Energy Recovery and Resource Optimization in Data Centers: Transforming Unused Power into Economic Value

Sowjanya Addu*

Computer Science & Engineering, Gokaraju Rangaraju Institute of Engineering Technology, Hyderabad, Telangana, India

Abstract

Data centers are among the most capital-intensive and energy-intensive facilities in the modern economy, yet much of the power they pay for goes unused. Electrical capacity is provisioned for peaks and redundancy that rarely occur, servers run at low utilization, renewable energy is curtailed on the grids that serve them, and the heat they produce is usually released into the air. Anioke (2026) describes unused electricity and capacity of this kind as stranded power. With global data center electricity use projected to rise from about 415 TWh in 2024 to around 945 TWh in 2030 (International Energy Agency [IEA], 2025), the cost of this waste is growing. This article examines how energy recovery and resource optimization can turn unused power into economic value. It identifies five value streams: recovering stranded capacity inside facilities, converting surplus renewable energy into low-cost computing, selling flexibility to the grid, optimizing energy costs through intelligent scheduling, and reusing waste heat. It also considers the role of enterprise analytics in capturing this value, the barriers that remain, and recommendations for operators, utilities, and policymakers.

Journal of Science, Technology and Social Transformation (2026)

Introduction

The economics of data centers are dominated by power. Anioke (2025) shows that the rapid growth of data centers is raising total electricity demand, creating new risks for grid stability, and complicating the transition to sustainable energy. In many regions, the scarcest resource for a new data center is no longer land or capital but a grid connection.

In this environment, unused power is not just an engineering inefficiency. It is a financial loss. Every megawatt of capacity that is provisioned but never drawn represents capital that earns no return. Every kilowatt-hour of renewable energy that is curtailed is revenue lost to a generator and clean power lost to the system. Every unit of waste heat released into the atmosphere is energy that could have displaced fuel elsewhere.

Anioke (2026) argues that stranded power from data centers has a range of possible uses and that finding them can ease grid pressure while supporting sustainability. This article takes an economic perspective on that argument. It asks how much value is locked up in unused power, how that value can be recovered, and what operators, utilities, and policymakers can do to capture it.

Corresponding Author: Sowjanya Addu, Computer Science & Engineering, Gokaraju Rangaraju Institute of Engineering Technology, Hyderabad, Telangana, India, Email: sowjanya1634@grietcollege.com

How to cite this article: Addu, S. (2026). Energy Recovery and Resource Optimization in Data Centers: Transforming Unused Power into Economic Value. *Journal of Science, Technology and Social Transformation* 2(4), 67-70.

Source of support: Nil

Conflict of interest: None

The Economics of Unused Power

To understand the value of recovering unused power, it helps to understand why so much of it exists.

Capital tied to provisioned capacity. Barroso et al. (2019) explain that a large share of the capital cost of a data center facility scales with its provisioned critical power capacity: the transformers, switchgear, uninterruptible power supplies, generators, and cooling plant needed to support the IT load. This capacity is sized for the maximum load the facility might ever draw, plus redundancy. If the actual load is much lower, the operator has paid for infrastructure that sits idle.

The gap between provisioned and actual power. Fan et al. (2007), studying large Google server clusters, found a consistent gap between the power provisioned and the power actually consumed, even at peak. Servers rarely draw their maximum power at the same time, so the combined peak of thousands of machines is well below the sum of their individual ratings. Capacity planned from nameplate ratings therefore strands a significant share of the investment.

Low utilization of computing resources. Barroso et al. (2019) also note that servers in many data centers spend much of their time at low utilization, and that servers consume a substantial share of their peak power even when lightly used. Low utilization means that both the capital invested in servers and the energy they consume deliver less useful work than they could.

Lost value on the grid. Outside the facility, wind and solar output is curtailed when grids cannot carry it. Kim et al. (2017) showed that significant stranded power persists in power systems and grows as the share of wind generation increases. Meanwhile, transmission and distribution networks sit well below their capacity for most hours of the year.

Discarded heat. Almost all the electricity a data center uses ends up as heat. In most facilities, this heat is removed by cooling systems and released into the air, even though it could be used for heating buildings or industrial processes (Wahlroos et al., 2018).

Value Stream 1: Recovering Stranded Capacity Inside Facilities

The most direct way to create value from unused power is to use the capacity a facility already has more fully.

Fan et al. (2007) showed that by measuring actual power use and setting limits on peak consumption, operators can safely deploy substantially more servers within the same provisioned power budget. This practice, called power oversubscription, relies on power capping, so that servers can be slowed briefly if total demand approaches the limit. The economic effect is significant. Because facility capital costs scale with provisioned capacity (Barroso et al., 2019), fitting more computing into the same capacity lowers the capital cost per unit of computing delivered.

Raising server utilization has a similar effect. Consolidating workloads onto fewer, busier machines lets operators retire or repurpose idle servers and reduces the energy wasted on lightly loaded equipment. Barroso et al. (2019) describe this as making better use of the warehouse-scale computer as a whole, rather than treating each server as a separate asset.

Anioke (2026) also points to the power capacity data centers hold in reserve for redundancy. With careful design, some of this reserve can support lower-priority workloads that can be shut down immediately if a failure occurs. The reserve remains available for emergencies but earns a return the rest of the time.

Recovering capacity in these ways also has a strategic value. Anioke (2026) notes that it can defer the need for new grid connections. With the IEA (2025) estimating that grid constraints could delay about 20% of planned data center projects, the ability to grow within an existing connection can be worth far more than the energy saved.

Value Stream 2: Converting Surplus Renewable Energy into Low-Cost Computing

The second value stream turns curtailed renewable energy into a productive input. When wind or solar output exceeds what the grid can absorb, prices in some markets fall to very low or even negative levels. A data center that can increase its computing at these times can buy power cheaply while giving renewable generators a market for output that would otherwise be lost.

Kim et al. (2017) examined how data centers could harness this stranded power. They found that simply building data centers next to wind farms did not reliably reduce stranded power and could even increase it in some cases. The benefit appeared when data centers were treated as dispatchable loads, able to raise or lower demand in response to grid conditions, and placed optimally across the system. Under this approach, flexible data centers reduced stranded power, lowered system costs, and ran at a large share of their capacity on energy that would otherwise have been wasted. Kim et al. (2017) concluded that this gives data center owners a significant economic incentive to adopt flexible computing.

Anioke (2026) identifies several practical uses for stranded renewable power, including flexible computing placed near wind and solar sites. Suitable workloads include AI model training, scientific simulation, large-scale data processing, and other batch tasks that do not need an immediate result. For these workloads, the main economic question is whether the savings on power outweigh the cost of hardware that may sit idle when surplus energy is not available.

Value Stream 3: Selling Flexibility to the Grid

The third value stream treats a data center's ability to adjust its demand as a product that can be sold.

Grid stress is concentrated in a small number of peak hours each year. Norris et al. (2025) examined the 22 largest balancing authorities in the United States and found that existing systems could integrate about 76 GW of new load if

that load accepted curtailment for an average of just 0.25% of the year. At 0.5% curtailment, the figure rose to 98 GW, and at 1.0% to 126 GW.

For a data center developer, this creates two sources of value. The first is faster access to the grid. A facility that agrees to reduce its demand during peak hours may be connected years earlier than one that requires firm capacity at all times. The second is direct payment. In many markets, customers that reduce demand on request are paid through demand response programs or capacity markets.

The IEA (2025) notes that many AI data centers are installing large battery systems to smooth the rapid power swings of their workloads, and that these batteries could also provide services to the wider grid if the incentives are right. A battery that earns revenue from grid services while also protecting the facility turns a cost center into a source of income.

Value Stream 4: Optimizing Energy Costs Through Intelligent Scheduling

Electricity prices and carbon intensity vary by hour and by location. Intelligent scheduling allows data centers to move flexible work toward the cheapest and cleanest power.

Radovanović et al. (2023) describe Google's carbon-intelligent computing system, which uses forecasts of next-day grid carbon intensity and of computing demand at each data center cluster. It then sets hourly limits on available capacity, delaying flexible tasks to hours when cleaner power is available. The system operates across Google's global fleet. It also lowers peak power demand at each site, which can reduce demand charges and the capacity a facility must reserve.

The same approach can target electricity prices directly. By shifting flexible workloads away from expensive peak hours and toward periods of low prices, often the same periods in which renewable output is high, operators can reduce energy costs without reducing the total work done.

Value Stream 5: Recovering and Reusing Waste Heat

The fifth value stream recovers energy that has already been used. Nearly all the electricity consumed by a data center is converted into heat. Instead of releasing it, operators can capture it and supply it to district heating networks, greenhouses, or industrial users.

Wahlroos et al. (2018) studied the potential for using data center waste heat in Northern Europe, where district heating networks are common. They found that operators, particularly in Finland, were planning to feed waste heat into district heating systems, and that when heat pumps are used to raise its temperature, the investment can pay back within the expected lifetime of the equipment. They also found that the environmental benefit depends on which fuel the recovered heat replaces.

The same study identified barriers. Data center heat is often low in temperature and its supply can be unpredictable, investment costs can be high, and business arrangements between data center operators and heating companies are often unclear (Wahlroos et al., 2018). Waste heat recovery therefore works best where a heat network or heat user is nearby and where long-term agreements can be put in place.

The Enabling Role of Enterprise Analytics

Capturing these value streams requires good information. Operators need to know how much capacity is provisioned and used, what power costs at each hour, how much flexibility they can offer, and how much heat they produce. They also need to link this information to financial decisions.

Enterprise systems can provide this link. Asuri (2019) examines how embedded analytics and modernized user interfaces within enterprise resource planning (ERP) systems can improve visibility and operational performance. Asuri (2023) describes a framework for integrating large language models with enterprise platforms to enable intelligent automation. Applied to data center energy management, such tools could bring power, asset, workload, and cost data together, allowing managers to see the financial value of stranded capacity and act on it routinely. For example, analytics could show the cost per unit of computing at each site, track revenue from flexibility programs, and flag servers or racks whose provisioned power is consistently unused.

Anioke (2025) notes that fair cost allocation is central to the sustainability of data center growth. Better enterprise data also supports this goal, by making clear which costs data center operators impose on the grid and which value they return to it.

Barriers to Capturing Value

Several barriers limit how much value can be recovered from unused power.

Reliability and risk. Data center operators place uptime above almost everything else. Oversubscription, reserve capacity use, and flexibility commitments all carry some risk and depend on reliable monitoring and control (Fan et al., 2007).

Capital intensity. AI data centers are very capital-intensive, so operators are reluctant to leave expensive hardware idle (IEA, 2025). This limits how often they can pause work to follow cheap power or provide flexibility.

Market design. Many interconnection processes offer only firm capacity, and many markets do not pay fairly for flexible demand or allow facilities to buy curtailed power directly.

Location. Sites with abundant curtailed renewable energy, and sites with heat users nearby, are not always where data centers can be built.

Information gaps. Many facilities lack detailed power and cost data, and many organizations treat energy, facilities, and IT as separate budgets, hiding the full value of recovery.

Recommendations

Seven steps would help data centers turn unused power into economic value:

1. **Measure and price unused capacity,** linking power data to financial reporting through enterprise analytics (Asuri, 2019, 2023).
2. **Provision power based on measured demand,** using oversubscription and power capping to recover stranded capacity (Barroso et al., 2019; Fan et al., 2007).
3. **Raise server utilization** through workload consolidation and shared infrastructure (Barroso et al., 2019).
4. **Operate flexible workloads as dispatchable loads,** using surplus renewable power where it is available (Anioke, 2026; Kim et al., 2017).
5. **Participate in flexibility markets** and pursue flexible interconnection agreements (IEA, 2025; Norris et al., 2025).
6. **Adopt price-aware and carbon-aware scheduling** for flexible workloads (Radovanović et al., 2023).
7. **Plan for waste heat recovery** where heat networks or users are nearby, supported by clear long-term agreements (Wahlroos et al., 2018).

Policymakers can support these steps by reforming interconnection rules and market designs, so that flexible demand, curtailed renewable power, and recovered heat can be bought and sold on fair terms (Anioke, 2025).

Conclusion

Unused power is one of the largest hidden costs in the data center industry. Capacity is provisioned for peaks that rarely occur, servers run below their potential, renewable energy is curtailed on the grid, and heat is released into the air. Each of these represents economic value that is currently lost. Energy recovery and resource optimization offer ways to capture it. Oversubscription and higher utilization recover capital tied up in stranded capacity. Dispatchable computing turns surplus renewable energy into low-cost work. Flexibility earns faster grid access and direct revenue. Intelligent scheduling lowers energy costs and emissions together. Waste heat recovery turns a by-product into a product. None of these strategies is free of risk, and each depends on good data, sound contracts, and supportive market rules. Taken together, however, they show that the most valuable new source of power for many data centers may be the power they already have.

References

Anioke, C. (2025). How data centers affect the electrical grid: Implications for electricity demand, grid stability, and energy sustainability. *SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology, 17*(2), 61–66.

Anioke, C. (2026). Stranded power from data centers and possible uses. *Journal of Emerging Intelligence and Engineering Technologies, 3*(1), 41–49.

Asuri, R. (2019). Transforming legacy enterprise resource planning systems through SAP Fiori expansion and embedded analytics to enhance user experience and operational performance in SAP ECC. *SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology, 11*(2), 176–190. https://doi.org/10.18090/samriddhi.v11i02.13

Asuri, R. (2023). A strategic enterprise framework for integrating large language models with SAP Business Technology Platform to enable intelligent automation and operational excellence. *International Journal of Humanities and Information Technology, 5*(1), 87–103. https://doi.org/10.21590/ijhit.05.01.7

Barroso, L. A., Hölzle, U., & Ranganathan, P. (2019). *The datacenter as a computer: Designing warehouse-scale machines* (3rd ed.). Springer. https://doi.org/10.1007/978-3-031-01761-2

Fan, X., Weber, W.-D., & Barroso, L. A. (2007). Power provisioning for a warehouse-sized computer. In *Proceedings of the 34th Annual International Symposium on Computer Architecture* (pp. 13–23). Association for Computing Machinery. https://doi.org/10.1145/1250662.1250665

International Energy Agency. (2025). *Energy and AI*. IEA. https://www.iea.org/reports/energy-and-ai

Kim, K., Yang, F., Zavala, V. M., & Chien, A. A. (2017). Data centers as dispatchable loads to harness stranded power. *IEEE Transactions on Sustainable Energy, 8*(1), 208–218. https://doi.org/10.1109/TSTE.2016.2593607

Norris, T. H., Profeta, T., Patiño-Echeverri, D., & Cowie-Haskell, A. (2025). *Rethinking load growth: Assessing the potential for integration of large flexible loads in US power systems* (NI R 25-01). Nicholas Institute for Energy, Environment & Sustainability, Duke University. https://nicholasinstitute.duke.edu/publications/rethinking-load-growth

Radovanović, A., Koningstein, R., Schneider, I., Chen, B., Duarte, A., Roy, B., Xiao, D., Haridasan, M., Hung, P., Care, N., Talukdar, S., Mullen, E., Smith, K., Cottman, M., & Cirne, W. (2023). Carbon-aware computing for datacenters. *IEEE Transactions on Power Systems, 38*(2), 1270–1280. https://doi.org/10.1109/TPWRS.2022.3173250

Wahlroos, M., Pärssinen, M., Rinne, S., Syri, S., & Manner, J. (2018). Future views on waste heat utilization – Case of data centers in Northern Europe. *Renewable and Sustainable Energy Reviews, 82*(Part 2), 1749–1764. https://doi.org/10.1016/j.rser.2017.10.058