co-located with new or existing generation.

F. Geographic Concentration and Regional Planning Needs

Approximately 80% of U.S. data-center load is concentrated in fifteen states [23], creating geographic clustering that intensifies local grid stress and complicates cost allocation for needed transmission upgrades. Regional planning processes should account for the cumulative impacts of these clusters, ensuring long-term infrastructure investments reflect expected load growth patterns.

G. Performance Standards, Operational Requirements, and Tariff Modernization

Beyond infrastructure build-out, data centers require clear performance standards to ensure they ride through short voltage or frequency disturbances and provide sufficient workload information for operational planning. Standards bodies can help by developing minimum ride-through, ramp-rate, and telemetry requirements. At the same time, tariff structures and cost-allocation rules need modernization: inconsistent rate designs and interconnection requirements across regions create uncertainty for both utilities and developers. More transparent and consistent cost-sharing frameworks would improve predictability and reduce delays.

H. Final Verdict: Toward a More Coordinated National Framework

Taken together, current developments highlight the need for more coordinated policy frameworks as AI-driven loads expand rapidly. Today's environment remains incomplete and regionally inconsistent. Strengthening coordination among federal, state, and regional entities would improve development timelines, enhance reliability outcomes, and provide greater clarity as large loads continue to scale nationwide.

IV. Role and Actions of Data Centers

Not all data centers impose identical electrical characteristics on the grid. General-purpose data centers, which support services such as web hosting, storage, and enterprise applications, typically exhibit relatively stable and predictable load profiles with moderate power density. In contrast, GPU-heavy AI data centers rely on accelerator-dense clusters that introduce significantly higher power density and rapid load

979-8-3315-8138-1

variability due to frequent transitions between training, inference, and idle states. These differences lead to distinct electrical behaviors, particularly in terms of ramp rates, power quality, and cooling demand, and therefore require differentiated consideration in grid planning and operational studies. Developers can mitigate these impacts through workload smoothing, improved forecasting, high-fidelity modeling, efficient cooling strategies, and on-site resources that enhance flexibility and resilience.

A. Smoothing Workloads and Managing Rapid Ramps

Rapid load changes are one of the most challenging aspects of AI-driven data centers. Workload-smoothing algorithms can reduce sudden ramp events by distributing compute demand more evenly across clusters. [24] demonstrates that adjusting workloads through an MDP-based stochastic optimization framework—using uncertainty models for both energy and demand—can significantly improve renewable-energy matching, outperforming previous methods by up to 55%. Data centers adopting similar approaches can substantially reduce the volatility of their demand on the grid.

B. Improving Predictability Through Machine-Learning Forecasting

Machine learning–based forecasting enables data centers to anticipate short-term fluctuations by analyzing telemetry, hardware utilization, thermal conditions, and historical workload patterns. Sharing aggregated or anonymized forecast outputs with utilities enhances situational awareness without compromising proprietary information. When integrated into real-time control systems, these forecasts allow operators to proactively limit abrupt load changes and prepare auxiliary systems before transitions occur.

C. Providing High-Fidelity Models for Grid Integration

High-fidelity modeling is essential for accurate assessment of data-center behavior during interconnection studies. Electromagnetic transient (EMT) models capture sub-cycle converter dynamics, UPS switching characteristics, cooling-motor behavior, and pulsed compute loads. Providing validated EMT or hybrid models enables utilities and RTOs/ISOs to perform more realistic stability analyses. Developers who deliver detailed models early in the interconnection process can significantly accelerate study timelines and reduce uncertainty.

D. Enhancing Energy Efficiency of AI Hardware and Workloads

AI workloads are increasingly energy intensive, but efficiency improvements can mitigate their impact. Using lighter or task-specific AI models, deploying more efficient accelerators, and optimizing workload scheduling reduce peak power clustering and total energy use. Better job orchestration also smooths compute demand across time, lowering the strain on the local grid and improving overall facility efficiency.

E. Optimizing Cooling Systems and Flexible Auxiliary Loads

Cooling represents the second-largest [10] share of a data center's electricity use. Advanced approaches such as hybrid air–liquid systems, direct liquid cooling, immersion cooling, and AI-enabled predictive cooling can materially reduce auxiliary load. Operators can further improve grid compatibility by designing cooling systems capable of flexible demand profiles, shifting or staggering loads to avoid peak periods or stressed grid conditions. Apart from the load burden on the grid, cooling systems can also influence grid performance. Large cooling motors and pumps are often controlled using variable frequency drives (VFDs), which may introduce harmonic distortion and fast power variations into the local network. In addition, motor starting and transient events associated with chillers, pumps, and auxiliary cooling equipment can cause short-term voltage fluctuations, particularly in high-density AI data centers. These effects highlight the need to account for cooling-related electrical behavior when evaluating overall grid impacts of large data center facilities.

F. Leveraging On-Site Energy Resources for Reliability and Flexibility

On-site resources such as battery energy storage systems (BESS), solar PV, wind, gas turbines, and fuel cells can provide essential flexibility and resilience. Integrating BESS within UPS or PDU architectures supports seamless ride-through and frequency response while offering zero-emission alternatives to diesel backup. Developers can collaborate with utilities to dispatch storage strategically during peak periods, align operations with wholesale-market signals, and reduce reliance on the bulk grid. Large-scale battery systems also enable load shifting and can reduce the need for continuous generator use. Coordinated use of storage enhances both grid support and the facility's operational resilience.

G. Exploring Nuclear Options and Emerging Baseload Solutions

Small Modular Reactors (SMRs), typically ranging from 50 MW to 300 MW, are emerging as long-term baseload options for AI clusters. Recent industry initiatives illustrate that this transition is moving from concept to early deployment. For example, Google and Kairos Power have agreed to develop a 500 MW fleet of molten-salt SMRs by 2035 to supply 24/7 carbon-free electricity to Google data centers [25]. Co-locating SMRs with large compute hubs could reduce transmission requirements and offer stable, carbon-free power, though regulatory frameworks and siting processes must continue to evolve before widespread deployment becomes feasible. From a grid-integration perspective, co-located or dedicated SMRs would need to comply with existing grid protection and operational requirements, including fault ride-through capability, coordinated protection schemes, and clear islanding and reconnection procedures. Depending on the interconnection configuration, SMRs may operate as grid-connected baseload resources or in hybrid islanded modes supporting critical data center loads. These considerations highlight the importance of aligning

979-8-3315-8138-1

SMR deployment with established protection standards and operational coordination frameworks to ensure reliable and secure integration.

V. Conclusion

AI-era electricity demand is expanding at an unprecedented pace, reshaping the operational, planning, and regulatory requirements of the electric power system. Utilities must adapt their interconnection processes, forecasting methods, stability analyses, and infrastructure investment strategies to accommodate large, fast-changing loads whose development timelines often outpace traditional planning cycles. Regulators and standards organizations play an equally critical role by modernizing policy frameworks, harmonizing state and federal requirements, and establishing consistent expectations for transparency, performance, forecasting, and cost allocation. At the same time, AI data centers must improve their operational compatibility with the grid by enhancing workload predictability, providing high-fidelity dynamic models, increasing efficiency, and deploying flexible on-site resources that support reliability.

Taken together, these three pillars (utilities, regulators and standards bodies, and data centers) form an interdependent ecosystem that must evolve in parallel to ensure reliable and sustainable integration of digital-era loads. Addressing the visibility gaps, interconnection challenges, and coordination barriers identified in this paper will be essential to maintaining grid resilience as data-center and other emerging large loads continue to scale. Continued collaboration, technical innovation, transparent data-sharing practices, and forward-looking regulatory structures will be central to powering the digital era while safeguarding the stability of the Bulk Power System.

VI. Acknowledgment

This paper is supported by National Science Foundation under awards ECCS 2338555 and OSIE 2330582, and Department of Energy under award number DE-CR0000117.

References

[1] U.S. Department of Energy and Lawrence Berkeley National Laboratory, "DOE Releases New Report on Rising Data Center Electricity Demand," 2024. [Online]. Available: https://tinyurl.com/rbddntmj

[2] E. D. Vann, "Data Centers and the Electric Grid: Trends, Challenges, and Policy Considerations," *Congressional Research Service*, Rep. R48646, Aug. 15, 2024. [Online]. Available: https://www.congress.gov/crs-product/R48646

[3] North American Electric Reliability Corporation (NERC), "NERC Alert Level 2: Large Loads," *North American Electric Reliability Corporation*, Atlanta, GA, USA, Feb. 2025. [Online]. Available: https://www.nerc.com/globalassets/programs/bpsa/alerts/2025/nerc-alert-level-2--large-loads.pdf

[4] U.S. Energy Information Administration, "Data centers and AI are driving significant increases in U.S. electricity demand," *Today in Energy*, June 3, 2024. [Online]. Available: https://www.eia.gov/todayinenergy/detail.php?id=61364

[5] U.S. Department of Energy, "Fact of the Week #1334: By 2030, the U.S. Will Need 28 Million EV Charging Ports," Mar. 18, 2024. [Online]. Available: https://tinyurl.com/d2wypnv5

[6] U.S. Department of Energy, "Electrolyzer Installations in the United States," DOE Hydrogen and Fuel Cell Technologies Office, Rep. DOE/EE-24001, Sept. 2024. [Online]. Available: https://www.hydrogen.energy.gov/docs/hydrogenprogramlibraries/pdfs/24001-electrolyzer-installations-united-states.pdf

[7] U.S. Energy Information Administration, "Manufacturing Energy Consumption Survey (MECS)," 2024. [Online]. Available: https://www.eia.gov/consumption/manufacturing/

[8] Idaho National Laboratory, "Navigating Integration: Addressing Challenges of Interconnecting Data Centers and Other Large Electric Loads," *Idaho National Laboratory*, Idaho Falls, ID, USA, Nov. 2025. [Online]. Available: https://csdet.inl.gov/content/uploads/45/2025/11/Navigating-Integration-report.pdf

[9] S. Zhao, R. Li, and Y. Zhang, "AI Data Centers and the Future of Grid Stability: A Comprehensive Review," *arXiv preprint* arXiv:2509.07218, 2025. [Online]. Available: https://arxiv.org/abs/2509.07218

[10] A. Author *et al.*, "Electricity Demand and Grid Impacts of AI Data Centers: Challenges and Prospects," 2024. [Online]. Available: https://www.researchgate.net/publication/395352695

[11] Executive Office of the President, "Advancing United States Leadership in Artificial Intelligence Infrastructure," *Federal Register*, vol. 90, no. 12, Jan. 17, 2025. [Online]. Available: https://tinyurl.com/3cuc4nm9

[12] The White House, "Accelerating Federal Permitting of Data Center Infrastructure," July 2025. [Online]. Available: https://tinyurl.com/38x3ur8m

[13] Texas Legislature, "Senate Bill No. 6 (SB 6): Texas Energy Insurance Program and Electric Grid Reliability," Austin, TX, 2023. [Online]. Available: https://legiscan.com/TX/text/SB6/id/3182174

[14] R. Behrens, "Texas law gives grid operator power to disconnect data centers during crises," *Utility Dive*, Jun. 17, 2024. [Online]. Available: https://www.utilitydive.com/news/751587/

[15] Prince William County Planning Office, "Prince William Digital Gateway Project," 2025. [Online]. Available: https://www.pwcva.gov/department/planning-office/pw-digital-gateway

[16] M. Anzilotti, "Judge voids Digital Gateway rezoning in Prince William County," *WTOP News*, Aug. 19, 2025. [Online]. Available: https://wtop.com/prince-william-county/2025/08/judge-voids-digital-gateway-rezoning-in-prince-william-county

[17] J. Amy, "Georgia governor vetoes tax breaks for data centers amid strain on power grid," *Associated Press News*, May 8, 2024. [Online]. Available: https://apnews.com/ee11e95f3a5ac9f6e9019401c35c2e88

[18] Hogan Lovells, "State Legislative and Regulatory Initiatives to Address Concerns Caused by the Surging Power Demand," *Hogan Lovells Publications*, Aug. 2024. [Online]. Available: https://www.hoganlovells.com/en/publications/state-legislative-and-regulatory-initiatives

[19] DLA Piper, "An AI Blueprint for Future Congressional Action," *DLA Piper AI Outlook*, 2024. [Online]. Available: https://tinyurl.com/tarr7bx2

[20] Federal Energy Regulatory Commission, "FERC Orders Action on Co-Location Issues Related to Data Centers Running AI," Sept. 19, 2024. [Online]. Available: https://tinyurl.com/y4fbmmzy

[21] Federal Energy Regulatory Commission, "Chairman Rosner's Letter to RTOs/ISOs on Large Load Forecasting," Jan. 22, 2025. [Online]. Available: https://tinyurl.com/3rfp3656

[22] North American Electric Reliability Corporation, "Large Load Task Force (LLTF) Initiative," 2024. [Online]. Available: https://www.nerc.com/comm/RSTC/Pages/LLTF-Initiative.aspx

[23] Wisconsin Policy Research Institute, "Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption," 2024. [Online]. Available: https://www.wpr.org/wp-content/uploads/2024/06/3002028905_Powering-Intelligence_-Analyzing-Artificial-Intelligence-and-Data-Center-Energy-Consumption.pdf

[24] Q. Huang, Y. Wang, and S. Santoso, "Grid Impacts of Data Center Integration: Insights and Analysis," *IEEE Power and Energy Magazine*, vol. 22, no. 5, pp. 46–61, Sept.–Oct. 2024. [Online]. Available: https://ieeexplore.ieee.org/document/10738452

[25] S. Cohn, "Google Bets Big on Nuclear: Inks Deal with Kairos Power for 500 MW SMR Fleet to Power Data Centers," *POWER Magazine*, Oct. 10, 2024. [Online]. Available: https://tinyurl.com/4twr3msu

979-8-3315-8138-1