September 2024

Could Data Centers be the Catalyst for Modernizing the U.S. Electric Grid?

Teri Viswanath
Lead Economist,
Power Energy and Water

Jeff Johnston
Lead Economist,
Digital Infrastructure

Inside…

Key Points:

- A new age of electricity growth is here. Data center demand growth is the most visible sign of the changing times, testing the ability of utility suppliers to keep pace.
- The fundamental mismatch in infrastructure build-out – between the time it takes to build a data center versus the time to build the electricity infrastructure – demands novel solutions.
- Creative adaptation has always been the name of the game but this time around some of the more innovative solutions are actually coming from consumers – at a scale and pace that will have lasting impacts on the industry.
- Indeed, data centers could shift data processing workloads to times or geographies where the grid is less congested. What's more, AI holds the promise on reducing the remaining energy consumption for other consuming sectors.
- Data centers are just the first wave in the fourth Industrial Revolution, launching the world into a full-fledged digital economy.

The new age of rising electricity demand tests supply response

The Industrial Revolution – both as an historic event and a concept – irreversibly tied economic growth with energy demand, with new factories and machinery creating a global race for securing energy supplies.[1] Start-up activities, related to new industrial or economic cycles, tend to inflate energy intensity. Competition and the need to reduce product input costs eventually flatten, then reduce the amount of energy consumed to achieve the same level of output.

An analysis of decades-long U.S. electricity demand growth reflects this tendency, with three distinct electricity growth cycles corresponding to major transformations in the broader economy – or what the World Economic Forum labels Industrial Revolutions 1-3 *(Exhibit 1).* Each cycle is explained by expansive electrification or hyper-growth, followed by a period of slower consolidation and then declining growth. Yet, there was a break in this cycle about two decades ago, with the loss of energy-intensive manufacturing, massive efficiency gains and a slowing economy. Now, that chapter appears to be ending, with the acceleration of the digital economy (Industrial Revolution 4), reshoring of American manufacturing and an "electrify everything" moment ushering in new cycle of surging consumer demand. The question for the next decade is not whether electricity demand will grow, but rather to what scale and will this source of new growth stick around?

It's hard to pin-point the exact moment when our collective mindset changed about electricity growth but the last annual outlook from the U.S. Energy Information Administration (EIA) EIA Annual Energy Outlook release might have been that inflection point. Back in March 2023, the agency kept to its no-load-growth mantra, predicting just 0.07% of annual electricity demand growth through 2033. By the time the North American Electric Reliability Corporation (NERC) released its updated Long-term Reliability Assessment in December, the bottom-up tally from grid operators persuaded the habitually conservative regulator to double its decade-long demand projections to 1.24%. But it is NERC's comments from its last installment that are particularly telling – "electricity peak demand and net energy growth rates are increasing **more rapidly than at any point in the past three decades.**" More recently, a host of respected market observers have joined the chorus with even higher predictions (the S&P Global estimate at 1.8% and Goldman Sachs at 2.4%), succinctly mapping out the genie's evacuation route from the bottle *(Exhibit 2).*

EXHIBIT 1: Demand is surging for the first time since the 1990s

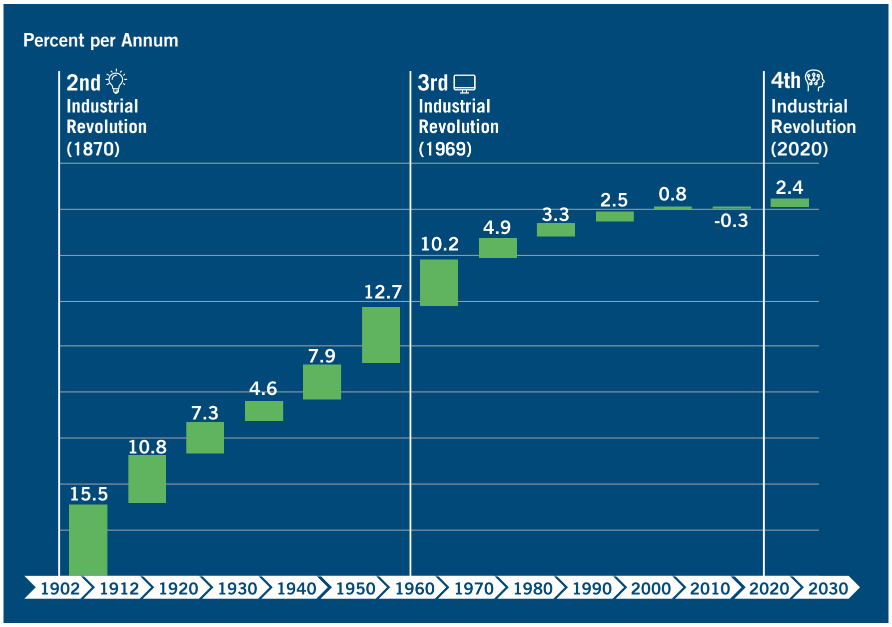

Source: Pre-1950, Edison Electric Institute; Post-1950 EIA

EXHIBIT 2: U.S. electricity consumption, actual and projected

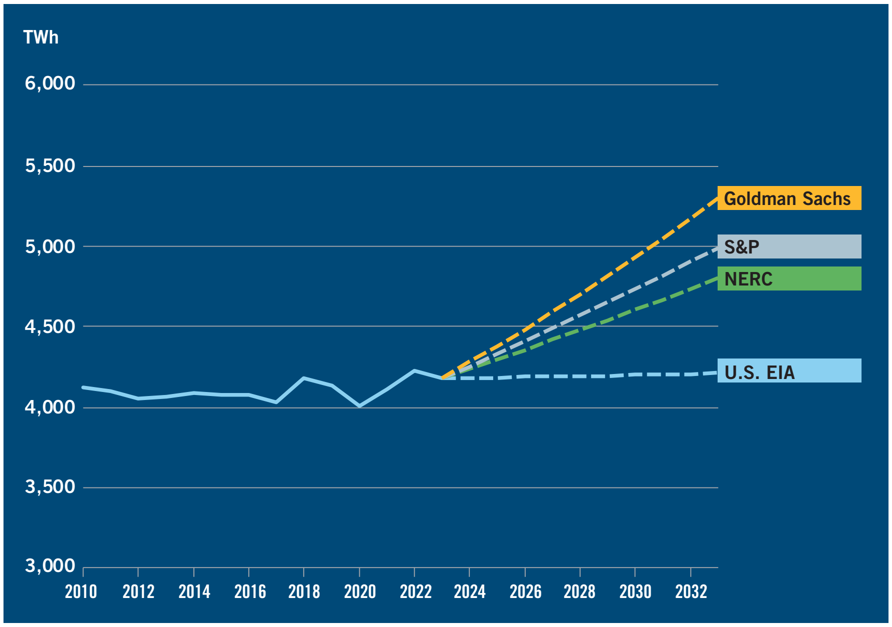

Source: CoBank analysis

The coming tidal wave of commercial sector energy demand

From our review of market projections as well as the real-time estimates coming in from EIA, the most visible sign of changing consumer behavior is coming from the commercial sector. Indeed, we believe that the emerging mega-loads from Gigawatt Data Center Campuses will ultimately test the ability of utility-suppliers to keep pace in this new age of rising electricity demand.[2] It is the trend toward building larger data centers with greater performance requirements – along with the sheer number of facilities in development – that places a novel burden on the industry. Data center energy use in the country has already doubled over the past three years, with researchers from the U.S. Department of Energy's Lawrence Berkeley National Laboratory expecting the segment to grow sharply by 2027 *(Exhibit 3).*

The lab's preliminary guidance, issued ahead of its soon-to-be-released detailed national analysis on behest of the DOE, suggests that data center electricity use has risen from 175 terawatt hours (TWh) in 2021 to about 375 TWh presently. Emerging power requirements for artificial intelligence, however, will likely push consumption past 500 TWh by 2027. Based on these estimates, the electricity demand from the nation's data centers would double in just four years, going from an estimated 4% of total electricity generated annually in 2023 to roughly 8% in 2027.

EXHIBIT 3: Annual electricity demand

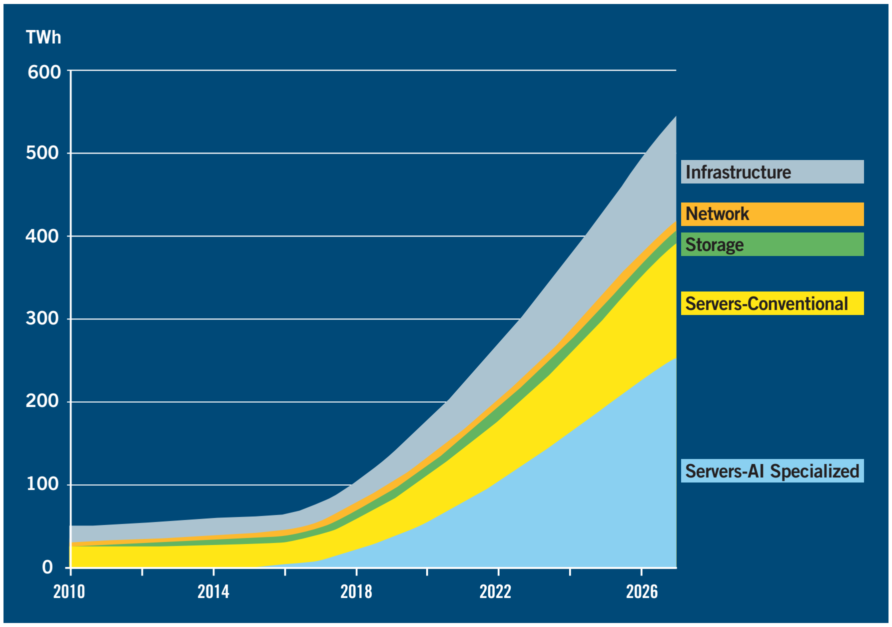

Source: DOE Lawrence Berkeley Labs

EXHIBIT 4: Risk area summary 2024-2028

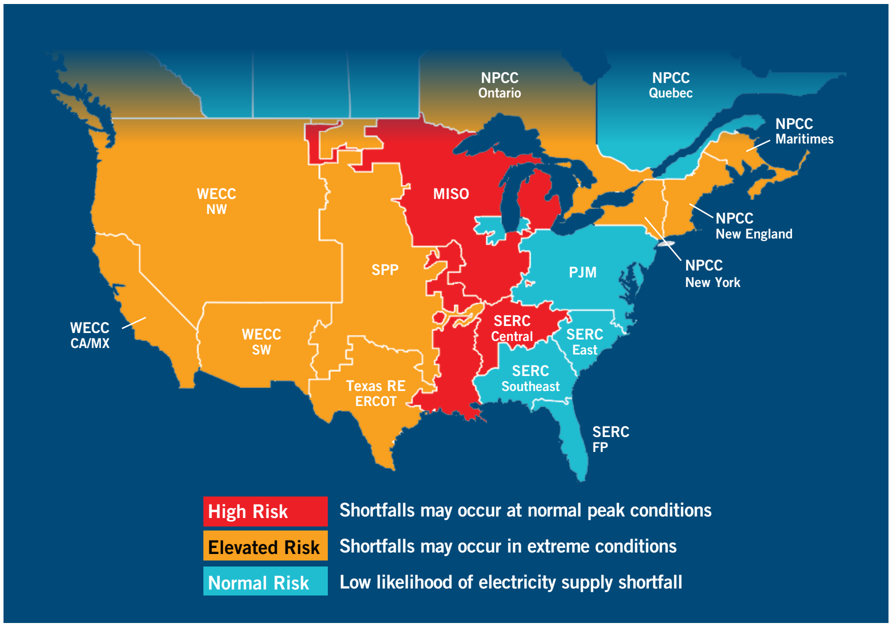

Source: NERC

Commenting on the relative preparedness of the country to meet this surge in demand, NERC cautions that the North American bulk power system is on the cusp of large-scale growth, **bringing reliability challenges and opportunities to a grid** that was already facing unprecedented change. In other words, this new data center demand is going to strain an already strained grid[3] *(Exhibit 4).* The **reliability challenges** part of NERC's message has been well received and amplified in the media (NYTimes, WSJ, Washington Post, NPR, Fox News). Yet, the **opportunities** the agency hints at have received far less attention. This is a mistake.

Arguably, the grid has been in a constant state of change for the past 140 years, with many moments in that history that feel much like the current one. Consistently, creative adaptation has occurred as the grid has evolved to meet the challenge of the day. But this time around, some of the more novel solutions are coming from consumers and the very data centers whose AI needs are propelling the change.

The emerging blueprint for sustainable data center development

Utilities consider large-load power connection agreements as long-term obligations – commitments, if you will, that might extend beyond the life of a data center. A sizeable proportion of the infrastructure developed to serve emerging industrial loads of the 1960s, 70s and 80s proved a bad investment – with efficiency gains, on-site generation and offshoring ushering in a grid exodus in the 1980s, 1990s and 2000s. Consequently, the hefty capital costs involved, protracted development lead-times and the need to keep power plants operating long enough to make a reasonable return on investment justify asking the question of whether these emerging data centers will have staying power. The short answer is "yes" but it is based upon growing evidence that digital transformation is driving economic growth in the country and that growth will support more expansive upstream supply investments. Or, more simply stated: The same sort of economic cycles that drove past electricity demand growth and hefty grid investments are occurring today.[4]

Energy + Environmental Economics (E3) Consulting recently described the grid's growing pains related to data centers as the following:

*"**If the transforming grid is a traffic jam during highway construction, then data centers are a large convoy of trucks with urgent deliveries pulling into the on-ramp.** This confluence of factors creates a gridlock, where utilities and regulators are overwhelmed working to modernize and decarbonize the grid, while managing queues of generators and new loads seeking interconnection, all bottlenecked at the same constraint. This may lead to suboptimal outcomes if grid decision-makers only see limited near-term options such as delaying new large loads interconnections and/or delaying retirement of existing fossil fuel generators."*

So, if we agree that meeting these new, large loads is important for achieving the next cycle of economic growth and that decarbonization of the grid will continue, how do we break the gridlock that E3 highlights? Acknowledging that there is not an easy solution waiting in the wings, we do see promising examples of where the industry is headed and how data centers might play an essential role in solving this current existential crisis.

Meeting digital infrastructure needs while protecting existing members

Focusing first on the point of local data center connection and traveling upstream, we mined our earlier podcast conversation with Northern Virginia Electric Cooperative's (NOVEC) CEO, Dave Schleicher for specific guidance here.

NOVEC is the largest cooperative in Virginia serving across six counties, predominantly Loudoun county and Prince William county just on the western side of the Washington D.C. metro area. Loudoun county's "data center alley" is home to the world's largest concentration of data centers, with more than 25 million square feet currently in operation. As a result of the rapid development in Northern Virginia, NOVEC has had to develop an agile response to the rapid progression in digital infrastructure development – from small, enterprise-owned and operated data centers to colocation services and now hyper-scalers. Facilities that were maybe 15 megawatts

EXHIBIT 5: Average lead times to build key load and capacity assets

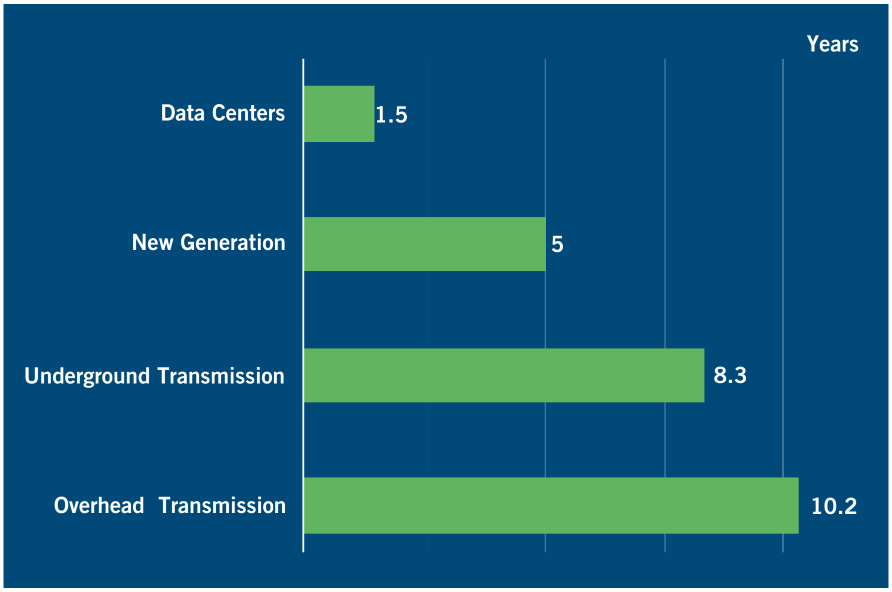

Source: IEA/Third Way; GridStrategies

a decade ago have now evolved into gigawatt campuses, requiring multiple substations and posing a different set of problems for a distribution utility that could find itself overwhelmed by overnight growth.

This discussion was particularly insightful on how to manage and accrue system benefits from this type of utility-scale growth without placing the existing membership in harm's way. Dave talked about streamlining or standardizing the intake process with data centers to achieve better collaboration and the need for financial safeguards to limit greater member exposure to upstream system costs of servicing these mega-loads. This translates to collecting interconnection costs upfront and establishing rigorous pass-through processes for costs associated with accessing grid supply. This conversation stressed the importance of having a deep understanding of the data center power needs, the timing around energizing the facility and potential flexibilities around supply, etc.

More recently, Duke Energy, a North Carolina-based utility, announced similar data center cost-sharing strategies to ensure these emerging consumers pay their fair share of grid improvements. To address potential distribution losses related to a mismatch in demand, Duke's new data center tariff will include take-or-pay terms for supply and up-front infrastructure build out payments. That is, Duke will require "minimum take" clauses that compel the facilities to pay for a certain amount of power regardless of how much they use. These contract terms are intended to provide a buffer between the growing infrastructure build-out required for these new loads and rising power bills for existing customers in the network.

Why accelerating the G&T upstream response matters

Yet, even if we solve new point load challenges at the interconnect, data centers may consume large quantities of energy that could threaten grid reliability further upstream if left unmanaged. And, as we move beyond distribution to transmission and generation, this is perhaps where the real systemic challenges lie. It's not just that cost allocation becomes more complex (it does),[5,6] but there is also a fundamental mismatch in the lead time required to construct a data center versus building new generation or wiring new large-scale transmission lines *(Exhibit 5)* – the latter representing the largest impediment to additional grid supply. Indeed, should all the proposed capacity awaiting in a transmission queue get built, the nation would double the amount of generation currently installed *(Exhibit 6)*.[7] The slow-walked development in our power delivery systems since 2015 *(Exhibit 7)* not only represents the primary challenge to meeting new electricity demand growth, but is also costly to existing consumers.[8]

To circumvent the siting and permitting,[9] as well as the far-flung who-will-pay-for-it challenges, more targeted spending is occurring to address some of the known transmission constraints for serving data centers.[10] Last December, PJM approved an "unprecedented" transmission package totaling $5.1 billion with roughly half the budget earmarked to address the incoming load from 7.5 GWs of new data centers in Virginia and Maryland. This funding is incredibly important, but it was nearly two years in the making and most of the spending will be aimed toward brownfield development – that is, using existing facilities and rights of way rather than greenfield development (which is new construction on land without existing transmission lines). Moreover, the quick-fix technologies currently being deployed are unlikely to meet the full scope of national transmission needs required to accommodate the more expansive growth down the road.

Further sidestepping these existing grid challenges is the interesting development of building data centers closer to captive supply. A spectacular example of this occurred earlier this year with Amazon's quiet acquisition of a data center linked to the Susquehanna Steam Electric Station in Pennsylvania (the country's sixth-largest nuclear plant). Talen Energy developed the co-located, direct-connect 960-megawatt data center last year – a sort of Field of Dreams "build it and they will come" development strategy for the independent power producer that appears to have paid off. While the industry has witnessed these types of supply strategies on a smaller scale, the Talen-Amazon transaction opens a broader discussion about what hyper-localization looks like through a large utility-scale lens.

EXHIBIT 6: U.S. installed electric generating capacity vs. interconnection queue

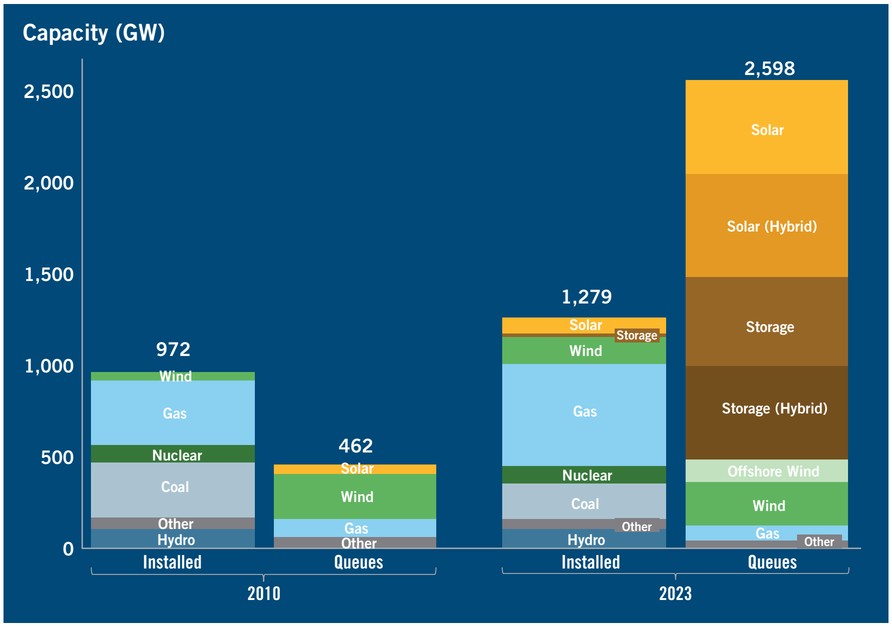

Source: DOE Lawrence Berkeley Labs

EXHIBIT 7: Buildout of high-voltage transmission lines, 2010-2023

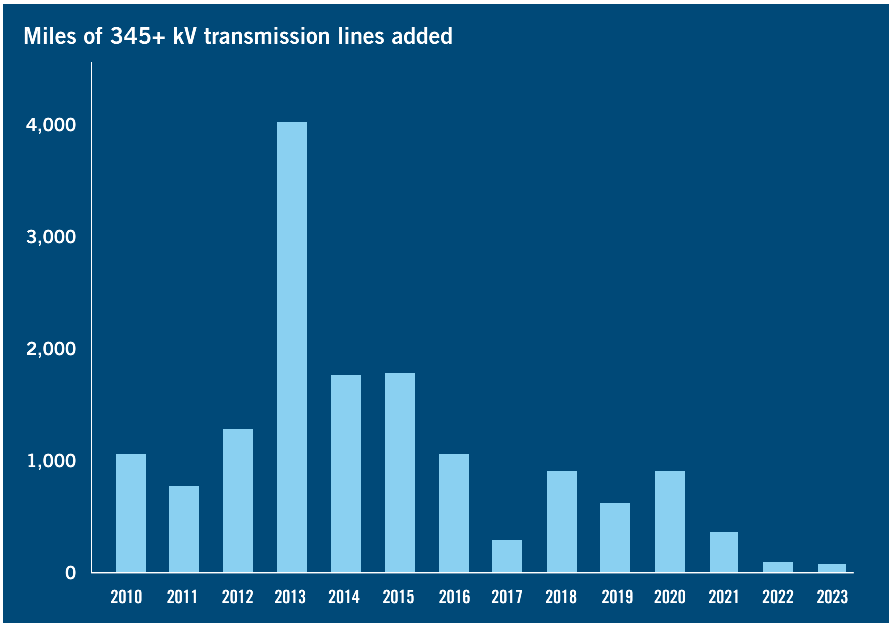

Source: IEA/Third Way; GridStrategies

Not to be outdone, Microsoft just inked the largest-ever power purchase agreement with Constellation for exclusive power supply sourced from the formerly shuttered Three Mile Island nuclear plant. Expounding on the importance of restarting these formerly shuttered resources, Joe Dominguez, Constellation Energy's president and CEO stressed, "powering industries critical to our nation's global economic and technological competitiveness, including data centers, requires an abundance of energy that is carbon-free and reliable every hour of every day, and nuclear plants are the only energy sources that can consistently deliver on that promise."

Certainly, the key advantage of this approach is avoiding "gridlock" but there are disadvantages of having a gigawatt-scale data center campus detached from alternate grid supply solutions. After all, the redundant design the grid affords is exactly the redundancy and high availability that data centers require. Specifically, there is the promise of multiple power sources, expansive connectivity and backup systems – the very impetus around the greater grid development of the 1960s and 1970s.

Yet, given the upstream supply-side challenges mentioned, data centers are taking it upon themselves to find their own solutions[11] – some of which are naturally occurring as the industry itself evolves and matures.

Location, location, location: How power constraints rework data center geography

One such strategy is for these mega-consumers to be much more intentional about macro-grid demand response opportunities. While generative AI is driving greater compute power, it also lends itself to greater geographic diversity because of the way that AI workloads are now architected. Let's unpack this…

Generative AI can be broken down into two core components: learning and active inference. The learning aspect of generative AI is incredibly compute-intensive, and therefore consumes an enormous amount of energy. Meta CEO Mark Zuckerberg recently disclosed that its next iteration of Llama (Llama 4.0) – the company's collection of open-source large language models first released in February 2023 – will require 10x more computing power to train the model. Training an AI model involves processing billions of parameters which effectively makes the system smarter. And as new AI models are released, the number of training parameters is increasing exponentially. For example, Llama 3 was released in April with 8 billion parameters – in July, Llama 3.1 was released with whopping 405 billion parameters.

Unlike AI inference, AI training does not require low latency network performance. As a result, it can be done in "remote" data centers that are far away from where applications are being run, possibly closer to where sustainable power supply exists. Increasingly, this development translates to infrastructure builds taking place outside traditional data center markets (Northern Virginia, Dallas, Chicago, Phoenix and Northern California).[12] That (theoretical) flexibility could also extend to when those training workloads are run to better align with intermittent renewable energy resources – process more workloads when the sun is shining and wind is blowing.

The other important factor is how operators virtualize their workloads within and between data centers. The concept of virtualization uses software to create an abstract layer over computer hardware that enables the division of a single computer's hardware components into multiple virtual machines. This technology enables operators to dynamically move workloads between computers based on factors including storage and compute resource availability, latency requirements and of course, power availability. Despite these technological advancements, they will not provide much relief to the energy complex in the near term.[13] There is the possibility that building greater "spare capacity" within the system could more efficiently use the current grid as it is configured. [14]

Where might new development take place? Prior off-shoring of the U.S. energy-intensive manufacturing during the 1980s, 1990s and 2000s has potentially opened an interesting option to utilize the existing transmission backbone developed in non-metro regions of the country – a story that has already played out in the Northwest, where server farms replaced aluminum smelters.[15] But, there

EXHIBIT 8: 2024 typical power shares

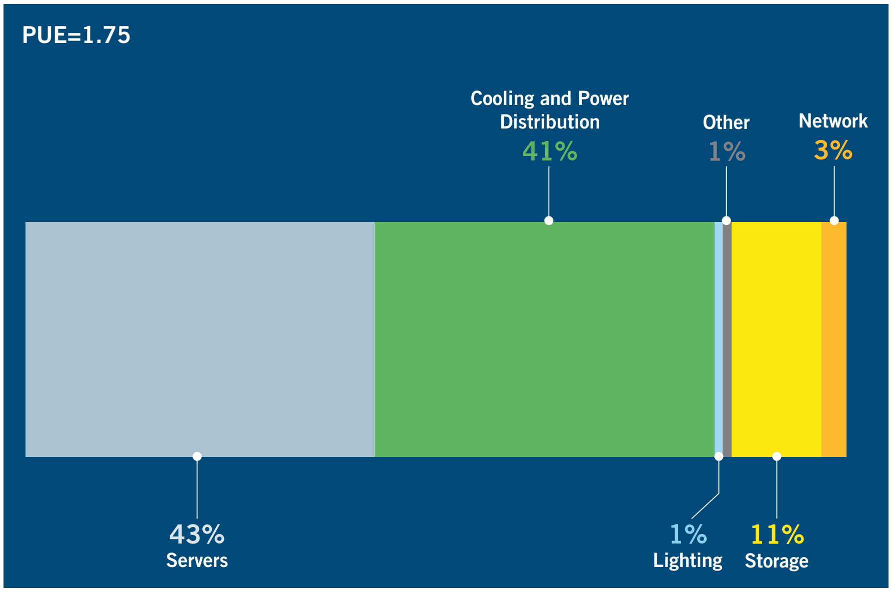

Source: NVIDIA

are also interesting pilot projects that could proverbially take these data centers off-shore, such Denmark's plan to connect giant offshore wind farms to new artificial "energy islands" housing data centers.

Over the next several years, novel approaches of scouring the grid for existing excess power capacity in non-traditional locations, greater behind-the-meter generation solutions and more expansive opportunities for load shifting are the pathways data centers will likely pursue to supplement their normal utility power systems.

Next step: Reduce energy waste, enhance efficiency across the stack

The next obvious step toward self-sufficiency will be to invest heavily in reducing energy waste. Afterall, the least costly form of electricity is the kilowatt-hour not consumed, so efficiency is going to play an increasingly important role toward fulfilling future supply requirements.

In the early 2000s there were similar concerns that the pace of electricity consumed for data centers was on a collision course with scarce energy resources. Lawrence Berkeley Lab researchers noted that electricity use was doubling every five years, due to explosive growth in both the number and density of data centers. Yet, something surprising happened by 2005: these large data centers made significant advances in operating more efficiently, which mostly kept demand constant for the next decade. Servers improved their power scaling abilities, thus reducing power draw during idle periods or when at low utilization. Efficiency improvements in storage, network and infrastructure likewise contributed to the efficiency gains that slowed electricity demand growth from this sector. Yet, in today's generative AI era, when workloads and the systems running them have changed dramatically, there is a greater need to drive energy efficiency. Arguably, new ways of measuring data center efficiency amidst surging energy demands are needed, targeting the two primary areas of consumption – the twin engines of powering servers or the "compute power" and the need to keep that equipment cool *(Exhibit 8).*

The demand for computational power is skyrocketing with generative AI, with the semiconductor industry finding itself approaching a new S-curve – with greater capital now being spent on advancements in chip design, materials, and architectures. In 1965, semiconductor legend Gordon Moore developed Moore's Law, which posits that the number of transistors on an integrated circuit will double every two years with minimal rise in cost (and energy consumption). This law held true throughout the last several decades of the semiconductor industry but over time it has become challenging to continue its trajectory due to transistor miniaturization. For example, Nvidia's AI Hopper architecture packs more transistors on its chip versus previous versions, and it costs more. Despite these challenges, the industry is still moving ahead with more energy efficient generative AI chips. What's more, with the passing of the **CHIPS and Science Act** policymakers have taken the first critical step in addressing the nation's **ability to maintain Moore's Law leadership.**

There are already promising signs or green shoots that the industry investment will pay off. Nivida announced its latest generative AI architecture dubbed Blackwell. The Blackwell platform promises to reduce Large Language Model (LLM) interference operating cost and energy by up to 25x over its previous platform, while it massively boosts the chips' computing capabilities. While this is great news for the energy market, the ultimate impacts are still an unknown. Will future AI models' compute and learning parameter requirements offset the efficiency gains from Blackwell and similar platforms? Or will these gains be able to bring into balance the projected gap between energy supply and demand? It's simply too early to tell at this point.

Of possibly lesser importance in the larger scheme of power use will be greater facility efficiency gains in cooling these overworked servers. Simply speaking, computing creates heat. Data centers, which are filled with servers, generate lots of heat. As much as 40% of electricity is being used to keep these facilities operating within a 64.4 to 80.6 degrees Fahrenheit (18 to 27 degrees Celsius) range. Yet, for every 1°F increase in server inlet temperature data centers can save ~ 5% in energy costs. Consequently, raising server inlet temperatures has been an important directive for many data centers to reduce electricity demand, with next-generation equipment rated for a maximum inlet temperature of 95 degrees Fahrenheit. The way these facilities are cooled is also being re-envisioned with an increasing number of data centers adopting liquid rather than traditional air-cooling equipment. While air cooling has been the standard due to its relative flexibility and ease of use, it is failing to keep up with the heat output in modern server farms.[16] Liquid technology innovations – from transferring heat to a liquid far from the source to immersion cooling – promise to use just 20% of the current cooling requirements.

Over the past two decades, data centers have nearly doubled their energy efficiency. Yet, given electricity supply limitations, greater emphasis will need to be placed here. With this in mind, the industry itself has suggested changing the ways in which efficiency is measured by taking compute performance into consideration.[17]

More interesting still is that changing perceptions around energy efficiency might permeate beyond the data center walls to the rest of us. If we think about the E3 analogy of highway construction, the solutions for data centers suggest that they might bypass the on-ramp or travel more intentionally. Yet, the fact that the utility industry itself is investing in generation AI at the fastest clip of any other industry, other consuming segments might see important efficiency gains. Meaning that AI holds the promise of luring existing drivers off the congested grid highway. By automating and optimizing energy management – such as adjusting thermostats, charging EVs or managing passive use appliances – the 4% (data centers) could influence energy efficiency for the 96% (the rest of us).

Lastly, if we consider prior economic cycles, where hyper-growth was followed by consolidation and then a decline in electricity demand, there are reasons to believe that this super-charged data center development cycle will wash-rinse-and repeat in the digital age, with more growth ahead.

The age of rising demand could usher in more efficient, affordable and reliable electricity

It might be said that the mass development of data centers might be arriving at just the right time. . .to make critical investments in the grid and more efficiently use existing infrastructure, as well as to catalyze a new era of energy efficiency. Sizeable investment from these well-capitalized consumers could potentially break the gridlock currently stalling the massive infrastructure needed, potentially rewiring the industry in a manner that has broader benefits.

Right now, there is a terrific opportunity for the utility industry to be intentional in how data center demand is met, setting an important precedent for handling future load growth in a way that better supports the grid. There is a unique opportunity to invest in the broad array of solutions starting with consumer efficiency and demand management and working upstream to provide ample, reliable and secure power in an age of rising electricity

demand. It is important to realize that even if the demand growth for data centers flattens or even reverses, it will set utilities on the right path for preparing for the subsequent wave of electricity growth driven by strong industrial policies and electrification of everything. That said, establishing priorities will remain critical, requiring the collective grid — from the consumer to the distribution system and beyond — to collaborate on solutions to avoid unintended consequences and optimally capture long-term benefits.

Notes

[1] As economic engines revved, energy consumption increased. During periods of energy resource scarcity, these same economies contracted, with governments re-prioritizing access to new supplies and fortifying energy security policies. Yet, the historical relationship between the world's developed economies and energy — measured by energy consumption required for marginal GDP growth or energy intensity — is nuanced and is still evolving, with no single blueprint for achieving affordable, sustainable energy security.

[2] According to EIA, commercial electricity demand in the 10 states experiencing rapid development of large-scale computing facilities increased by a combined 42 billion kWh between 2019 and 2023, representing growth of 10% in those states over that four-year period.

[3] According to the agency, much of the country is now at risk of running short of power due to this combination of higher power demand, generator retirements and the potential for replacement resources to "fall short of capacity and energy needs."

[4] One of the more lasting impacts of COVID was the acceleration toward this digital economy. McKinsey estimates that the pandemic accelerated digital transformation by about 6 years in the U.S. What's more, this sector of the economy is having a profound impact on top-line economic growth. The U.S. Commerce Department's Bureau of Economic Analysis recently took a first stab at estimating the size of the country's digital economy. The government arrived at a $2.6 trillion valuation for America's "digital GDP" in 2022 or a little over a tenth the size of that year's $25.7 trillion estimate for the total economy— were the U.S. digital economy an independent country, it would have ranked eighth in the world. More interesting still is the rate of change of the digital economy compared to the country's overall top-line growth. In 2022, digital economy real value added grew 6.3%, compared to total U.S. real GDP growth of 1.9%. Expectations of continued blistering growth means that a growing proportion of GDP in the future will be based on a digital platform. Indeed, some estimate that 70% of the new economic value globally will be created on digitally-enabled platforms over the next decade.

[5] The process of cost-allocation "generally works neatly for distribution system poles and wires and the smaller-scale local transmission projects. . .large-scale transmission projects are far more difficult to execute because they can cross multiple jurisdictions, utility systems, market operator borders, transmission planning regions, and the regulatory boundaries of state public utilities commission." "Assessing Electric Transmission's Cost Allocation Dilemma", Center for Strategic International Studies (CSIS), Cy McGeady, 6 October 2023.

[6] Recently issued FERC Order No. 1920 requires transmission planning regions to develop long-term planning using a common set of standards. The regions then are required to establish a mandatory default cost allocation methodology while allowing for states to provide input on an alternative cost allocation methodology. Costs to be assigned must be commensurate with benefits so that no entity can be forced to pay more than the benefits they receive. While this rule might have created a 'bright line' for cost allocation, it is being widely challenged, with the Supreme Court's recent decision on the Chevron deference providing greater traction for the appeal.

[7] Berkeley Lab's updated count of new proposed power plants seeking transmission interconnection as of the end of 2023 revealed 11,600 projects representing 1,570 gigawatts (GW) of generator capacity and 1,030 GW of storage actively waiting in the queue. The gridlock for these projects has increased nearly eight-fold over the last decade and is now more than twice the total installed capacity of the existing U.S. power plant fleet.

[8] Just how big is this problem? The DOE's latest installment of their "National Transmission Needs Study", found that to meet high-load growth by 2035, within-region transmission to integrate new assets must increase by 128% and inter-region transmission by 412% (a quintupling of the nation's current interregional transfer capacity, respectively). Even if approved and sanctioned, grid construction still faces supply chain hurdles, including lead times of more than a year and a 400% price surge for large-power transformers, compounded by a shortage of specialized steel.

[9] There are enormous challenges in addressing the sort of transmission development that will be required to fundamentally satisfy both growing demand and new sources of supply – an exercise of putting the genie back into the bottle. Unlike other kind of long-distance infrastructure regulation and planning, there is a century-long history of patch-work development completed by a myriad of local and regional overseers that are not easily sidelined. Re-asserting the federal government's role in overseeing transmission development took a leap forward with FERC's transmission rulemaking over this past year – the first effort in more than a decade to specifically address key issues that are preventing more expansive greenfield work. Yet, state regulators, advocacy groups and others have already initiated litigation, suggesting that a quick fix in the form of these new existing federal regulations is probably not going to happen.

[10] What's more, there are real concerns that the very targeted way in which spending is now occurring might not be the most efficient. There is little argument that high-voltage regional or inter-regional transmission lines provide real value to ratepayers in terms of grid services per dollar spent. Yet, researchers at Brattle found that of the approximately $25 billion in transmission investments between 1996 and 2020, 90% were driven by local reliability needs or quick-fix solutions for delivery problems already identified.

[11] While somewhat novel for both industries, there are numerous examples where whole industries have relocated to secure low-cost energy supply. The rising cost of natural gas in the early 2000s forced the closure of many small fertilizer companies, with increased domestic reliance on fertilizer imports.

[12] At present, data center capacity is highly concentrated, with fifteen states accounting for 80% of the national data center load. Virigina can boost of the greatest concentration of data center capacity with these consumers accounting for a quarter of that state's electric load. But, changing requirements are altering the landscape, with greater flexibility in siting these facilities beyond Tier One metros where power is severely limited to Tier Two locations such as Columbus and Reno where power is more likely to be available and beyond. Oklahoma, Georgia, and Mississippi are further up-and-coming markets for data centers due to available power.

[13] There is a supply - demand imbalance in the market as most data centers are already running at maximum capacity, and the new capacity coming online has already been leased.

[14] The idea of 'spare capacity' within the upstream oil market is not a new concept; however, increasingly this idea is being conceptually applied to renewable energy resources and now possibly downstream consumers.

[15] Low-cost electricity lured the aluminum companies to the Northwest in the 1940s and 1950s; increased prices for raw materials, volatile electricity prices, and worldwide competition with other, often lower-cost production and supply, eventually drove them out of business.

[16] Liquids have a higher heat capacity and can thus absorb heat more efficiently than air does. Water for example can absorb heat some 30 times faster than air due to its greater thermal conductivity.

[17] And here too, the industry has been pushing the envelope with Jules Verne-like subsea experimentation with Microsoft's Project Natick underwater data center.

CoBank's Knowledge Exchange Division welcomes readers' comments and suggestions.
Please send them to KEDRESEARCH@cobank.com.