Ceres

An Investor Guide to Addressing Electric Power Risk from Data Center Growth

August 2026

Table of Contents

1. Introduction

Data centers are moving forward at extraordinary speed. In so doing, they are urgently trying to build and procure power amid rising electricity demand and surging prices. This pace has clear implications and has fueled a national pushback, with local communities backing bills in almost every state to halt or slow data center construction. Ceres is neither supporting or opposing these data centers, but as data centers proliferate in the U.S., it is crucial that new facilities are built in ways that are fair, affordable, efficient, and clean, while protecting the electric grid, local watersheds, community interests, residents, and local businesses and involving communities in the proposed build out. This report identifies challenges for electricity companies created by data center growth, proposes solutions, and provides a guide investors can use to engage with electricity companies.

A related resource, **An Investor Guide to Addressing Water Risk in the High-Tech Sector Value Chain** (available to Valuing Water Finance Initiative signatories) focuses on the significant water dependencies stemming from data center operations—including indirect water use tied to electricity generation—and the semiconductor supply chain. Additional resources, including Ceres' **Sustainable Solutions for Data Centers** and research on data centers' water impacts, can be found on **Ceres' website**.

Data centers range from small edge facilities serving localized needs to hyperscale campuses supporting cloud computing and AI training. Hyperscale data centers dominate capacity growth and account for the largest loads, often measuring tens to hundreds of megawatts per site, with emerging campuses targeting gigawatt-scale demand. Cloud service providers such as Amazon Web Services, Microsoft Azure, and Google Cloud lead the hyperscale market. Meta and Apple are two other significant power users in the U.S. Enterprise, colocation, edge, and containerized data centers represent smaller load profiles but collectively contribute to growing system complexity.

Data centers currently account for 4% to 5% of U.S. power demand and load is expected to rapidly increase over the next decade. Utilities are struggling to accurately forecast additional power needs because there are many factors that can impact the actual power required. For example, a data center's actual power draw is often lower than its nameplate capacity, particularly in early years while it is ramping up operations. This is significant because it is the nameplate capacity that informs utility planning. In addition, not all data centers in the utility interconnection queue will come to fruition. Despite these uncertainties, the U.S. electricity system is entering a phase of unprecedented growth. Given the magnitude of power requirements, electricity companies shoulder much of the risk of the AI boom, and they need to enact the available solutions to mitigate that risk.

2. Risks to Electricity Companies

An electricity company's business model plays a major role in shaping its risk profile, whether it is a regulated electric utility (integrated or wires-only), or whether it participates in competitive markets through merchant generation, contract generation, wholesaling, retailing, or any combination of these activities. Although some utility sector companies also engage in gas distribution or retailing, our focus is on electricity-related risks.

2.1 Affordability and Community Opposition

As data centers become more prevalent around the country, community impacts (land use, pollution, water stress, electricity costs) are becoming more obvious. Affordability is top of mind, and the potential for data centers to contribute to higher electricity costs for all other customers is a serious concern. Typically, large electric customers pay for the upgrades needed to interconnect to the distribution system. On the other hand, transmission grid and generation buildout costs are socialized across all customers. Note that in some cases, the additional revenue from data center electricity use offsets the costs of upgrades, so the impact of data centers on electricity costs must be analyzed site by site.

Consequently, community opposition to data center development is growing. Data Center Watch estimates 48 data center projects totaling over $156 billion were stalled or abandoned over the course of 2025 as opposition hardened into zoning, permitting, and legislative barriers. Additionally, opposition is increasingly entering the political arena. In March 2026, Senators Bernie Sanders and Alexandria Ocasio Cortez proposed legislation to pause all new data center construction nationwide until AI safeguards are in place, citing rising electricity costs, pollution, and community harm.

Moratoriums pose a relevant regulatory risk for utilities planning to supply power to data centers, since political intervention could slow or cap load growth beyond what utilities expect. Committing to building infrastructure ahead of signed non-revocable contracts would be an especially risky strategy for a utility company since the costs of unnecessary assets would be borne by customers rather than the company and shareholders. This could eventually harm shareholders, since commissioners consider customer dissatisfaction when determining future approvals related to capex, tariffs, cost recovery, and return on equity.

2.2 Electric Infrastructure Delays

Time to power has become one of the main pressures driving data center development, as delays in securing electricity can materially affect data center operators' revenue, competitive positioning, and ability to justify the high valuations embedded in today's AI-driven markets. Many data center developers are considering alternatives to grid power if utilities cannot provide power fast enough. It's critical that electricity companies reduce the interconnection timeline and minimize delays in generation or grid upgrades to avoid data centers building their own power, which is likely to include on-site fossil generation. Besides the impact on the community, utilities would lose that resource for the rest of the grid and would lose the investment opportunity.

Furthermore, if electricity company projects have delays, the longer project timelines mean potentially missing the shareholder value created by either the commercial bilateral contract, wholesale markets, or the regulated asset base opportunity. When timelines slip because of unexpected delays in project execution, there can be direct financial consequences and broader reputational damage that can harm the ability to win future deals.

The value creation at stake can stretch into billions of dollars for a utility building new transmission lines. While costs vary significantly by location and capacity, major lines can be millions of dollars per mile. An example closely tied to data center growth is AEP, Dominion, and FirstEnergy's joint 765kV 415-mile project in Virginia, West Virginia, and Maryland, which has an updated cost estimate of $1.5 billion. For integrated utilities, building baseload capacity represents an investment opportunity of about $2.8 billion per gigawatt. Furthermore, building renewables and energy storage is cleaner and often faster to build with greater value added since fuel costs are converted to capital costs.

2.3 Data Center Utilization

There will likely be economic pressure to make data centers more efficient as competition and electricity-related operating costs increase. Efficiency gains across chips, cooling systems, models, and inference workloads can all lower power use below expectations, which presents system and asset-specific risk due to uncertainty in recovering the costs of electric infrastructure built to serve data centers.

Over time, newer, more efficient data centers may displace older ones. Overbuilding power generation for specific data centers could result in cost recovery challenges due to underutilization.

Electricity companies with a responsible capital allocation strategy should consider more than just asset lifetime in assessing the cost recovery risk of serving a data center. Assuming ongoing baseload operation throughout the lifetime is a financially dangerous assumption because it ignores ongoing improvements in data center efficiency.

2.4 Counterparty Credit

Through long-term contracts, especially when dedicated infrastructure is involved, electricity companies are exposed to the credit risk of their customers. While hyperscalers have investment grade credit ratings, which provides some confidence, there is still risk for electricity suppliers because:

- Investment grade hyperscalers are not the entire data center market
- Contracts signed by a subsidiary or special purpose vehicle do not automatically become obligations of the parent or wider group without special provisions, so the credit risk of the entity signing the contract may be significantly different than the parent
- Credit quality can change over time
- General unsecured claims have considerably higher risk

Electricity companies need credit assessment processes, standards, and contractual protections against counterparty credit risk, and investors need more information to ascertain that their companies are managing credit risk exposure well.

2.5 Water

Around half of data centers in the United States today are fully or partially powered by power plants located in water stressed regions. Ceres' research shows that annual indirect water use from a rapidly growing data center cluster in the Phoenix region (Arizona) may increase by almost 400%, emphasizing the significance of embedded water use in energy procurement.

The water use embedded in purchased electricity depends on the mix of energy sources supplying the grid, the design of generation units, and production of fuels. For example, renewable energy sources relying on wind and solar have a lower indirect water footprint, while coal, natural gas, and nuclear power generation require significant water for cooling and steam production. In addition to the volume of water used, the process used to produce the fuels can also impact water resources. For example, water contaminated in the production of natural gas used for electricity generation is an issue in water constrained regions such as the Permian Basin in the Southwestern U.S.

Electricity companies must consider the availability of water and the impact that data center demand will have on their operations, including curtailment and shutdown. The National Renewable Energy Laboratory has documented historical incidents where water resources have affected power plant operations.

2.6 Carbon Reduction Goals

Many electricity companies have established carbon reduction targets as a key mechanism for managing climate-related risks and opportunities. If data center load growth results in higher utilization of older, less efficient gas and coal generation plants, the electricity companies may not be able to meet their goals. While this may not have a short-term direct impact on electricity company profitability, it may have indirect impacts due to reputational damage, reduced good will from regulators and customers, and reduced investor confidence. An example of this is Southern Company's utilities citing data center driven demand as a reason to extend the life of coal fired generation previously slated for retirement, raising concerns about the company's ability to meet its net-zero commitments and increasing environmental and public health risks.

It should be noted that data center growth can also undermine the commitments of technology companies to address the risks of climate change, creating reputational and investor risks that extend beyond the power sector. For example, a single off grid, gas-powered data center campus in West Virginia supplying compute capacity to Microsoft could increase the company's annual emissions by 44%, despite Microsoft's public goal of becoming carbon negative by 2030. At full buildout, the 8 GW facility would emit the equivalent of nearly six million cars' worth of CO_2 annually, highlighting how speed-to-power decisions, particularly those relying on behind the meter fossil generation, can directly conflict with stated targets for reducing climate risk even among the tech companies leading on managing these impacts.

2.7 Grid Reliability and Resiliency

Data centers create unique reliability challenges since they operate continuously, scale rapidly, and cluster in ways that create concentrated stress on local and regional grid infrastructure. Serving hyperscale demand requires significant new generation capacity and coordinated transmission and distribution upgrades to physically deliver power to load centers. When generation additions and grid buildout lag power demand, the system operates with smaller reserve margins and more frequent grid congestion. The North American Electric Reliability Corp., the nonprofit regulatory authority, has raised the alarm around load growth creating reliability issues across several regions in its Long-Term Reliability Assessment.

Grid operators have documented a growing number of close calls over the past five years, with sudden, large shifts in electricity use requiring emergency operator intervention to prevent outages. These issues reflect a grid that is operating with a lower margin for error than it was designed to accommodate. In addition, the power draw from data centers is inherently variable. These rapid fluctuations can destabilize local grid conditions, particularly when multiple large facilities are connected to the grid near each other. Utilities need to deploy more assets and technology to manage these risks, and companies that have not invested in these tools face heightened operational risk exposure.

3. Solutions

No single solution will provide a silver bullet. We anticipate that a combination of solutions will be needed and that solutions used will vary by company and site.

Table 1: Mapping Risks to Potential Solutions

Solutions		Risks						
		Community Opposition	Electric Infrastructure Delays	Data Center Utilization	Counterparty Credit	Water	Carbon Reduction Goals	Grid Reliability and Resiliency
Generation	Clean Energy	✓	✓			✓	✓	
Generation	Hourly Matching						✓	
Generation	On-Site Generation		✓					✓
Generation	Responsibly Produced Natural Gas						✓	
Contract Design	Flexible Interconnection	✓	✓	✓		✓	✓	
Contract Design	Tarriff Structures	✓		✓	✓			
Demand Side Management	Energy Efficiency	✓	✓	✓		✓	✓	
Demand Side Management	Demand Response	✓	✓	✓		✓	✓	
Demand Side Management	Distributed Generation	✓	✓	✓		✓	✓	✓
Demand Side Management	Virtual Power Plants	✓	✓	✓		✓	✓	✓
Grid Assets	Smart Grid Technologies						✓	✓
Grid Assets	Energy Storage		✓				✓	✓
Planning	Integrated Resource Planning		✓	✓			✓	
Planning	Transmission Planning		✓	✓			✓	

3.1 Generation

Clean energy

Building out new solar, wind, and energy storage resources, whether owned by the utility or procured through a power purchase agreement (PPA), can be a useful strategy for electricity companies to meet data center demand. Critically, this approach addresses data centers' requirements for speed to power because these resources can generally be brought online faster than new fossil fuel plants.

For vertically integrated utilities, investing in clean energy creates value because the return they are allowed to earn by regulators on capital investments exceeds their actual cost of equity. In addition, with this strategy, utilities convert operating expenses (such as fuel costs, which can only be recovered at cost) into capital expenses for assets, which utilities are allowed to earn a return on. Berkshire's subsidiaries and Xcel have successfully used this approach.

Additionally, this approach supports the goals that the data center and electricity companies may have for lowering emissions and adopting clean energy. It is important to note that the actual impact on system-wide emissions depends on whether it is existing, new, or restarted power generation. Contracting existing generation does not reduce emissions if the existing clean generation is already being fully utilized. But, building new or restarting mothballed, clean generation in lieu of fossil generation would improve the emissions profile of the systems.

As an example, NextEra Energy's partnership with Meta, covering 2.5 GW of solar and battery storage across multiple grid regions, demonstrates how clean energy and storage can be scaled quickly to meet data center demand and support Meta's goal of matching 100% of its operations with clean electricity. To get power quickly and minimize emissions, utilities, independent power producers (IPPs), and contract generators should build or contract new, clean generation.

Hourly matching

The standard approaches to clean energy procurement include purchasing renewable energy through long-term PPAs bundled with renewable energy certificates (RECs) or through purchasing RECs separately. Under the existing GHG Protocol, emissions inventories are calculated on an annual basis without 24/7 matching. This is an imperfect but reasonable approach for cleaning up data center load, considering that the main sources of commercially competitive electricity (wind and solar) are intermittent. This methodology allows a utility or data center owner to offset its annual electricity consumption with renewable generation, even though it does not ensure that clean electrons are flowing to the load during every hour. However, this type of annual accounting may mask continued reliance on fossil generation and exposure to location and time-sensitive electricity price risk.

Ideally, generation procurement structures would enable 24/7 carbon-free energy matching, making clean energy available and matched to load on an hourly basis. To minimize emissions, utilities, IPPs, and contract generators could explore options to supply data centers with hourly carbon-free energy matching technologies, potentially including new nuclear and geothermal in the longer term.

On-site generation

Speed-to-power constraints and interconnection queue backlogs have driven data center customers to explore on-site generation as a bridge or complement to grid electricity, and more recently as a replacement to grid electricity. On-site deployments are typically fossil fueled technologies that provide firm power such as reciprocating internal combustion engines, aeroderivative gas turbines, and fuel cells.

These technologies come with fuel price and emissions risk. Utilities and IPPs can help mitigate these risks by pairing on-site thermal technologies with on-site energy storage and on-site or nearby renewables like solar, wind, geothermal, or hydro. On-site or nearby renewables offer meaningful benefits and can improve community acceptance. If the output of solar, wind, and energy storage cannot meet a data center's full power needs across the year, an integrated approach that combines clean energy resources with firm on-site generation may be the most practical way to ensure reliability, reduce emissions, manage risk, and maintain project momentum compared to using older, less efficient fossil fuel plants.

Responsibly produced natural gas

As a last resort, if gas-fired sources of power are adopted, either in front of the meter or behind the meter, the generation owner should strive to source responsibly produced natural gas that minimizes methane emissions and leakage during production. Data from the U.S. Environmental Protection Agency shows that the highest-emitting gas producers leak methane at 20 times or more than the rate of the most efficient ones, meaning the sourcing decision itself has significant climate consequences.

By prioritizing certified, low methane natural gas produced under rigorous measurement, reporting, and verification standards such as the Oil and Gas Methane Partnership 2.0, power producers can meaningfully reduce the lifecycle emissions of the electricity they generate. This is increasingly important because more than 40% of generation serving data centers today already comes from natural gas, which means utilities and developers must not only ensure cleaner future gas supply but also address emissions from current gas-based generation. Taking these steps could also serve as a competitive advantage for power producers competing for data center contracts with technology companies looking to minimize their value chain emissions.

3.2 Contract Design

Flexible interconnection agreements

A potentially powerful tool available to utilities managing data center load is the use of demand response provisions in large load contracts, also known as flexible interconnection agreements. Rather than treating data centers as passive loads that must be always served in full, utilities could require large customers to accept reduced or interrupted service during grid emergencies or peak demand events as a condition of interconnection. This could be a relatively modest operational concession, given that large facilities typically have significant backup generation that can carry loads through brief grid stress events. However, it will be important to consider the data center's location since it may not be feasible to deploy backup diesel generation in non-attainment regions where air pollution levels persistently exceed national air quality standards.

Some industry experts have proposed that instead of using backup generation, data centers should reduce their loads during peak times either by shifting workload at a location to another time or by shifting workload to a different data center. This approach has been rejected by some data centers since they are expected to never be offline. However, EPRI's DCFlex project, which includes partners such as Google, Meta, Microsoft, Duke Energy, and PJM, is actively testing demand response and workload-shifting frameworks that include treating data center uninterruptible power supplies as grid resources.

If the utility can reduce its system peak demand, it will need less new grid infrastructure and generation, which would speed up interconnection and minimize costs for other customers. It may or may not reduce emissions depending on what resources are dispatched off-peak. Utilities could consider embedding flexibility provisions into large data center service agreements as a condition of interconnection, not merely as an optional add-on.

Tariff structures

For regulated electric utilities, the details of tariff structures matter. One of the most consequential decisions utilities faces is how to allocate the financial risk of demand forecasts that may not materialize. It is a risky strategy to build infrastructure in anticipation of forecasted load and expect to recover costs from all ratepayers if the load growth does not materialize. This socializes the speculative risk of data center development across residential and commercial customers who derive no direct benefit from the infrastructure. This risks the goodwill of regulators that are critical to rate-regulated utilities.

A more appropriate approach requires data center customers to bear meaningful financial exposure if their actual load falls short of contracted levels. Most states are implementing or considering large load tariffs to manage the costs and uncertainties of servicing data centers and to protect residential customers. Utilities should design large load tariffs and service agreements that include provisions requiring customers pay for a baseline amount of service, interconnection cost responsibility, and contractual termination penalties sufficient to recover stranded costs if a project is canceled. For IPPs and contract generators, the equivalent protection is robust "take-or-pay" provisions in PPAs. For all electricity companies, structuring contracts to manage counterparty risk is prudent.

3.3 Demand-Side Management

Energy efficiency

Energy efficiency should be the first solution implemented because reducing electricity consumption from residential, commercial, and industrial customers directly reduces the baseline load that utilities must serve, freeing up transmission and generation headroom for new large loads. This avoids building new generation or grid capacity that may be underutilized.

Demand response

A similar approach to minimizing new generation or grid infrastructure is to reduce the amount of electricity required at peak times. Demand response programs incentivize customers to reduce the amount of energy they use during these peak times either by using less energy overall (such as increasing the temperature for air conditioning) or by shifting that electricity use to another time of day (such as running the dishwasher at night). Besides reducing the cost of upgrades and freeing up grid capacity, demand response could reduce emissions if the overall amount of energy use is reduced or if the use is shifted to a time of day when generation produces lower emissions.

Distributed generation

Distributed generation, including rooftop solar and distribution-connected solar farms, can come online faster than large generation sources, helping meet data centers' speed-to-power needs. If significant quantities of distributed generation are added, it could reduce or eliminate the need for new large generation sources and potentially displace fossil generation. Pairing distributed generation with energy storage amplifies these benefits, while providing local backup power, voltage and frequency support, and faster recovery from outages that strengthen reliability.

Virtual power plants
Distributed generation, controllable load, and energy storage can be aggregated over many sites and customers to create a virtual power plant (VPP). VPPs can be dispatched to help meet system needs (or market signals), similar to traditional power plants, and help free up grid capacity and avoid expensive upgrades. Funding from data center developers could be used to sponsor community demand-side management programs. Recent examples include Google's energy efficiency fund for residents in the county surrounding its West Memphis data center in Arkansas and its funding of a VPP in PJM. In addition, Digital Realty has community solar agreements in Illinois.

3.4 Grid Assets

Smart grid technologies
Modernizing the grid with digital intelligence unlocks the full potential of energy efficiency, demand response, distributed energy resources, and virtual power plants. On the grid operator side, real-time sensors, automated switching, and AI-driven energy management systems give utilities the insight needed to optimize dispatch and respond to disruptions in real time. On the customer side, smart meters let large customers such as data centers participate in demand response and adjust consumption based on grid signals. These tools also enable more precise integration of variable renewables, reducing curtailment and system-wide emissions. They are also beneficial for reliability because they help grid operators detect and respond to rapid, synchronized swings in power draw caused by data centers before they cascade into outages. Technologies such as grid-forming inverters, flywheels, and synchronous condensers deliver similar reliability benefits by maintaining the precise, instantaneous balance between generation and load that grid stability requires, especially as more variable resources come online.

Energy storage
Energy storage is a critical distributed energy resource and grid asset. It allows excess clean generation to be stored so that it can be used during peak demand periods, reducing the need for fossil peaker plants. From a reliability standpoint, distributed storage provides a critical buffer against sudden load fluctuations, helps maintain voltage and frequency stability, and enables faster recovery from local outages.

3.5 Planning

Integrated resource planning
Vertically integrated utilities use the integrated resource plan (IRP) to develop, communicate, and seek approval for (if needed) their long-term generation strategy to regulators, investors, and the public. IRPs rely on power system models that find the optimal combination of resources that meet capacity, reliability, and affordability needs. The outcomes of these models are sensitive to assumptions on cost and the amount of clean energy that can be built in a year. How companies model clean energy resources is indicative of whether they are fully pursuing the clean energy build opportunity within state and federal policy settings.

In a period of extraordinary demand and policy uncertainty, the quality of IRP scenario analysis becomes especially important. Among several variables, the load growth assumption is important to test. The base case is most important since this has the greatest influence on the planned build, but the low case is also important since it reveals if and how the utility would adapt if data center growth were materially lower than the base forecast. Investors should treat a single-scenario IRP or one that lacks a credible low-demand case as a red flag. Investors should also watch for three further warning signs: limited disclosure of the assumptions utilities are using to model the costs of building new generation, assumptions about the costs of clean energy that diverge significantly from third-party estimates or peers, and modeling limits that are overly restrictive of clean-energy deployment.

Transmission planning

The single largest bottleneck to serving data center load and achieving a cleaner grid is transmission capacity. Interconnection queue backlogs are driven in part by inadequate transmission, particularly within and between areas that have significant clean energy generation and load centers.

FERC Order 1920 finalized in 2024 requires transmission providers to conduct long-term scenario planning over a 20-year horizon, explicitly accounting for large load growth scenarios, including data centers. This is a significant regulatory development since it creates the framework for proactive transmission planning that leads to more optimal investment, rather than reactive, project-by-project expansion.

As best practice, utilities should be active participants in regional and inter-regional transmission planning processes and should invest in transmission that enables both data center growth and clean energy integration. This includes modeling of lower cost upgrades of existing transmission lines described below when planning is controlled internally and advocating for this when planning is the responsibility of an ISO/RTO or a state transmission planning authority.

Software options such as dynamic line ratings can achieve moderate increases in capacity and be delivered in mere months. New greenfield transmission can take five to 15 years to come online. Fortunately, as noted by GridStrategies, there are many cost and time effective options for upgrades to the existing transmission grid to increase capacity. Important solutions include:

- Advanced conductors, which can increase capacity by 50% to 110% in 18 to 36 months.
- Double circuiting, which can increase capacity by 80% to 100% in two to four years.
- Voltage uprates, which can increase capacity by 100% to 300% in three to five years.

The solutions described in this report require deliberate strategic choices by electricity company management teams, and those choices are often visible in regulatory filings and processes and capital expenditure plans before they are reflected in financial outcomes. Institutional investors can engage electricity companies directly on these questions to shape strategy.

4. Shareholder Engagement

The following questions are designed to guide investors in direct engagement conversations with electricity company management. Ceres can work with investors to develop company-specific questions. These are organized by business model and are structured to be specific to elicit substantive answers.

4.1 For Regulated Utilities

On clean energy (for integrated utilities only)
Given that building clean energy can generate long-term shareholder value, how does the company plan to pursue this approach for meeting data center-driven demand, and how is this reflected in your subsidiaries' resource plans?

On water risk (for integrated utilities only)
Have you conducted a comprehensive water risk assessment that assesses your water dependencies in relation to areas of water stress? How frequently is this assessment updated and what tools are used?

On demand management (for all regulated utilities)
What percentage of your projected data center interconnection load is covered by demand response arrangements? How much have you been able to increase the utilization of your assets through demand management? Before approving new supply-side capital expenditure to serve data center load, what demand-side resource analysis have you conducted for the affected region? What is the maximum feasible contribution of energy efficiency, distributed solar, distributed storage, and virtual power plant programs for freeing up capacity for new large loads?

On overbuild risk (for integrated utilities only)
What is your IRP's low-demand scenario for data centers? What are the financial consequences if that scenario materializes? How would generation asset operations and lifetimes change?

On overbuild risk (for all regulated utilities)
What contractual protections do you have in place to manage counterparty risk and ensure that speculative demand does not create costs borne by other ratepayers? Do the states that you operate in have dedicated large load tariffs that protect existing customers, or are they in the process of implementing these?

On transmission planning (for all regulated utilities)
How do you reflect faster and cheaper options for upgrading existing lines, in addition to greenfield projects, in transmission planning processes? Do these planned upgrades support increasing adoption of clean energy?

4.2 For Independent Power Producers

On market exposure
Considering the combined generation and retailing businesses, are you a net seller or net buyer in wholesale electricity markets in average and peak conditions? What share of your installed capacity in competitive markets is currently contracted long-term versus exposed to capacity market and energy market prices? How has the recent price surge affected your contracting strategy?

On counterparty risk
What clauses to protect against counterparty risk do you include in bilateral agreements with data center customers?

On revenue and volume risk
Do your contracts include take-or-pay provisions that provide revenue certainty if a customer's load fails to materialize as projected?

On clean energy positioning
What is your strategy for positioning your generation fleet to serve hyperscalers who have committed to carbon-free energy? What percentage of your proposed data center contracts are structured around hourly CFE matching rather than annual accounting?

4.3 For Contract/PPA Generators (including Renewable Developers)

On contract structure
Do your data center PPAs include take-or-pay provisions that protect revenue if contracted customers fail to build on forecasted timelines, or use the data center less than expected? What is your largest single counterparty concentration?

On clean energy standards
Are you structuring new data center contracts around 24/7 hourly carbon-free energy matching or annual accounting?

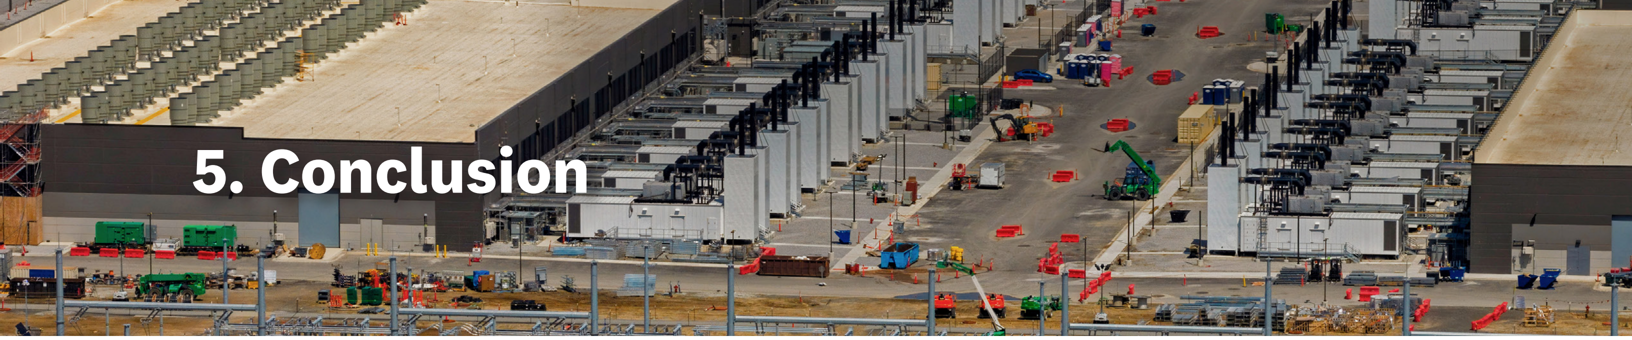

5. Conclusion

The AI data center boom has arrived at a moment when the U.S. electric grid is simultaneously under pressure to become cleaner, modernize, and expand at a pace it has not attempted in generations—all while keeping electricity affordable. For electricity companies, this is the defining strategic challenge and opportunity of the next decade.

The electricity companies that will create lasting value for investors are those that protect themselves and their ratepayers from speculative data center growth and treat this moment as a catalyst for the transformation the grid needs: accelerating the buildout of clean energy, modernizing the transmission and distribution grid, and deploying flexible resources that make the system more resilient. While the utilities that overly rely on fossil fuels to meet data center growth, build infrastructure on speculative demand without adequate contractual protection, and allow the burden of grid investment to fall on residential ratepayers will face financial and reputational risks.

The questions in this report are a starting point for investor-corporate engagement on data center growth. The goal is to establish that electricity companies must demonstrate they are managing data center growth in a way that is financially disciplined, responsible, and fair to the full range of customers and communities they serve.

About Ceres

Ceres is a nonprofit advocacy organization working to accelerate the transition to a cleaner, more just, and resilient economy. With data-driven research and expert analysis, we inspire investors and companies to act on the world's sustainability challenges and advocate for market and policy solutions. Together, our efforts transform industries, unlock new business opportunities, and foster innovation and job growth—proving that sustainability is the bottom line. For more information, visit ceres.org.

Acknowledgments

Lead Authors

Colette Lamontagne, Senior Director, Electric Power
Olivia Schadler, Manager, Energy

Thanks also to the colleagues at Ceres who provided invaluable editorial review and design support with this project, including Rob Berridge, Steven Clarke, Miranda Eisenman, Heather Green, Giselle Hancock, Sarah Hicks, Rebecca Hoffman, Kirsten James, and Kelly Trombley.

Ceres would like to thank As You Sow, Mercy Investment Services, and Trillium Asset Management for their thoughtful feedback as external reviewers.

This publication was made possible by the generous support provided by Omidyar Network.

The views or opinions expressed in the report do not necessarily reflect those of the organizations listed above.