



Considerations for Distributed Edge Data Centers and Use of Building Loads to Support Large Interconnections

Aditie Garg, Chuck Booten, Otto VanGeet, Shanti Pless, Katelyn Stenger, Judith Vidal, Roderick Jackson, Luke Lavin, and Kumaraguru Prabakar

National Renewable Energy Laboratory

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated under Contract No. DE-AC36-08GO28308**

This report is available at no cost from
NREL at www.nrel.gov/publications.

Technical Report
NREL/TP-5500-96700
November 2025



Considerations for Distributed Edge Data Centers and Use of Building Loads to Support Large Interconnections

Aditie Garg, Chuck Booten, Otto VanGeet, Shanti Pless, Katelyn Stenger, Judith Vidal, Roderick Jackson, Luke Lavin, and Kumaraguru Prabakar

National Renewable Energy Laboratory

Suggested Citation

Garg, Aditie, Chuck Booten, Otto VanGeet, Shanti Pless, Katelyn Stenger, Judith Vidal, Roderick Jackson, Luke Lavin, and Kumaraguru Prabakar. *Considerations for Distributed Edge Data Centers and Use of Building Loads to Support Large Interconnections*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5500-96700. <https://www.nrel.gov/docs/fy26osti/96700.pdf>.

**NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated under Contract No. DE-AC36-08GO28308**

This report is available at no cost from
NREL at www.nrel.gov/publications.

Technical Report
NREL/TP-5500-96700
November 2025

15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored by NREL for the U.S. Department of Energy (DOE), operated under Contract No. DE-AC36-08GO28308. Funding provided internally by NREL. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

This report is available at no cost from NREL at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.osti.gov.

Cover photos (clockwise from left): Josh Bauer, NREL 61725; Visualization from the NREL Insight Center; Getty-181828180; Agata Bogucka, NREL 91683; Dennis Schroeder, NREL 51331; Werner Slocum, NREL 67842.

NREL prints on paper that contains recycled content.

Acknowledgment

The authors would like to thank Killian Mckena, David Sickinger, Chioke Harris, Xin Jin, Achilles Karagiozis, and John Farrell for their time in reviewing this report and for providing valuable feedback that strengthened its quality. The authors also gratefully acknowledge the support of the Building Technologies and Science Center in the development of this report.

Executive Summary

The rapid expansion of artificial intelligence (AI) and machine learning is driving unprecedented electricity demand from data centers. It is predicted that by 2030, 90% of AI workloads will be inference-based, requiring interconnection of multiple low-latency edge data centers (<20 MW) sited closer to end users—often on already constrained distribution feeders (Infrastructure Masons 2025). Although individually small, these loads can aggregate to large loads per feeder, straining infrastructure, creating multi-year interconnection delays, and driving up customer costs.

This paper proposes a data center-focused grid-integration framework that combines feeder hosting capacity analysis with building energy efficiency, building load flexibility, and waste heat reuse to expand effective feeder and substation headroom. Such approaches can reduce interconnection delays, lower costs for ratepayers, and accelerate AI-ready infrastructure deployment.

Table of Contents

Executive Summary	iv
1 Introduction	1
2 Understanding the Current Landscape	3
2.1 Research Questions	5
3 Industry Needs and Nationwide Grid Challenges	6
3.1 Developer Needs	7
3.2 Grid Needs and Impacts	7
3.3 Customer Needs and Impacts	7
4 Solution: Capacity Mapping Using Building Load Data and Waste Heat Reuse	9
4.1 Data Center-Focused Hosting Capacity Maps	9
4.2 Demand Reduction Strategy: Energy Efficiency and Responsive Loads.....	10
4.3 Demand Reduction Strategy: Data Center Heat Reuse	11
4.4 Example Case for Building Load Reduction, Grid Responsive Loads and Waste Heat Thermal Reuse	12
5 NREL Capabilities	16
References	17

List of Figures

Figure 1. 4-MW edge data center located in close to proximity to major research institutions and technology companies.....	2
Figure 2. Data center infrastructure in the United States, 2025	3
Figure 3. Inference and enterprise type data centers in a distribution network	5
Figure 4. Emerging edge data centers (includes current, under construction, and planned) using Baxtel data (https://baxtel.com).....	6
Figure 5. Illustrative example of a feeder hosting capacity map for a feeder within Massachusetts.....	9
Figure 6. An upgrade scenario for energy consumption reductions modeled for commercial buildings comparing baseline energy consumption with inclusion of building envelope upgrade and efficient electrical equipment in a building.....	10
Figure 7. An illustrative example for feeder peak load reduction with building flexibility.....	11
Figure 8. Building load management with waste heat reuse for data center in the vicinity of a neighborhood	12
Figure 9. Current grid capacity with interconnection queues of data centers	13
Figure 10. Updated grid capacity with building energy efficiency, flexibility, and waste heat reuse	14
Figure 11. Potential data center-focused hosting capacity map with building load solutions and waste heat recovery.....	16

List of Tables

Table 1. Three Types of AI-Related Data Centers.....	4
Table 2. Available Feeder and Substation Capacity	12
Table 3. Building Solutions for Increasing Grid Capacity.....	14

1 Introduction

The rapid growth of artificial intelligence (AI) and machine learning (ML) is driving a surge in electricity demand from data centers. Today, roughly 90% of AI workloads are training-based and only 10% are inference-based,¹ but by 2030 this ratio is expected to flip to 90% inference workloads (Infrastructure Masons 2025; Vincent 2025). Unlike training, inference requires low-latency processing close to end users, often making edge data centers (<20 megawatts [MW]) the preferred option.²

These edge data centers are increasingly being sited on distribution feeders, where available capacity is already constrained (EY-Parthenon 2025). Although individually small (often <5 MW), aggregated deployments can account for a large amount of load on a single feeder, creating unprecedented strain on local infrastructure. This rapid growth has major implications:

- **Infrastructure costs and delays:** Interconnection of new data center loads typically requires significant feeder and substation upgrades, which are expensive and can take 1–10 years to complete (Skidmore 2025). For example, in Virginia, data center interconnection delays of up to 7 years have been reported (Skidmore 2025).
- **Customer impacts:** Higher interconnection and upgrade costs are often passed through to ratepayers, and sudden load spikes from inference workloads can undermine grid reliability. For example, in Virginia, customers are already paying an additional \$276 per year on their electricity bills due to data center interconnection (Penn and Weise 2025).
- **Planning challenges:** Unlike hyperscale³ data centers that connect to the transmission grid, edge data centers must rely on the limited headroom of distribution feeders and substations, requiring new strategies to integrate these large, latency-sensitive loads.

Figure 1 illustrates a 4-MW edge data center that was recently deployed in Raleigh, North Carolina. This data center is located in close proximity to residential and commercial buildings (American Tower 2025).

¹ “Inference” is the model performing its desired task—using its training to make decisions/predictions.

² “Edge data centers” refer to relatively small data centers located closer to end users than traditional data centers defined further in Table 1.

³ Hyperscale data centers range in size from 100 MW to a few gigawatts (GW).

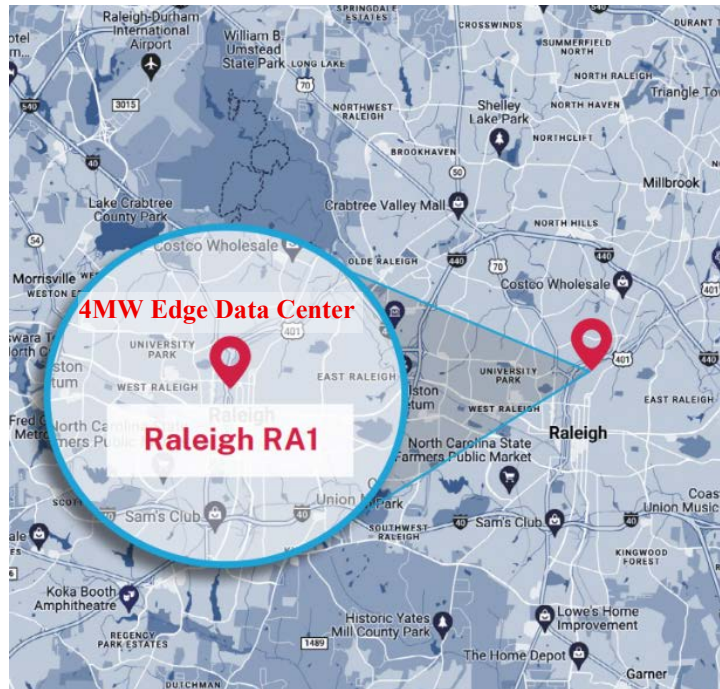


Figure 1. 4-MW edge data center located in close to proximity to major research institutions and technology companies

Figure from American Tower (2025)

When considering an edge data center like the one in Raleigh, latent flexibility exists within nearby residential, commercial, and industrial buildings that could help alleviate the costs and pressures noted above. Energy efficiency, demand-responsive building loads, and waste heat reuse networks represent untapped solutions to reduce peak demand and expand effective feeder capacity—potentially expediting interconnection and lowering customer costs.

Hosting capacity maps, widely used in interconnection studies, could be extended to include these building-side and thermal flexibility solutions, revealing new opportunities to site and integrate edge data centers with minimal costs and grid upgrades.

This white paper considers these opportunities and proposes a grid-integrated approach to siting, planning, and operating data centers up to 20 MW. By coupling feeder-level hosting capacity analysis with demand-side flexibility and waste heat reuse, utilities and developers can unlock additional headroom, reduce interconnection delays, and limit cost impacts to customers—all while enabling the rapid deployment of AI-ready infrastructure.

2 Understanding the Current Landscape

Lawrence Berkeley National Laboratory’s 2024 report on data center energy use shows that data center load growth has tripled over the past decade and is likely to double or triple again by 2028. The report found that 4.4% of total U.S. electricity was consumed by data centers in 2023, likely rising to 6.7%–12% by 2028. The report estimates that U.S. data center electricity demand may reach 74–132 GW by 2028 (Shehabi et al. 2024).

To get a sense of the nationwide infrastructure required for data centers, Figure 2 shows current U.S. data centers, transmission lines, and fiber optic lines. Note that currently operating data centers are shown as yellow circles, data centers being built are orange circles, and proposed data centers are white circles.

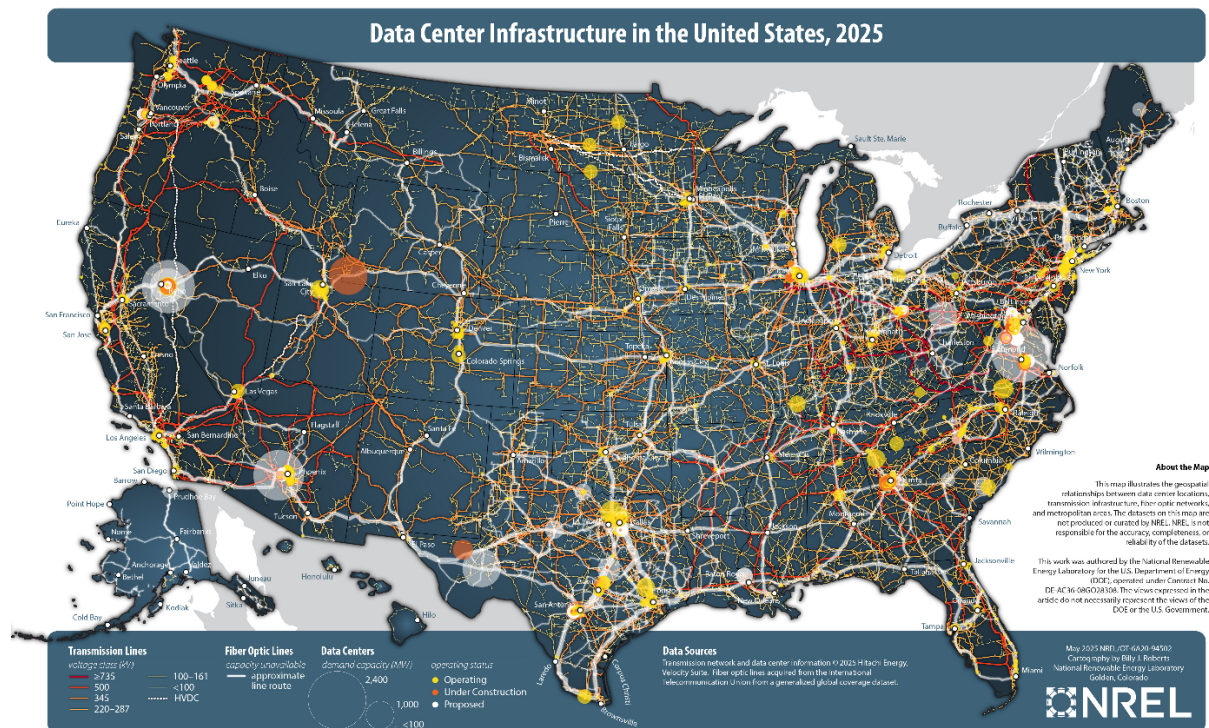


Figure 2. Data center infrastructure in the United States, 2025

Figure by Billy Roberts, NREL (2025)

As we consider the current and future landscape of data centers and how to approach powering them, Table 1 presents three broad categories of AI-related data centers: (1) hyperscale training centers, (2) enterprise/co-location centers, and (3) edge centers.⁴

⁴ DOE has prioritized research into temporal and spatial flexibility for AI-centric workloads (training and inference) through testbeds and grid-integrated operations (DOE 2024a; Shehabi et al. 2024).

Table 1. Three Types of AI-Related Data Centers

	Hyperscalers	Enterprise/Co-Location Data Centers	Edge (Inference Type, Telco and Commercial) Data Centers
Size	Large scale (>100 MW up to a few GWs)	Medium scale (5–100 MW)	Individually small (<5 MW)
Location	Centralized and spatially spread generally on <i>transmission grid</i>	Integrated into large facilities or campuses, typically intended for supporting the business or organization internally on a distribution substation	Highly localized on distribution feeders, with large aggregated demand
Latency Requirement	High latency	Low or medium latency	Low latency
Benefit Potential From Building Loads	Low: Massive adoption of energy efficiency programs across multiple feeders and substations can cumulatively free up transmission and generation capacity needed for large loads. This largely depends on the scale of adoption of programs that would benefit data centers training large AI models	Medium: Data centers that are connected within a distribution network (substation or feeder level) may benefit from the additional headroom resulting from building flexibility and waste heat reuse	High: These data centers are smaller, latency-sensitive facilities located within a distribution feeder closer to end users; they would benefit from building flexibility and waste heat reuse

Currently, about 90% of power demand growth in data centers is driven by AI model training. By 2030, however, the majority of workloads (nearly 90%)—are expected to shift toward inference that use AI trained models for inferencing (Infrastructure Masons 2025). This shift will fundamentally reshape digital infrastructure and accelerate data center deployment at distribution grid, enabling ultra-low-latency interactions between people and machines across communities worldwide. To interconnect and integrate these edge data centers (up to 20-MW Inference and Enterprise Data Centers) highlighted in Figure 3, the distribution network needs enough capacity and headroom available at the feeder and substation level. Considering the placement of such data centers within a distribution network opens opportunities for building loads to support data center interconnection ranging in size up to 20 MW.

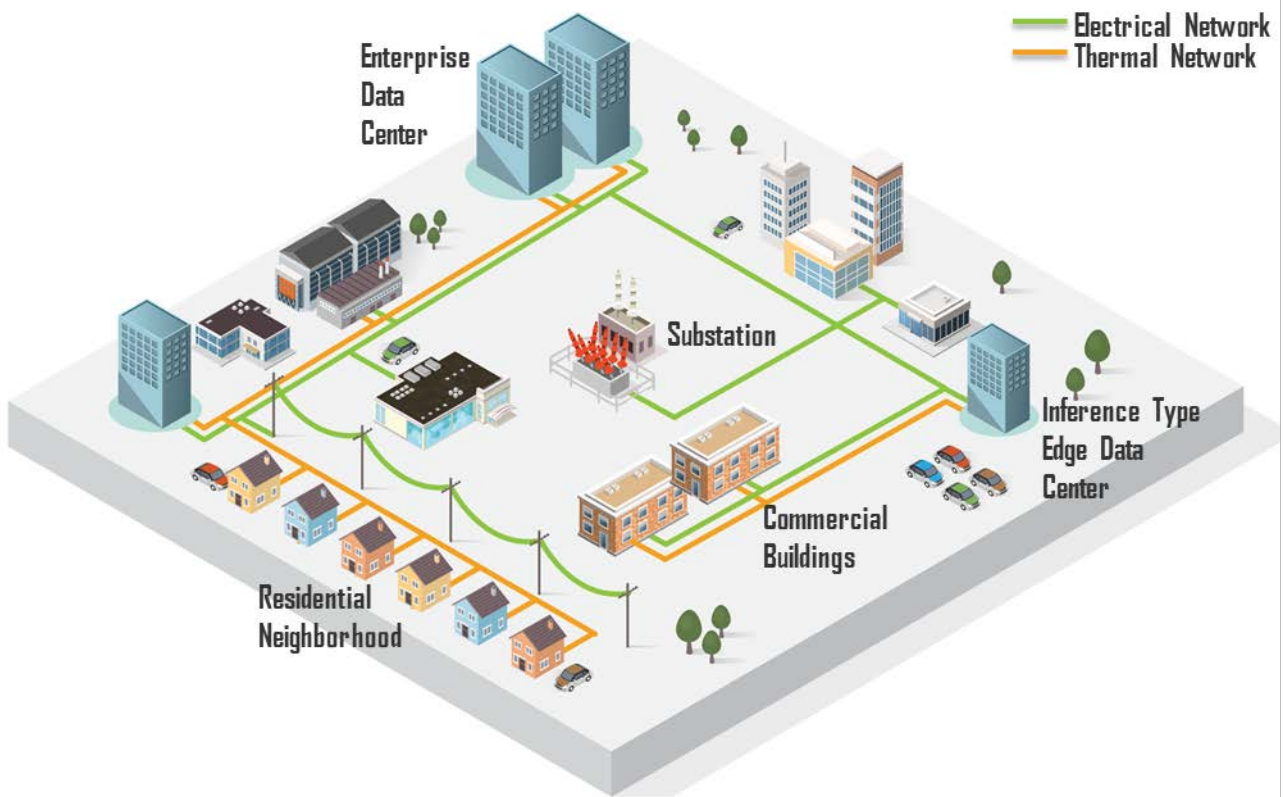


Figure 3. Inference and enterprise type data centers in a distribution network

Figure by Fred Zietz, NREL

2.1 Research Questions

This white paper addresses many of these factors for data center siting and decision-making. The goal is to support edge data center growth by answering the following questions:

- What is the impact of increasing data center interconnection on the grid, on customers, and on the needs of industry?
- What are the factors that affect interconnection of data centers?
- What are the concerns of a data center developer interconnecting new data centers at the edge of the grid?
- What are the proposed solutions to support data centers for an affordable and reliable power grid?

3 Industry Needs and Nationwide Grid Challenges

Nationwide, utilities are projecting substantial load growth due to AI and increased data center interconnection requests (Wilson et al. 2024). Using data from Baxtel, Figure 4 shows smaller-scale data centers emerging across the country, with interconnection sizes varying by region.

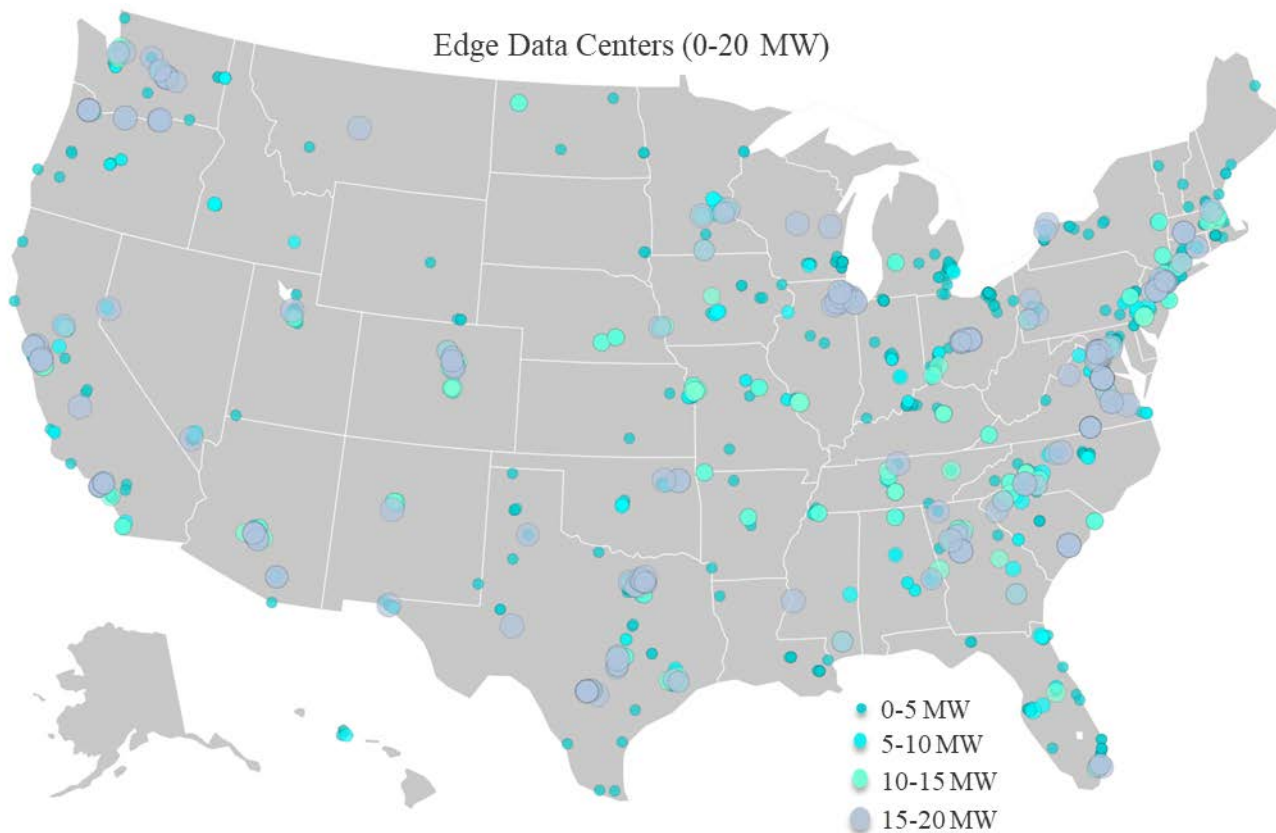


Figure 4. Emerging edge data centers (includes current, under construction, and planned) using Baxtel data (<https://baxtel.com>)

Edge data center growth, while economically promising, presents a critical risk: Without tools to manage and strategically integrate new loads, utilities face mounting interconnection delays, grid upgrade requirements, increasing customer costs, and potential reliability shortfalls. Moreover, there are multiple factors that play an important role in siting edge data centers. These factors include, but are not limited to:

- **Time to power the data center:** Speed of interconnection and ability to deliver power on short timelines.
- **Land use/co-location possibilities:** These centers are proliferating rapidly and are increasingly being sited in commercial and suburban zones—where available grid hosting capacity is often limited.
- **Connectivity network structure:** Access to high-speed fiber optics to support low-latency requirements.

- **Proximity:** Geographic proximity to population and load centers for reduced latency.
- **Water availability:** Access to water for cooling, where required.
- **Available grid capacity:** Availability of substation and feeder headroom to support large, concentrated loads.
- **Available demand capacity:** Ability to manage sudden or variable high-load conditions without compromising grid reliability.
- **State and local taxes:** Taxes significantly affect the total cost of owning and operating a data center, including property, income, and utility-related taxes. Higher taxes can make a location less economically attractive, influencing decisions on where to site and interconnect a data center (Magnum et al. 2020).

Development and integration of data centers with the grid requires a three-way alignment between the developer, grid, and customer, ensuring each party's requirements and needs are met. The following subsections describe the needs of the developer, grid, and customer.

3.1 Developer Needs

Interconnection timelines: Long interconnection queues reduce profitability and delay the ability to deliver services, creating risks for both developers and their customers, hence, affecting the productivity and losses to manufacturing industry.

Energy efficiency: data centers are looking for efficient cooling solutions for server, and racks. An efficient solution would help developers reduce total peak demand requirement, lowering interconnection requirements and potentially accelerating approvals (Kozubal et al. 2017, Geet et al 2024)

3.2 Grid Needs and Impacts

Limited grid infrastructure: Even relatively small-scale data centers (up to 20 MW) can create significant capacity expansion challenges due to limited grid capacity in the distribution network. Utilities have reported feeder-level clusters reaching 30 MW on distribution systems, exceeding infrastructure capacity (EY-Parthenon 2025).

Affordability: Edge data centers are often sited near population centers, where limited grid capacity forces costly upgrades that can cascade into higher rates for all customers.

Interconnection process: Unlike larger data centers (greater than 100 MW) with access to transmission networks, edge and enterprise data centers (up to 20 MW) depend on constrained distribution feeders. This mismatch leads to long interconnection timelines and expensive reinforcements.

3.3 Customer Needs and Impacts

Affordability: Grid upgrades to support new data center load increase capital costs, which are often socialized through higher electricity rates. This could create affordability issues for residential, commercial, and industrial customers at the edge of the grid.

Reliability: AI-driven inference workloads can drive sudden and unpredictable spikes in demand. These fluctuations stress local feeders, reducing service reliability and potentially affecting customer comfort and productivity.

Without coordinated planning, the growth of data centers could stress already constrained grids, undermining reliability and affordability. Thus, to enable integration of edge type data centers, utilities could review solutions that enhance grid capacity by utilizing nearby building loads for load reduction strategies. The following sections introduce the concept and illustrate an example with the load reduction strategies.

4 Solution: Capacity Mapping Using Building Load Data and Waste Heat Reuse

4.1 Data Center-Focused Hosting Capacity Maps

To enable rapid yet responsible deployment of edge data centers, utilities and developers need access to data center–focused hosting capacity maps. These maps should build on existing feeder hosting capacity tools but go further—by integrating data center-specific requirements such as thermal limits and workload type, and overlaying energy efficiency and load reduction potential from surrounding building loads. This will enable planners to assess not just where there is available headroom, but also where active load coordination (such as pre-cooling buildings or using flexible loads such as water heaters to shift peak load; see Section 4.2) can create additional capacity to host data centers without traditional infrastructure upgrades. Further, by identifying neighborhoods with high potential for data center thermal heat energy reuse (e.g., commercial corridors or multifamily residential zones, see Section 4.3), planners can co-optimize electrical and thermal integration for enhanced efficiency and resilience.

To illustrate this approach, Figure 5 shows the hosting capacity for a region within Massachusetts. It highlights a specific section of the feeder that is currently heavily constrained, but where possible opportunities exist for building loads to create additional headroom capacity. As noted previously, although each edge data center draws modest power, clusters in urban neighborhoods can overload feeders. Utilities are reporting aggregate peak loads at the feeder, which often trigger expensive upgrades and long interconnection delays (O’Neil 2025; Wang Efram Nora et al. 2024).

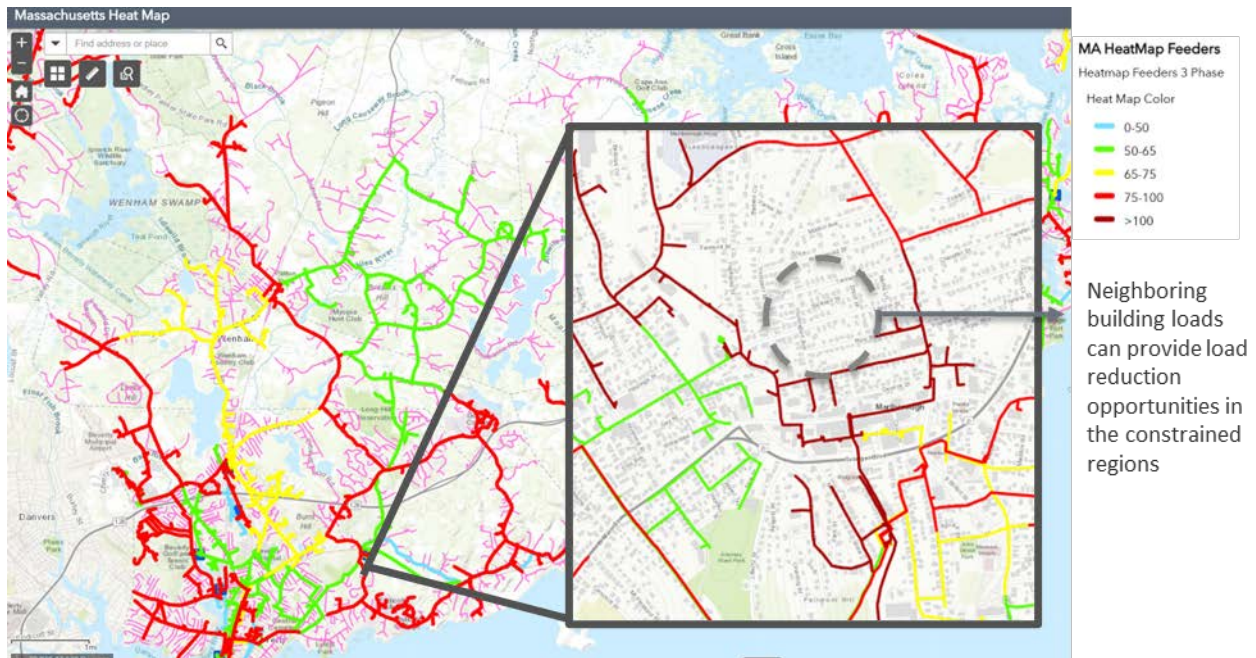


Figure 5. Illustrative example of a feeder hosting capacity map for a feeder within Massachusetts

Map from National Grid (<https://systemdataportal.nationalgrid.com/MA/>)

4.2 Demand Reduction Strategy: Energy Efficiency and Responsive Loads

Targeted building energy efficiency and demand flexibility strategies can significantly reduce peak feeder load, enhancing grid capacity headroom. Figure 6 illustrates an upgrade scenario for commercial buildings that could reduce energy consumption of a building by approximately 10% by incorporating building envelope measures. The energy consumption can be reduced further by demand flexibility and building peak load reduction measures to around 18%–20%.⁵

Incorporating energy efficiency with improved controls and retrofits typical residential and commercial buildings can reduce 15%–30% of their peak electric load, as highlighted in Figure 7 (DOE 2024b).

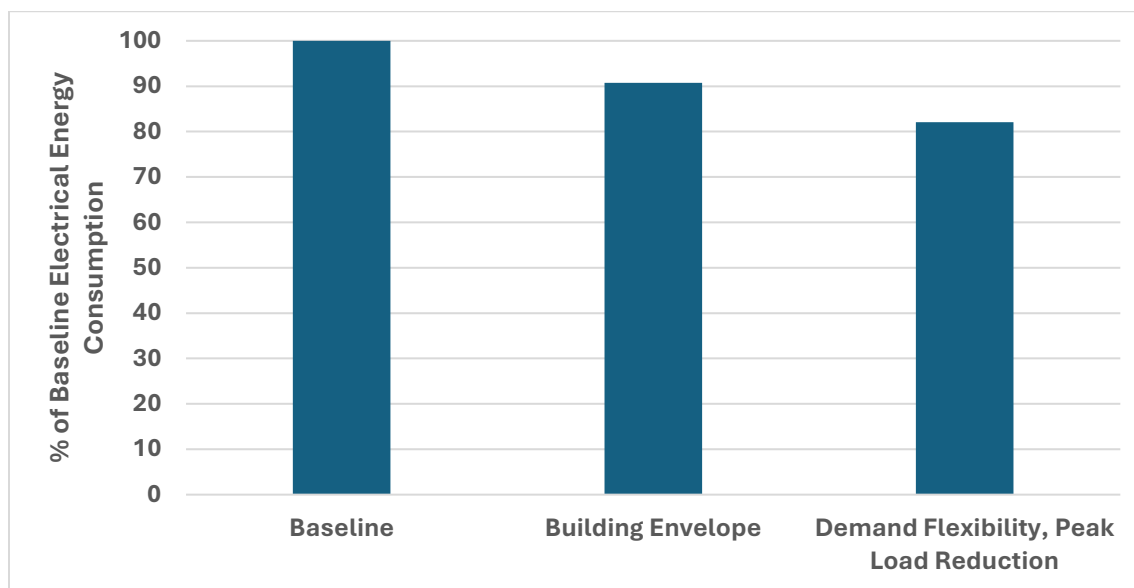


Figure 6. An upgrade scenario for energy consumption reductions modeled for commercial buildings with inclusion of building envelope upgrade and demand flexibility measures in a building.

⁵ This analysis was carried out using NREL’s ComStock™ (<https://comstock.nrel.gov/>) database for the complete United States, considering buildings with only electric consumption (gas and other fuel types of buildings were not considered). The building envelope scenario included wall insulation upgrades, roof insulation, and windows to align with target values per climate zone. The demand flexibility and peak load reduction scenario included combined lighting and thermostat setbacks with demand flexibility measures.

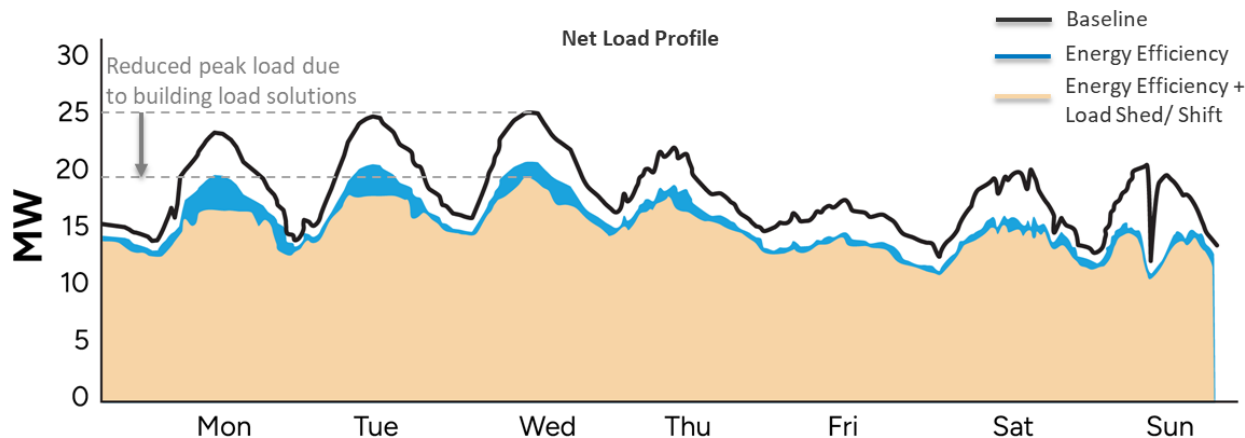


Figure 7. An illustrative example for feeder peak load reduction with building flexibility

Figure by Fred Zietz, NREL

When deployed at scale across neighborhoods, aggregated reductions can meaningfully increase distribution feeder capacity, enabling the hosting of inference data centers with less strain on the grid. Moreover, load shifting strategies that utilize flexible building resources, such as HVAC, water heaters, and energy storage (both electric and thermal), can also help shift the peak load to non-peak hours based on grid conditions. Coordinated dispatch of flexible resources, demand reduction operations, and building-side load control may achieve a higher peak load reduction for a feeder. An example case is illustrated in Section 4.4.

4.3 Demand Reduction Strategy: Data Center Heat Reuse

Edge data centers represent an opportunity to enhance energy efficiency and reduce local electric demand through heat reuse due to their proximity to end users. By capturing and repurposing low-grade thermal energy for neighboring building heating and domestic hot water needs, data centers can reduce stress from the electric load required for space and water heating. Examples include distributed storage (electric or thermal) and hybrid HVAC systems that can utilize multiple forms of energy (e.g., data center heat, fossil fuels, geothermal or air-source heat pumps) and potentially use them for reducing load while simultaneously maximizing cost-effectiveness of buildings. Load reduction efforts in buildings could also include redesigning buildings with retrofits and/or new construction to store data center heat including high thermal mass buildings using advanced concrete formulations and/or rapid high thermal mass construction techniques (prefabricated concrete walls and foundations) and/or designing contained hot aisle that allows reuse of waste heat in data centers (Sheppy, M. et al. 2011).




Figure 8. Building load management with waste heat reuse for data center in the vicinity of a neighborhood

Figure by Fred Zietz, NREL

As highlighted in Figure 8, by leveraging these efficiency approaches combined with thermal use potential and responsive building loads, utilities and planners can widen distribution feeder headroom without costly infrastructure upgrades. Such solutions would enhance the hosting capacity capability of distribution networks, thus enabling edge data centers less than 20 MW sited in locations with high aggregated demand while still maintaining safety margins for the grid and data centers as well.

4.4 Example Case for Building Load Reduction, Grid Responsive Loads and Waste Heat Thermal Reuse

Consider a distribution network where a feeder rated at 40 MW has 8 MW of available headroom, yet it is facing an interconnection request from an inference data center of 15 MW. At the substation level, a transformer rated at 100 MW has currently 15 MW of available headroom capacity, yet is facing an interconnection request from an enterprise data center of 30 MW. Table 2 summarizes the existing capacity and interconnection requests, while Figure 9 illustrates the current grid availability.

Table 2. Available Feeder and Substation Capacity

	Feeder	Substation
Rated Capacity	40 MW	100 MW
Available Headroom	8 MW	15 MW
Data Center Interconnection Queue	Inference: 15 MW	Enterprise: 30 MW

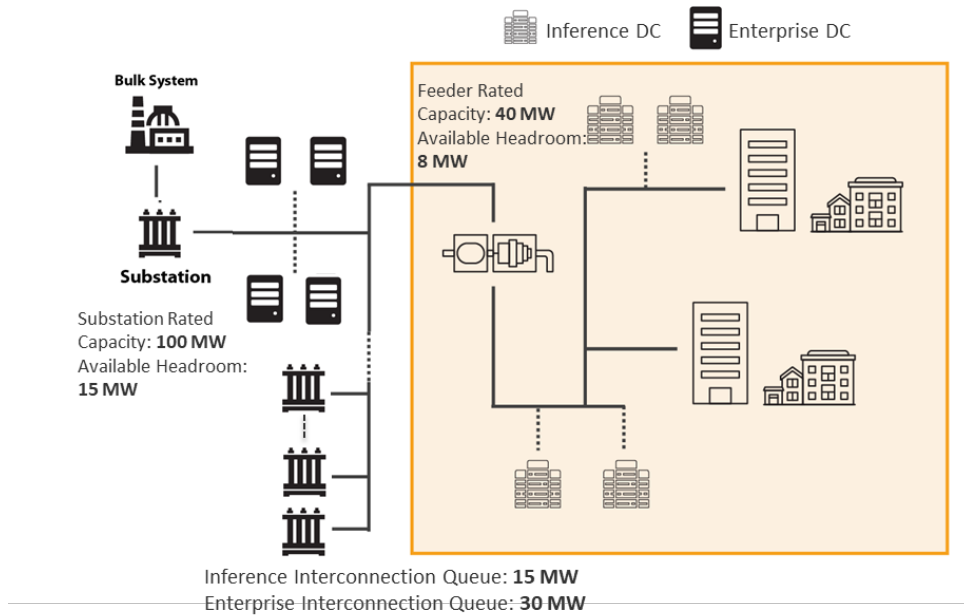


Figure 9. Current grid capacity with interconnection queues of data centers

Given both the feeder and substation do not have enough grid capacity to interconnect the data center, the building load solutions described in the previous section—energy efficiency, load flexibility, and waste heat reuse⁶—can be leveraged to create additional headroom. As shown in Table 3, implementing these solutions increases feeder capacity by 10 MW (totaling the new headroom capacity of 18 MW), enabling interconnection of the requested 15 MW inference data center. At the substation level, aggregated building solutions across feeders add 15 MW of capacity, raising total available headroom to 35 MW, sufficient to interconnect the 30 MW enterprise data center (Figure 10).

⁶ The proposed waste heat reuse is proposed to be utilized by nearby residential and industrial buildings. It must be noted that although the waste heat from the data center would be used for heating these buildings, some heat reuse applications have seasonality and intraday variations that would need to be considered in more details to understand its full utilization.

Table 3. Building Solutions for Increasing Grid Capacity

Increased Capacity Factors	Feeder	Substation
Building Energy Efficiency	2 MW	5 MW
Building Load Flexibility	4 MW	5 MW
Waste Heat Reuse	4 MW	5 MW
Total Increased Headroom Capacity at Feeder	8 MW + 10 MW = 18 MW	15 MW + 15 MW = 30 MW

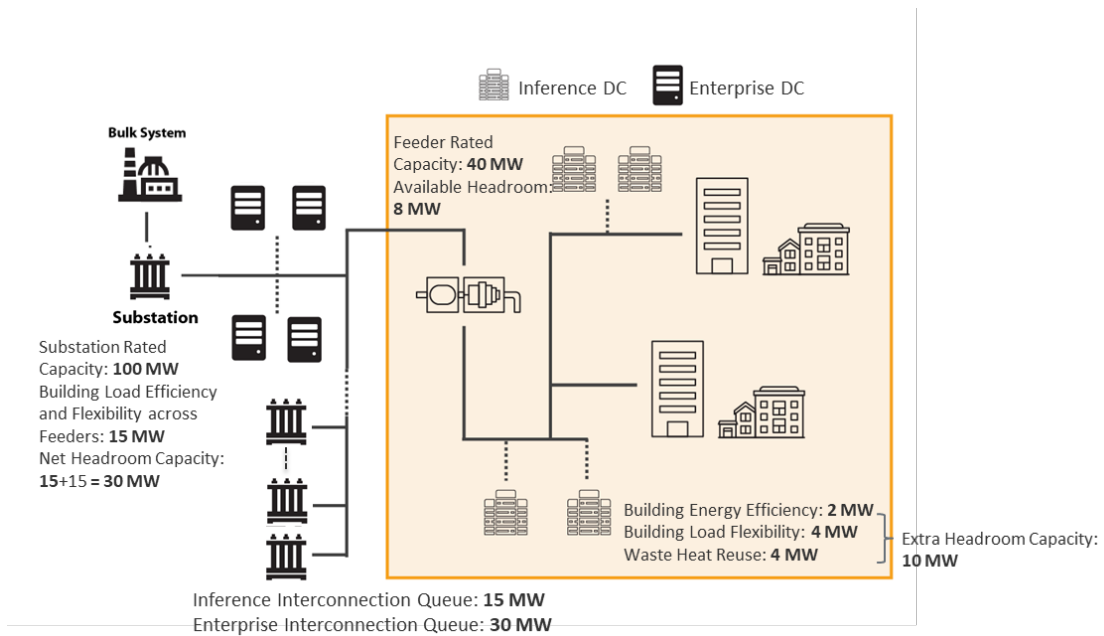


Figure 10. Updated grid capacity with building energy efficiency, flexibility, and waste heat reuse

Thus, adopting building energy efficiency, load flexibility, and waste heat reuse solutions across feeders within the distribution network increases the available headroom capacity at the feeder and substation. This approach allows the grid to accommodate these large loads while also avoiding infrastructure upgrades that would be costly and time consuming. This process satisfies the needs of major stakeholders as follows:

- Developer: Faster interconnection times, reduced costs
- Utility: Avoided infrastructure upgrade costs and times
- Customer: Reliable power with affordable data center interconnection solution.

Thus, these solutions are productive in not only increasing the capacity at a feeder end but also increasing substation-level capacity to interconnect larger data centers. These holistic approaches are expected to provide two-way solutions to reduce net electric demand on constrained feeders and also contribute to broader reliability and energy affordability goals.

5 NREL Capabilities

To ensure that the rise of distributed edge data centers supports rather than undermines the power grid, we must adopt a forward-looking, integrated planning approach. As mentioned in the previous section, this includes developing **data center-specific hosting capacity maps** that combine feeder-level capacity data with localized building energy efficiency, grid responsive loads and thermal reuse opportunities. Figure 11 illustrates a potential data center hosting capacity map that could be overlaid with the projected data center and highlights locations with available capacity due to building load flexibility and waste heat reuse capability. Such solutions can help utilities, planners, and developers identify zones where edge data centers can be strategically deployed without triggering major infrastructure upgrades.

Utilizing the potential solution map for edge data centers for a specific region NREL can enable investment-quality decisions through its rigorous validation processes and utilizing advanced integrated modeling platforms. Modeling frameworks such as **ComStock™** and **ResStock™**⁷ can further help pinpoint optimal siting zones where electricity and heat demands can be synergistically managed.⁷ Additionally, platforms like NREL's **ARIES**⁸ testbed can be used to simulate real-time coordination between edge data center loads and neighborhood-level flexible loads. These approaches ensure that AI innovation and economic development can proceed while preserving grid reliability and keeping energy affordable.

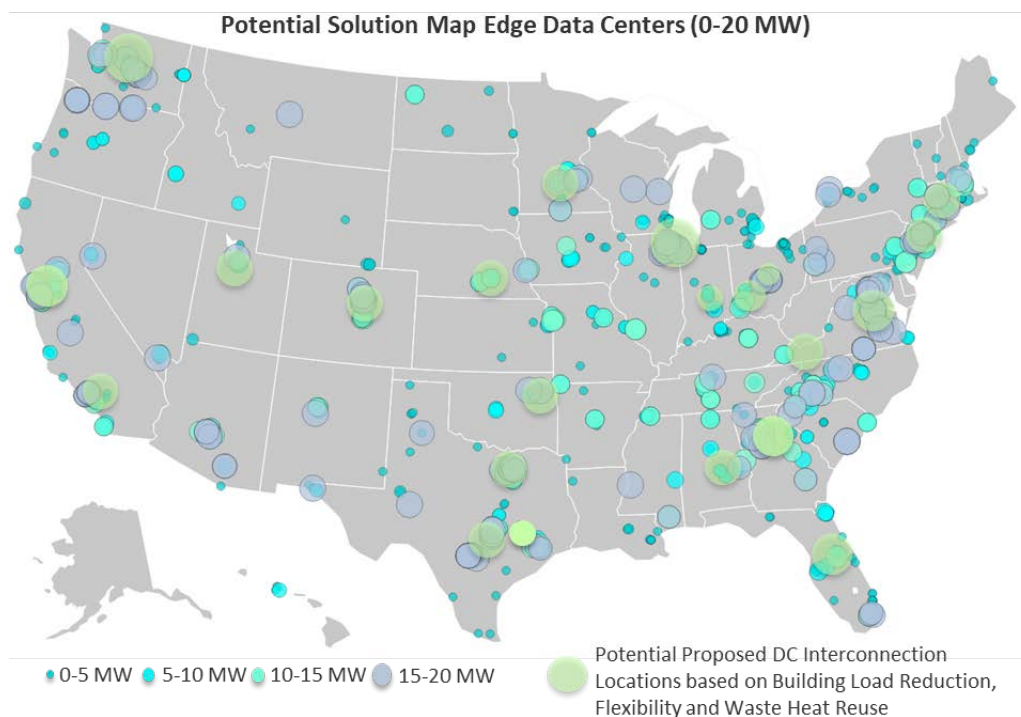


Figure 11. Potential data center-focused hosting capacity map with building load solutions and waste heat recovery

⁷ <https://comstock.nrel.gov/>, <https://resstock.nrel.gov/>

⁸ Advanced Research on Integrated Energy Systems, <https://www.nrel.gov/aries>

References

- American Tower. 2025. “Raleigh Data Center: American Tower Edge Data Centers.” https://pardot.americantower.com/l/25692/2024-04-26/7sykh4/25692/1714158551dlzDuWB0/atc_solutions_aggregation_edge_data_center_raleigh_nc.pdf.
- DOE. 2024a. Recommendations on Powering Artificial Intelligence and Data Center Infrastructure. U.S. Department of Energy. <https://www.energy.gov/sites/default/files/2024-08/Powering%20AI%20and%20Data%20Center%20Infrastructure%20Recommendations%20July%202024.pdf>.
- DOE. 2024b. *Decarbonizing the U.S. Economy by 2050: A National Blueprint for the Buildings Sector*. U.S. Department of Energy. <https://www.energy.gov/sites/default/files/2024-12/bto-decarbonizing-us-economy-2050-122724.pdf>.
- EY-Parthenon. 2025. “The new data center collaboration for utilities and developers.” <https://www.ey.com/content/dam/ey-unified-site/ey-com/en-us/campaigns/data-centers/documents/ey-energy-data-center-report-vf.pdf>.
- Geet, Otto Van, and Sickinger, David. 2024. “Best Practices Guide for Energy-Efficient Data Center Design.” Federal Energy Management Program. https://www.energy.gov/sites/default/files/2024-07/best-practice-guide-data-center-design_0.pdf
- Infrastructure Masons. 2025. “State of the Digital Infrastructure Industry: Annual Report 2025.” <https://imasons.org/wp-content/uploads/2025/05/State-of-the-Infrastructure-Industry-2025.pdf>.
- Kozubal, Eric J. "Innovation Incubator: LiquidCool Solutions Technical Evaluation. Laboratory Study and Demonstration Results of a Directed-Flow, Liquid Submerged Server for High-Efficiency Data Centers." Dec. 2017. <https://doi.org/10.2172/1413900>
- Magnum, A. Fletcher, David Zorn, Martina Arel, and Kate Driebe. 2020. “The Impact of Data Centers on the State and Local Economies of Virginia.” Prepared by Magnum Economics for the Northern Virginia Technology Council. https://biz.loudoun.gov/wp-content/uploads/2020/02/Data_Center_Report_2020.pdf.
- Moe, Allison, and Patrick Gibbs. 2023. “Equitable Electrification Analysis for Existing Buildings in Richmond, CA.” National Renewable Energy Laboratory for the U.S. Department of Energy’s Communities LEAP program. <https://docs.nrel.gov/docs/fy23osti/86954.pdf>.
- National Grid. 2025. “Massachusetts System Data Portal.” <https://systemdataportal.nationalgrid.com/MA/>.
- O’Neil, Connor. 2025. “Powering Data: NREL Partner Forum Puts Everything on the Table.” National Renewable Energy Laboratory. News story, June 11, 2025.

<https://www.nrel.gov/news/detail/features/2025/powering-data-nrel-partner-forum-puts-everything-on-the-table>.

Penn, Ivan, and Karen Weise. 2025. “Big Tech’s A.I. Data Centers Are Driving Up Electricity Bills for Everyone.” Aug. 14, 2025. *New York Times*.

<https://www.nytimes.com/2025/08/14/business/energy-environment/ai-data-centers-electricity-costs.html>.

PJM Inside Lines. 2025. “2025 Long-Term Load Forecast Report Predicts Significant Increase in Electricity Demand.” <https://insidelines.pjm.com/2025-long-term-load-forecast-report-predicts-significant-increase-in-electricity-demand/>.

Reyna, Janet, et al. 2025. ResStock Technical Reference Documentation v3.3.0. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5500-91621.

<https://www.nrel.gov/docs/fy25osti/91621.pdf>.

Roberts, Billy. 2025. “Data Center Infrastructure in the United States, 2025 (Map).” Golden, CO: National Renewable Energy Laboratory. <https://www.nrel.gov/docs/gen/fy25/94502.jpg>.

Shehabi, A., Smith, S.J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M.A.B., Holecek, B., Koomey, J., Masanet, E., Sartor, D. 2024. *2024 United States Data Center Energy Usage Report*. Lawrence Berkeley National Laboratory, Berkeley, California. LBNL-2001637.

<https://doi.org/10.71468/P1WC7Q>.

Sheppy, M., Van Geet, O., Pless, S., Donovan, K., Powers, C., & Lobato, C. (2011). *Reducing Data Center Loads for a Large-Scale, Low-Energy Office Building: NREL's Research Support Facility (Book)*: NREL (National Renewable Energy Laboratory).

<https://doi.org/10.2172/1031393>

Skidmore, Zachary. 2025. “US data centers face grid bottlenecks as regional operators delay upgrades.” *Data Center Dynamics*. <https://www.datacenterdynamics.com/en/news/us-data-centers-face-grid-bottlenecks-as-regional-operators-delay-upgrades/>.

Vincent, Matt. 2025. “Inside NVIDIA’s Vision for AI Factories: Wade Vinson’s Data Center World 2025 Keynote.” *Data Center Frontier*. <https://www.datacenterfrontier.com/machine-learning/article/55286658/inside-nvidias-vision-for-ai-factories-wade-vinsons-data-center-world-2025-keynote>.

Wang Efram, Nora, and Neal Elliott. 2024. “Turning Data Centers into Grid and Regional Assets: Considerations and Recommendations for the Federal Government, State Policymakers, and Utility Regulators.” ACEEE Policy Brief.

<https://www.aceee.org/sites/default/files/pdfs/Turning%20Data%20Centers%20into%20Grid%20and%20Regional%20Assets%20-%20Considerations%20and%20Recommendations%20for%20the%20Federal%20Government%20-%20State%20Policymakers%20-%20and%20Utility%20Regulators.pdf>.

Wilson, John, Zach Zimmerman, and Rob Gramlich. 2024. “Strategic Industries Surging: Driving US Power Demand.” Grid Strategies, LLC. <https://gridstrategiesllc.com/wp-content/uploads/National-Load-Growth-Report-2024.pdf>.