The new data center collaboration for utilities and developers

Moving from demand growth to value creation

The better the question. The better the answer.
The better the world works.

EY Parthenon
Shape the future with confidence

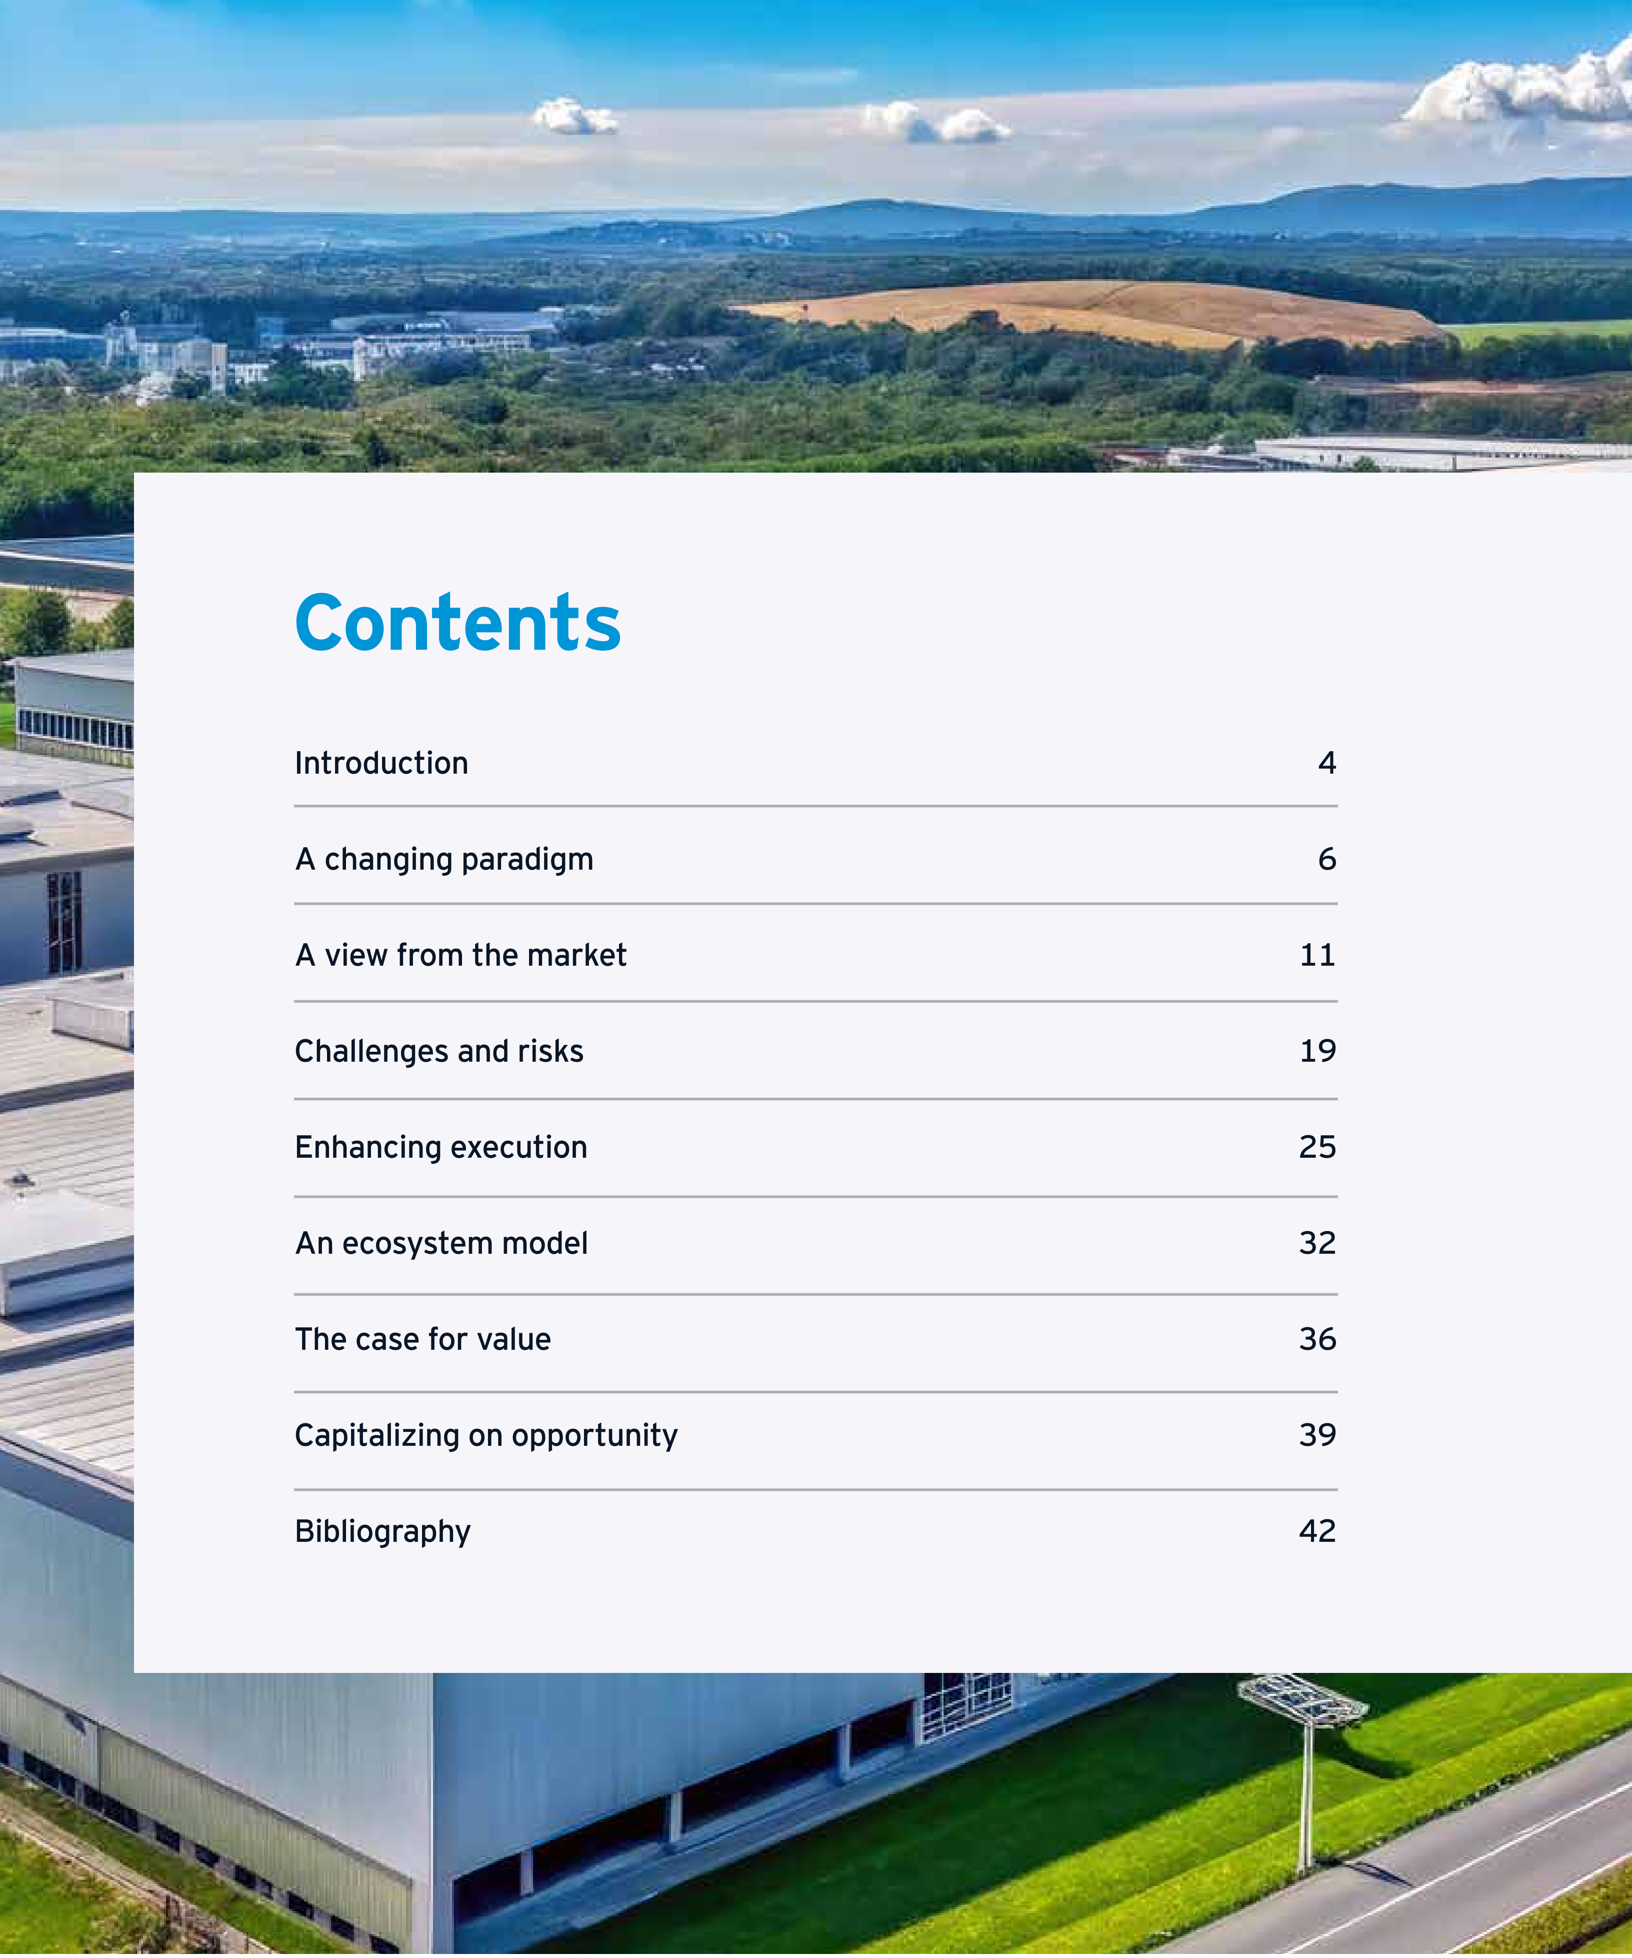

Contents

Acknowledgments

We thank the 18 utilities, data center sponsors and regional transmission organizations (RTOs) that were interviewed as part of the development of this document. Their perspectives provided valuable insights into the contemporaneous development and evolution of data centers and the emerging challenges related to future evolution of these projects.

The entities contributing to this point of view include:

- American Electric Power
- Arizona Public Service
- ComEd
- Digital Realty Trust
- Duke Energy
- Electric Reliability Council of Texas
- Entergy
- FirstEnergy
- Flexential
- Midcontinent Independent System Operator
- Montana-Dakota Utilities
- NiSource
- Northern Virginia Electric Cooperative
- Oklahoma Gas & Electric
- Oncor Electric Delivery Company
- PPL Corporation
- Southern Company
- WEC Energy

Their willingness to participate in direct interviews is greatly appreciated by EY-Parthenon teams.

The data center surge advances

Securing the promise of sustained growth

The electric sector has enjoyed past periods of high and transformative growth, but this pinnacle occurred in the 1950s and 1960s with the expansion of residential and commercial development, the emergence of more energy-intensive industrial businesses and the mass marketing of electricity-consuming products such as air conditioners, television sets and home appliances.

Since that period, electric demand gradually tapered off through 2000, when it flattened out or shrank leading into the COVID-19 pandemic. Demand levels subsequently regressed, taking several years to ultimately recover. Although dramatic change has enveloped the utility industry since the turn of this century, a meaningful demand growth trajectory (non-COVID recovery based) has not been a by-product — until now.

During the pandemic period, a major electricity demand trigger emerged from data center proliferation throughout much of the US. Following this emergence, a confluence of factors, such as a transition to advanced cloud and edge computing; fiber overlay; urban spread; 5G telecommunications; computer chip improvement; hyper- and mega-scalers; crypto mining; advanced analytics; electrification; and advanced manufacturing growth, have elevated both the focus on data computing power and overall demand growth. And now artificial intelligence (AI) training and inference are the dominating factor in forecast future energy consumption levels and growth. This continuing expansion of data center electricity demand has created intense market desire for "speed-to-power," i.e., rapid and on-time ramp-up of power supply to match market demand.

The need to serve this demand growth opportunity is creating new load, revenue and other challenges to traditional models, e.g., supply planning, transmission capacity, infrastructure build-out, application of regulatory tariffs (the approach to pricing and service conditions), and execution speed.

The projected 15% to 30% terawatt-hour (TWh) pace of annual growth in data center demand into 2030, and the scale over which demand is expanding, is unprecedented and creating an unexpected bonanza.[1] And the level of continuing annual growth in physical data centers (hundreds annually) and facility scale levels (megawatts (MW) to gigawatts (GW)) illustrates the level of future demand potential. (See Figure 1)

But this realized growth does not occur without risk incurred by developers (sponsors) and utilities (hosts) alike. Sponsors do not want to buy land and plan projects where they cannot strike an acceptable contract with a utility for power supply and infrastructure support. Sponsors also cannot be surprised about lags in speed-to-power timing from power supply constraints. Nor can hosts find themselves building a long-term asset lacking sufficient revenue for the life of the attributed power supply or transmission assets, i.e., creating a stranded asset. Similarly, hosts cannot create operating challenges — congestion, imbalances, concentration risk, etc. — for their system simply to accommodate a sponsor's preferred data center siting, ramp-up and scale.

Neither sponsors nor hosts are certain about the direction and longevity of data center expansion — will it be a shooting star with explosive energy for a short-term boom then bust, or a fundamental redefinition of sustained-and-growing sector power consumption that lasts several decades? Whatever the outcome, it is likely that increasingly larger data center projects will continue to be built and that this will accelerate over the rest of the decade for two reasons: First, the evolution of AI training is giving more computing power to sponsors; and second, the need to leverage AI's computing power is viewed as foundational to national security within the US and a source of competitive leadership, with data centers as a core component. These factors combine to solidify the near-term future of data center construction and the increasing value to be derived from massive AI inference.

The contours of this data center wave typically involve multiple parties related to infrastructure placement, specialized equipment installation and interconnection readiness, and multiple types of enabling participants, each with different objectives. In any single project, the involved parties may be sponsors, operators, financiers, original equipment manufacturers (OEMs), fiber providers, water companies, RTOs, construction contractors, municipalities, independent transmission owners, backup-capacity suppliers, hosts and state regulators.

From a physical perspective, the nature of data centers continues to evolve from smaller single-entity scale sites (less than 30 MW), multi-tenant co-locators (less than 50 MW) and larger crypto-based facilities (less than 100 MW) to hyper-scalers (more than 100MW) and mega-scaler generative AI locations (more than 1 GW). These mega-scale facilities are enabling projects of what the investment analyst community has coined the "Magnificent 7", a group of technology stocks that drove a significant portion of the markets returns since 2023.

While electricity demand from data centers has been substantially increasing, the utility industry has also been the beneficiary of higher demand from advanced manufacturing and due to supply chain reshoring and electrification. Together, it is estimated these factors will increase future power demand by 2.6% to 4.7% annually (and far more in the commercial customer class, where data centers reside). Demand growth was relatively flat after 2000 and is well below the projected growth.[2] (See Figure 1)

The future of data centers is likely to continue to morph from single projects to creating a network of participants that think about ecosystems — meaning, large multi-facility, multi-entity, multipurpose locations, intended to create and capture economies of scale and value for the principal data center sponsor, host and other participants.

Creating and preserving increased value-added contribution from either stand-alone or ecosystem models will be a fundamental future focus of sponsors and hosts. The expanding focus on data centers has already enabled an "investment unicorn" to emerge. How the future direction and stand-up of the data center phenomenon evolves will depend on how key participants shape their purpose, reimagine their objectives and collaboratively execute throughout the development and delivery cycles.

To provide additional insights to this study, interviews were conducted with a panel of 18 hosts, RTOs and sponsors who provided unique and direct insights to complement our research and analysis. These individuals shared their perspectives on a range of data center stand-up challenges and its future growth and sustainability.

Figure 1: Electric demand growth

Historical and forecast US annual power demand, 1950-2030E TWh

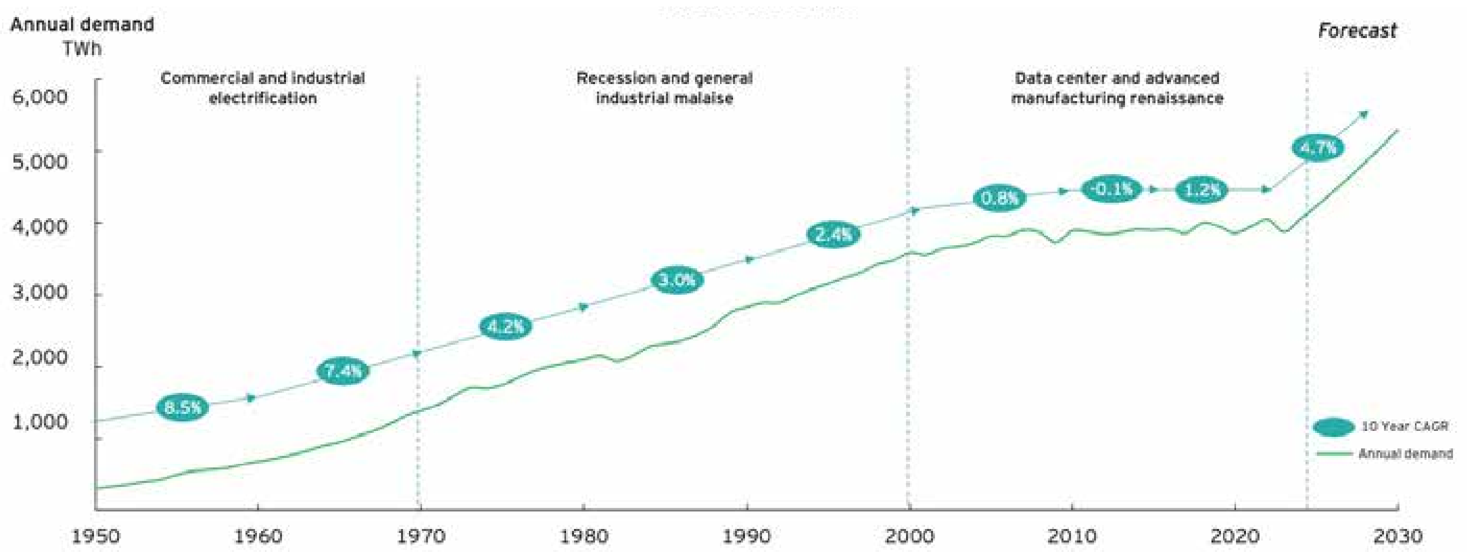

Source: Lawrence Berkeley National Laboratory; Grid Strategies; U.S. Energy Information Administration

A changing paradigm

Data centers have been a backbone of the digital information age since the proliferation of public internet access. However, as an asset class, data centers have evolved in purpose, scale and capability over the past 35 years, with the AI boom being the latest technology innovation facing data center sponsors and hosts and dramatically enhancing the value of computing. The continuous increase in system load has spurred hosts to frequently increase planned capital investment requirements well beyond integrated resource plan (IRP) levels developed within the same fiscal year, and often even when updated within the last two quarters.

This historic boom in data center demand brings significant questions, such as, "Who pays for incurred costs," "Where does new capacity come from" and "How can transmission be built quickly?" These questions are central to data center stand-up decisions but have not necessarily been fully addressed and answered for the future environment. Here, very large and expensive assets need to be contracted and constructed in a constrained power supply situation, where neither sponsors nor hosts know exactly how long they will be needed at expected levels.

As more sponsors expand facility computing requirements, and refine their use cases, and more industry catalysts emerge over time, data center proliferation will rapidly increase. Enterprise (on-premises) storage, colocation, cloud storage, cloud services (Infrastructure-as-a-Service (IaaS), Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), etc.), crypto currency mining and now AI training and inference actualization have led to a sharpening of data center purpose and advancement of facility capabilities such as high-density computing and power consumption use and efficiency. For example, growth in data center-related demand is fundamentally driven by increasing digitalization, bolstered by the broad adoption of cloud computing and network services. By the end of 2024, more than 85% of organizations adopted a "cloud-first" principle, firmly entrenching data centers at the heart of the digital economy.[3]

The observed increase in demand scale has significant implications for data center design, particularly regarding rack density and cooling systems to handle additional computational power requirements. These changes increase both construction and/or retrofit expenses and ongoing operational costs. The scale and intensity AI requires versus other common data center activities derives from how generative AI and large language models (LLMs) function. This impetus lies in how AI computing functions in three core categories: pre-processing, training and inference.[4]

- Pre-processing is the organization of the data set before the model initiates computations and is the least energy intensive of the three functions, as its function of cleaning, filtering and transforming data sets is a one-time event to prepare for training.
- Training is essentially teaching the model how to interpret data and often relies on building a complex neural network to process billions of parameters to predict (or infer) results, in much the same way as a human brain.
- Inference-driven demand derived from preparing for AI computational execution is just beginning and expanding, as both usage and intensity continue to grow from adoption and expansion. From deepfakes to movie visuals to auto-driving cars to predictive maintenance to medical diagnosis to financial and legal support, AI-use cases for both commercial and personal applications are just being adopted and integrated into day-to-day life.

Compared to traditional data centers, AI data centers consume vastly more power. ChatGPT queries are estimated to be six to 10 times more energy-intensive than web hosting or search queries.[5] This substantial computational power required for AI makes traditional data centers unsuitable for AI workloads. For example, several Mag 7 entities have rescoped a number of data center projects in 2022, and a representative noted that "supporting AI workloads at scale requires a different type of data center than those built to support our regular online services."[6,7] Additionally, data creation itself is estimated to expand at a 22% compound annual growth rate (CAGR) from 2024 to 2030, which has already caused private equity data center investment to run at a $100b clip.[8]

In addition, the rise in domestic advanced manufacturing is driven, in part, by the CHIPS and Science Act (CHIP), which since 2022 has provided approximately $53b toward semiconductor research and manufacturing, attracting nearly $400b in private investment. This historic level of investment is intended to enable the US to again become a leading global supplier of semiconductors aimed at producing approximately 30% of the world's semiconductors by 2032.[9]

Since 2020, US annual power demand has shown steady growth and progressed through pandemic stage recovery, with ongoing demand recovery expected to 2030. A similar outcome is expected to 2030 for annual peak demand growth as core growth drivers are expanding and data centers and advanced manufacturing are driving sustained new demand, albeit both at a 2% CAGR. (See Figure 2)

However, the overall magnitude of electric consumption increases projected creates a seachange impact on hosts, who will be required to secure capital investment, mitigate financial impacts and address rate cases. It also causes hosts to assess the implications of a sustained level of demand growth across future time periods, how long it can be sustained and how customers could be affected.

In 2025, data centers collectively are forecast to be responsible for about 8% of total US power demand; however, data center demand levels are forecast to increase to as much as 18% of total US power demand through 2030, driven by a confluence of cloud computing, AI and machine learning, and digital transformation across industries.

Figure 2: Forecast electric demand range

Total US annual power demand, 2020-2030E TWh | **Total US annual power demand , 2020-2030E GW**

Source: NERC; Lawrence Berkeley National Laboratory; Energy Information Administration; EY-Parthenon analysis

Forecast estimates of data center related demand are constantly changing among the investment banks, consulting firms and institutions — at increasing rates and levels — and will likely continue for the next few years as mega-scale generative AI-directed facilities increase in prevalence.

The tepid increase in US annual power demand prior to the pandemic period was consistently low because energy efficiency, global competition, tariffs, slow business development, industrial malaise and consumer sentiment constrained energy consumption. The post-pandemic period reversed these conditions but was about recovery, not really incremental growth. If the level of total annual demand growth, based on recent company estimates, continues at a roughly 3.5% to 4.5% pace on a sustained or greater basis, then US power demand levels will measurably and consistently be outpacing prior-year levels and further exacerbating needs for generation capacity.

Forecasts project that data center demand growth could be the major contributor to peak power demand, at 90 GW, by 2030, with an array of manufacturing expansion in automative, electronics, textiles, chemicals and pharmaceuticals to result in another 20G W of peak demand by 2030. Moreover, the growth in data centers and the semi-conductor industry is strongly interconnected, as advancements and breakthroughs in semiconductors reinforce this data center demand. Electrification of industries, businesses and transportation is expected to drive another 15 to 20 GW of peak demand.[10]
(See Figure 3)

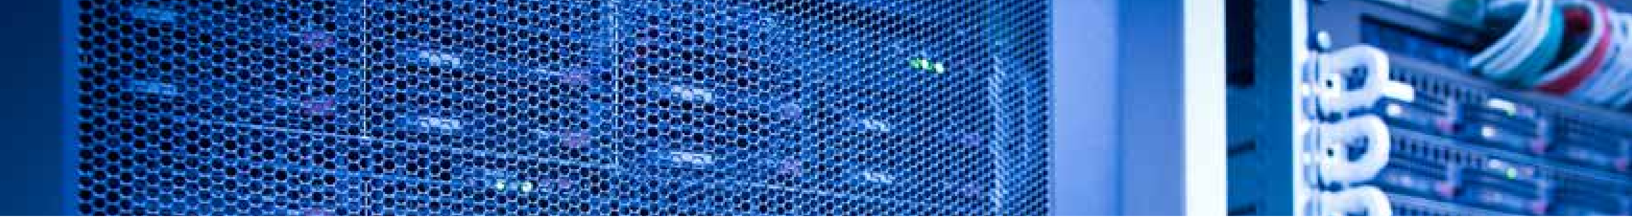

Figure 3: Demand contribution

US power demand growth CAGR, 2023[1]-2030%

US peak power demand growth forecast, 2030 GW

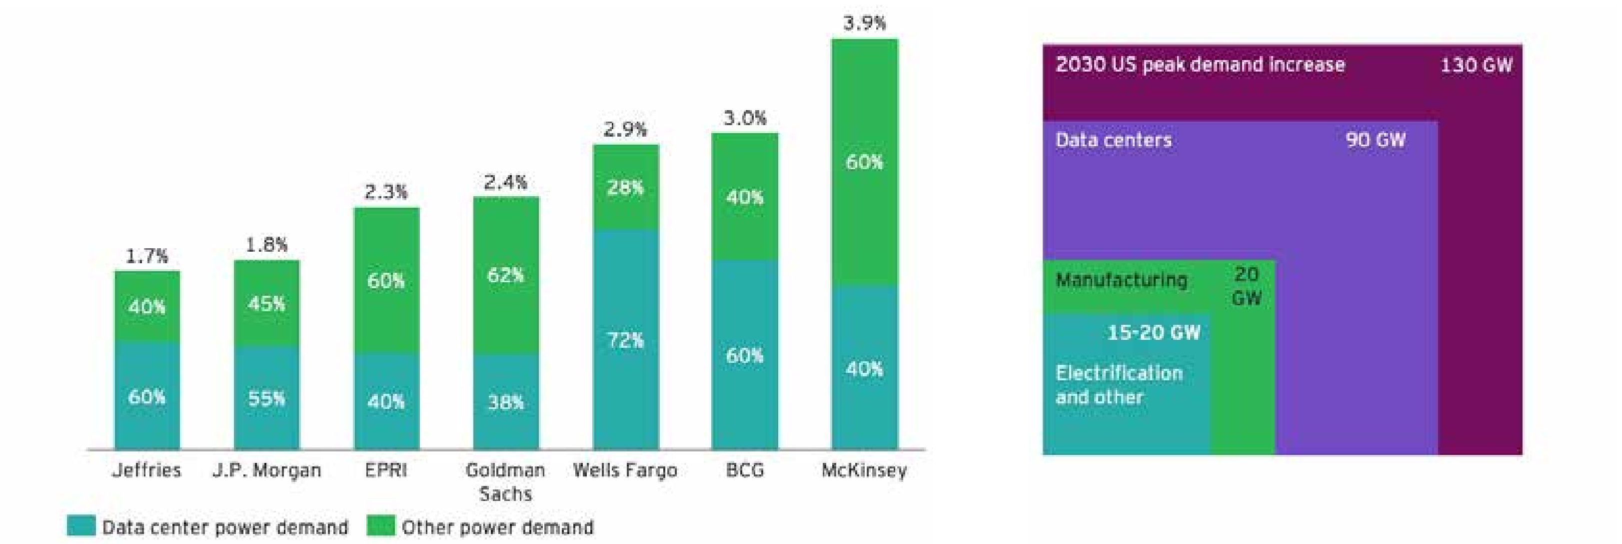

[11] For some analysts, the start year is 2024.

Source: BCG; Energy Information Administration; GridStrategies; Goldman Sachs; Jeffries; S&P CapitalIQ; Wells Fargo; EY-Parthenon analysis

At the US RTO level, these entities are currently inundated with data center requests for service, creating challenges to integrating large load requests into supply sources and transmission networks. For each of the past four years, the PJM Interconnection (PJM) summer 15-year peak load forecast CAGR has grown dramatically, from 0.3% in 2021 to 3.1% in 2025, resulting in an additional 25 GW by 2030. [11,12] Similarly, the Electric Reliability Council of Texas (ERCOT) interconnection queue has grown from 17 GW[13] in April 2022 to more than 52 GW[14] in March 2024, which represents nearly half of ERCOT's current total system load.

The impact of data centers on period-to-period peak demand growth is even more startling. This peak demand growth points out just how significantly annual load growth can affect sector requirements and outcomes. Highlighting the unpredictability and speed at which power demand can grow largely driven by data centers in many states is requiring hosts to constantly update their future capacity requirements.

The large variance in these predictions highlights the uncertainty of future data center power demand. This uncertainty is further complicated by the specter of generative AI consumption, which is just emerging and expected to be the biggest driver of data center demand in the US, estimated to comprise nearly 33% of all data center consumption in 2030, less than a decade after the 2022 introduction of ChatGPT. AI-dedicated power consumption is expected to accelerate at an even higher rate, approaching a 39% CAGR.[15]

The surge in new load requests from data center sponsors, the rapid expansion of AI and the continued use of traditional servers have also created uncertainties surrounding future data center growth, leading industry experts and analysts to predict a wide range of levels for data center load consumption and growth. One of the highest data center demand estimates, from Boston Consulting Group (BCG), projects data center demand growth of up to 491% from estimated actuals from 2024 to 2030.[16] A mid-range estimate from JP Morgan anticipates 223% growth through 2028, while even a conservative prediction by the Electric Power Research Institute (EPRI) forecasts US data center power demand growth could exceed 130% on the high end over the next six years.[17] (See Figure 4)

Figure 4: Future data center demand growth estimates through 2030

Data Center demand estimates, 2030E, GWs

J.P. Morgan's estimate is through 2028 only.

Source: company reports; EPRI; Grid Strategies, EY-Parthenon analysis

A surge in peak demand could exceed current power capacity levels in many areas across the country. This would require unprecedented investment in grid infrastructure to meet heightened levels of future demand already exacerbated by extreme weather events. Altogether, to meet the future power demand in 2030, RTO and utility planners have to increase planning and construction of current levels of new generation and transmission capacity several times over. Actual demand could be even higher due to the uncertainties surrounding the true scale of data center demand, as experienced by RTOs in the past few years.

Today, over 3,000 data centers exist in the US with legacy data centers (enterprise type) ranging from less than 5 MW to 10 MW. However, as a result of higher individual data center scaling over the past five years, average data center load capacity requirements have grown from 20 MW to 30 MW, to 80 MW to 100 MW and more, with emerging requests for service from data center sponsors now reaching and exceeding 1 GW.[18] (See Figure 5)

To put the size of a typical 100 MW data center into perspective, it consumes enough electricity to power more than 80,000 homes.[33] The dominance of hyper- and mega-scaler data centers will have a significant impact on the future size of the average data center. Proposals for recent mega-scale data centers exceeding 1 GW of capacity include facilities that can stretch almost a mile and include 4 million square feet in dimension.

Figures 5: Average data center scale change 1990 to 2025

Average data center power capabilities, 1990-2025+ MW

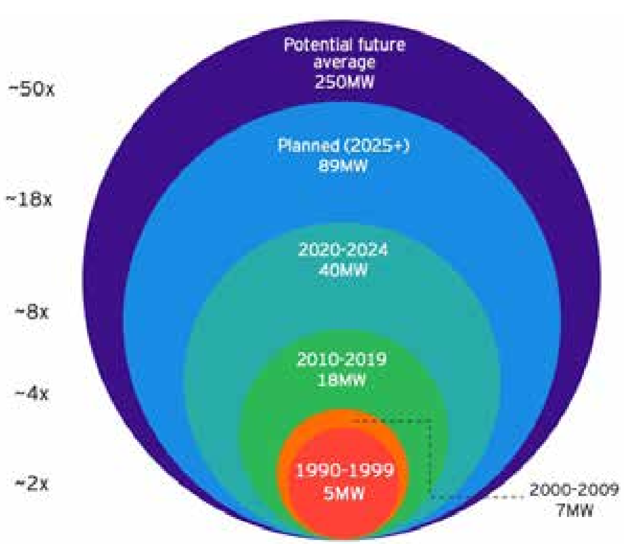

Source: EPRI, EY-Parthenon analysis

Markets including Northern Virginia, Dallas-Fort Worth, Chicago, Phoenix, Central Ohio and Atlanta have been the locations of choice for data center construction but have become saturated and increasingly less able to accommodate new requests. This has caused sponsors to select the next level of attractive markets (e.g., in Central Washington, Charlotte/Raleigh, Denver, Northern Indiana, Lehigh Valley, Louisiana, etc.) and smaller markets that may lack robust fiber infrastructure or skilled or even available labor. To illustrate absolute scale impacts in Virginia, which has become the largest data center market in the world, electricity demand from data centers increased by approximately 500% from 2013 to 2022.[19]

In newer markets, a 1 GW or more, mega-scale data center may pose fiber, water and power challenges. To cool servers, for instance, one estimate indicates roughly 2 million gallons of water per year could be needed to support cooling for an IT load averaging 1 MW per hour.[20] Going forward, other approaches – like liquid cooling – may supplant the need to access valuable local water resources.

Various providers, OEMs and start-ups across the data center participation spectrum are actively pursuing solutions to resolve these challenges through more energy-efficient algorithms, pruning and knowledge distillation or improving air-cooling technology to reduce reliance on water. However, the sector may already be approaching its next power efficiency horizon challenge, with declining facility Power Usage Effectiveness (PUE) and kWh per compute between 2016-2022, with modest improvement expected through 2030.

Further, the growth in request size causes several challenges for hosts. First, the ability of hosts to absorb increasingly larger load supply requests into existing capacity becomes further strained. While capacity availability still exists, larger and more data center requests are quickly absorbing existing reserves. Second, the increase in data center load scale adds concentration to a local geographic area and complexity to system configuration and balancing.

While the full ramifications of this new paradigm are incomplete, sponsors and hosts have been actively dealing with these issues over the last several years, with the path to power supply and transmission adequacy, cost and risk apportionment and data center stand-up and lifecycle production over an extended timeframe still uncertain. What is certain is that data centers are here to make a permanent mark on the electric grid infrastructure, the rapidly evolving digital economy and cityscapes across the country. Meeting the challenges and turning them into viable opportunities requires a collaborative effort between providers, communities, government agencies, sponsors, off-takers and hosts.

A view from the market

Just as impacts to energy and peak load have changed over time, data centers have also undergone significant transformation. From conventional localized enterprise models to the emergence of hyper-scale and mega-scale facilities (generative AI-focused) the nature and positioning of data centers have also shifted.

Five types of data centers — enterprise, cloud, colocation, edge and modular — are typically referenced with each designed for specific purposes. On-premises-based enterprise data centers are the most traditional customer segment, providing dedicated storage and application processing for a specific enterprise and are typically modest in size with as few as 10 racks or cabinets. Historically, these facilities have been sited on premises but advances in internet connection speed and security have made this previous requirement less common. However, enterprise facilities are now a shrinking segment of the market, with 80% of enterprises planning to shut down their data centers by 2025, due to construction and maintenance cost advantages of outsourced, tenant-based colocation.[21]

By 2027, enterprise data centers are expected to process about 25% of critical IT load worldwide, down from 40% in 2022 signaling moderate growth compared to other data center types. [22] Cloud data centers emerged as an off-premises solution to capture the cost advantage of scale and viability of cloud computing. Operated by a single cloud service provider, these data centers give both small and large enterprises the option for secure, flexible, cheaper and on-demand solutions for their data storage and application needs, eliminating the need to own and manage physical infrastructure.

Colocation data centers operate as a service for multiple tenants by leasing or partitioning their capacity into batches. A colocation data center provides the overall facility infrastructure, such as building, cooling, bandwidth, security, etc., while companies (tenants) are responsible for providing and managing their components, including servers and storage. Some colocation data centers also offer third-party IT services such as server and network management and are known as managed services data centers.

The colocation data center archetypes constitute three core subsegments: retail-enterprise, retail-network and wholesale and scale.

- The Retail-Enterprise segment involves leasing or selling data center space, power, cooling and network connectivity to smaller enterprise customers which do not have critical data/IT scale requirements needed to benefit from an in-house data center and are often less than 10 MW -15 MW. Customers of Retail-Enterprise providers are commonly looking for turn-key data center service and lower up-front costs.

- The Retail-Network segment primarily provides interconnection points for customers seeking network speed and optimization across their digital infrastructure. The customer base in this category is primarily in telecommunications, financial services, cloud and IT services and other network providers. Retail-Network facilities of up to 10 MW are similar in size to the Retail-Enterprise type but can be as large as five times that (50 MW), often in proximity, if possible, to a hyper-scaler campus

- Wholesale and Scale is the largest data center category, with capacity needs approaching 100 MW. Capacity is leased out in large batches to larger enterprise customers or cloud service providers. Customers of Wholesale and Scale data centers tend to be more sophisticated than their Retail-Enterprise counterparts, often managing their own connectivity and power needs. At the higher end, Wholesale and Scale facilities can grow to hyper-scale levels depending on their purpose. While there is no set definition for hyper-scale, these facilities tend to be massive, more than 1 million square feet and hold at least 5,000 servers, switches, racks and other networking and storage hardware.[23]

The introduction of edge and modular data centers further diversifies the data center landscape. Edge data centers are smaller facilities, strategically located closer to end-users, minimizing latency and supporting real-time data processing for the Internet of Things (IoT), autonomous vehicles, tele-medicine and other applications. Edge data centers will continue to grow as a segment as consumers and businesses demand less latency in their experiences (i.e., smart grid, on-line gaming and autonomous driving).

Modular data centers are prefabricated, portable systems with less than 2 MW, designed to be quickly deployed to address scale needs. They are ideal for temporary solutions or remote locations but have found a particular niche in addressing the extended timelines of full-scale data center development.

The fastest growing data center segment is hyper and mega-scalers, which provide dedicated support to the core business operations of large-scale cloud service providers. Post-pandemic, these hyper-scale data centers have grown to sizes exceeding 100 MW. While fit-for-purpose has always been a tenet to design and construction of data centers, hyper-scalers have variability in their specification requirements based on use case.

For instance, crypto mining data centers prioritize cheap power, while AI data centers require very high load reliability for their high-volume computational power and low latency. These differences initially drive placement decisions (power and land availability, development of reasonable tariffs for reliability, etc.), but also impact ongoing level of redundancy and likelihood of future scale increases for hosts.

AI mega-scale data centers present significant challenges related to capital investment, energy consumption and infrastructure requirements. These facilities are specifically designed to handle massive AI training models and inference applications, exceeding 1GW in load demand, which necessitates substantial capital investment of $10m to $12m per MW. [24,25]

Since the pandemic and with commercialization of LLMs, investments in mega-scale data centers by Magnificent 7 companies have surged. These investments are expected to accelerate in the latter half of this decade. For instance, Microsoft's Atlanta campus is expected to cost roughly $1.8b, and its Wisconsin facility approximately $3.3b to develop, while Alphabet plans to invest roughly $3.3b across two data centers in South Carolina.[26, 27]

Meta's Louisiana data center is expected to cost roughly $10b and require 4 million square feet.[28] Microsoft in January said it plans to invest $80b globally, half of it in the US, to develop data centers in fiscal 2025, which would eclipse its fiscal 2024 total capital expenditure (including non-data center investments). [29] Similarly, Amazon announced in 2024 that it plans to spend on building and operating data centers over the next 15 years.

Just as other demand-related estimates have continued to ratchet upwards, Magnificent 7 sponsor capital expenditures may also increase beyond expectations as more opportunities are identified and pursued. It should be noted that the capital expenditure levels noted above reflect only the committed project spend to be incurred pre-2030. The remainder for these announced projects will continue past 2030 as development and construction timelines are stretching out to 5-8 years when they involve new power supply and transmission construction. These amounts will also grow further as Magnificent 7 sponsors pursue nuclear generation as a power supply source to meet mega-scale project needs and provide a clean energy source.

By definition, 1 GW well exceeds a hyper-scaler footprint, which may require millions of square feet of land, resulting in zoning and community backlash. These facilities also require water for cooling systems to manage heat generated by AI processors and additional fiber. Further, ongoing investments in the latest hardware remain high compared to smaller colocation facilities to maintain computational efficiency in a rapidly evolving AI race.

Roughly 50% of data centers are in the aforementioned core markets in ten major data center hubs and are distinguished by their well-developed infrastructure, including robust fiber and water access, reliable power supply and access to transmission corridors. They also largely benefit from proximity to major business hubs, favorable business environments offering tax incentives and skilled workers.

Figure 6: Data center concentration

Data center locations, Totals numbers by state[1] as of March 20, 2025

[1] Includes operational, planned and under construction data centers

Source: Aterio; Data Center Map; Cloudscene

Virginia is home to one-sixth of US data centers, which consume more than 25% of the state's power.[30] Due to high concentration in these "core markets," the resulting strain on power capacity and some local pushback due to a range of concerns around aesthetics and fears of rising electricity costs, sponsors are increasingly expanding to "next tier" and "emerging" markets. These markets are becoming attractive due to lower costs, land availability for larger facilities and available power supply and rapid interconnection.
(See Figure 6)

Figure 7: Relative state data center load

Current and forecast load in top 15 states, 2023 vs. 2030E % of totalstate electricity consumption

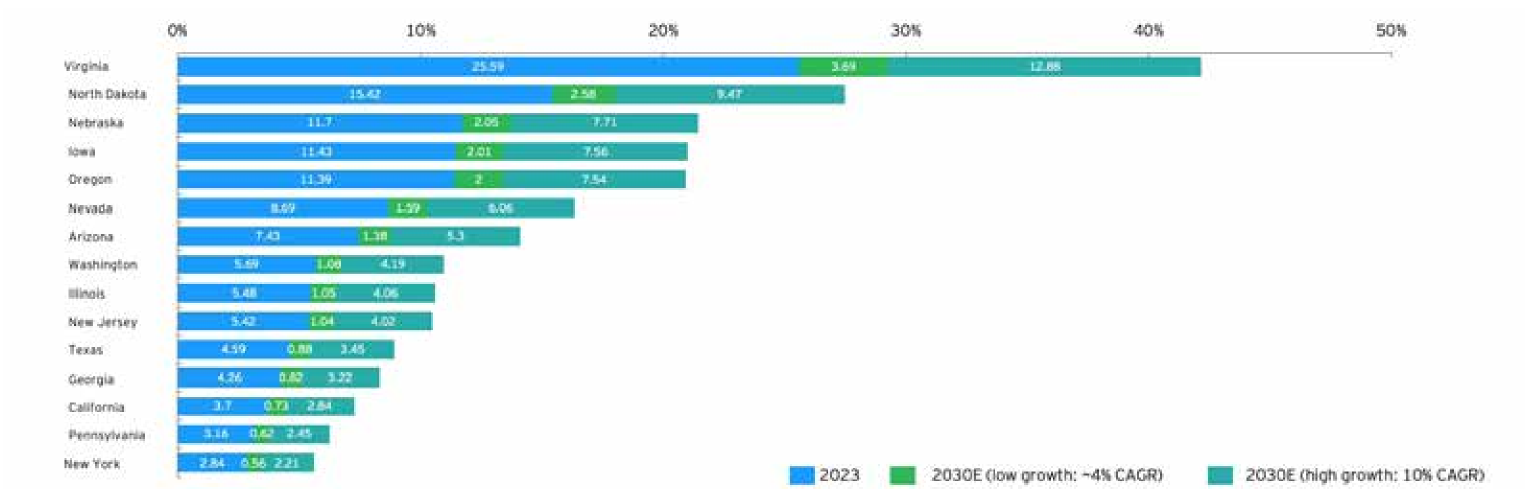

Source: EPRI

As demand for data centers intensifies, states are competing to attract sponsors by offering financial incentives for location, ranging from tax credits to waiving sales and use taxes on expensive equipment. Despite these efforts, demand is expected to remain high in primary markets. In 15 states where data centers comprise a higher share of total state electricity consumption, data center demand is projected to sustain its trajectory. According to EPRI, for example, by 2030 data centers could consume between 30%-45% of Virginia's power, 18%-31% in North Dakota and 15%-25% in Nebraska, Iowa and Oregon. Despite ongoing efforts by states and sponsors to diversify data center markets, certain clusters will continue to be major power consumers as they provide available and/or attractive infrastructure. (See Figure 7)

Despite this shift, CBRE data from the second half of 2024 indicates primary markets, such as Northern Virginia, Atlanta, Dallas, Phoenix and Chicago still lead new construction of colocation data centers by a factor of five (more than 5,800 MW under construction compared to roughly 1200 MW in next tier markets.) In terms of capacity, primary markets had about 11 GW of available inventory and under-construction capacity in H2 2024, compared to about 3 GW in secondary markets. (See Figure 8)

Figure 8: Existing and under construction data center capacity

Capacity in primary and secondary markets, H2 2024 MW

¹Primary markets: major data center hubs (Northern VA, Atlanta, Dallas, Phoenix, Chicago) are highlighted; CBRE includes Hillsboro, Silicon Valley, and NY Tristate
²Secondary markets: regions with a growing presence of data centers often with lower costs for real estate and power (Austin / San Antonio, Silicon Valley, Hillsboro, NY Tri State, Central Washington, Houston, Southern California, Charlotte / Raleigh, Denver and Minneapolis); numbers in the smaller markets are not displayed
Source: CBRE

As states actively compete to attract data center sponsors, seeking to capitalize on the influx of investments, the creation of construction jobs and the stimulation of the local economy, the most pressing challenge faced by sponsors is securing adequate power supply. Available power capacity is inadequate to meet needs in several states, as prior to data center proliferation, RTOs and hosts were anticipating flat to low demand growth for much of the decade.
Most planned generation was intended to replace coal capacity and support green energy commitments.

However, since early 2023, some hosts have seen 15% - 20% growth in commercial load demand and 10% -12% growth in data center demand alone. [31] To compound this increase in demand, new capacity has been slowed by RTO interconnection process timelines, putting pressure on both data center sponsors and hosts to find power supply and transmission from an increasingly constrained grid. This impact is just emerging as generative AI-focused hyper- and mega-scale data centers have only recently made requests for service and are still in front-end planning or early construction. Data from the most recent generation and transmission interconnection process timelines show a cumbersome and slow process compared to the speed required to respond to new data center demand.

By year end 2023, the average interconnection completion time from the request for service to agreement was generally between three and four years, while the length to commercial operation was generally between five and six years for most of the RTOs, while over eight years for CAISO. The quickest total timeline was ERCOT, but at over four years, even that is slow relative to prior typical data center capacity activation timelines. As a result of these backlogs, increase in interconnection requests and scaling of data center size in recent years, there are currently over 2,500 GW of active interconnection requests across the RTOs, twice the current total installed capacity in the US. (See Figure 9)

The RTOs, sponsors and hosts are working with the Federal Energy Regulatory Commission (FERC) to accelerate the interconnection timeline by addressing inefficiencies in every step of the process from interconnection request process structure to queue management, agreement, design and construction execution. For instance, PJM has implemented a cluster-based study approach to group projects together to identify grid upgrades and assign costs more efficiently and requires readiness deposits to discourage speculative projects. The new PJM process aims to reduce the entire interconnection timeline to just under two years, with MISO seeking an even faster process execution period. [32, 33]

In 2023, FERC transitioned from a "first-come, first-served" serial interconnection study process to a first-ready, first served cluster study process, as well as adding increased deposits and penalties for withdrawal to accelerate study completion timelines. [34] With the urgency of removing this extended interconnection backlog, the U.S. Department of Energy (DOE) announced in November 2024 that up to $30 million would be made available "to accelerate the interconnection process for new energy generation through the introduction of artificial intelligence techniques." [35] This funding is intended to develop partnerships between software developers, grid operators and energy project developers to modernize the interconnection application process.

Figure 9: RTO interconnection timelines

Estimated mean interconnection completion time by region, Through Year-end 2023 months

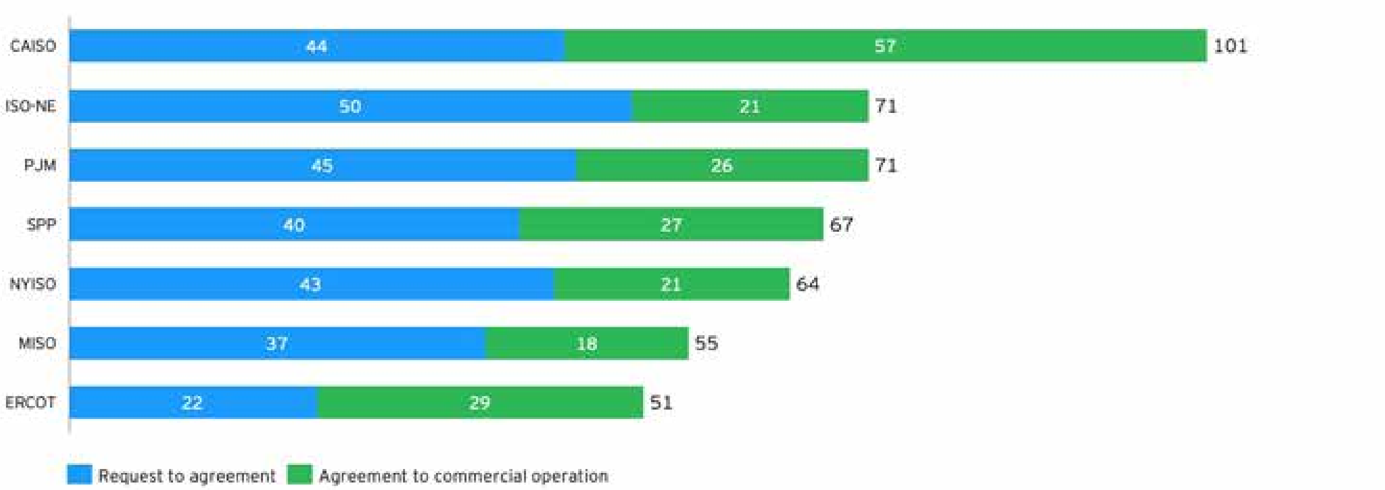

Source: Lawrence Berkeley National Laboratory; EY-Parthenon analysis

Even when power has been successfully secured for a site, data center stand-up is increasingly experiencing delays due to high demand. The growth in market demand has overwhelmed the value chain, causing supply chain congestion. As a result, supply chains are facing significant delays, with lead times for key data center components like transformers, non-specialized breakers and switchgear now extending to between two and three years.

Increased demand for semiconductors among downstream industries such as AI, 5G and edge computing is accelerating growth in chip demand by 29% by the end of 2026, double what semiconductor organizations anticipated, causing fears for another chip shortage post-COVID. [36] Even if power and related equipment can be sourced faster, labor availability is constrained to build these highly specialized facilities with nearly half of all technical data center workforce at or nearing retirement age. [37] All this culminates in execution challenges, extending project timelines —from one to three years pre-pandemic to five to eight years (or longer) today — adding complexity to a strained development environment.

These factors have increased the importance of hosts in the data center development process. Historically, hosts were not heavily involved in sponsor planning. However, the investment levels, scale, complexity and delays of hyper-scale data centers are drawing hosts closer to more engagement, as their planning processes are directly impacted by the scale and timing of critical facility decisions.

Sponsors are proactively seeking transmission study guidance for large load power requests is occurring, with hosts developing tailored and responsive strategy, governance and process changes to enable faster execution. However, broader and deeper collaboration is episodic between sponsors and hosts still in nascent stages. This is creating an opportunity for hosts to recognize how they can accelerate and enhance project planning while managing risks. Regardless, hosts are no longer just suppliers. They are now evolving to de facto partners so that projects meet their goals. [38]

When a data center is completed, sponsors need to implement their learnings, revise design features, enhance fundamental materials and leverage available technologies to maximize power usage effectiveness (PUE) – a key metric for measuring data center energy performance. A PUE value of 1.5 indicates that for every megawatt-hour (MWh) of energy consumed by server, process and computing equipment, an additional 0.5 MWh is used for non-IT purposes such as cooling, lighting and other overhead.[39]

Between 2006 and 2015, industry average PUE declined sharply from 2.5 to 1.6, and although improvements have plateaued since then, achieving lower PUE remains possible.[40] For instance, Google has achieved an average PUE of 1.1 across its data centers by employing efficient cooling systems, smart design, operational best practices and heat reuse.[41] This difference significantly impacts power use – a 100 MW data center with an industry average PUE consumes roughly 400 GWh more electricity annually than a Google data center of the same capacity. Key Magnificent 7 entities are pursuing methods to improve their operating efficiency and chase the efficient frontier. (See Figure 10)

Figure 10: Power usage effectiveness

Power Usage Effectiveness (PUE) over time, Best in class entity vs. Industry Average

Source: Company Data; Grid Strategies; EY-Parthenon analysis

Gains in computing efficiency don't necessarily result in permanent reduction in overall power demand due to Jevons Paradox. This theorem suggests that short-term resource efficiencies at the micro level ultimately lead to increased resource use at the macro level as cost reductions and design and operating enhancements are viewed as increasing the attractiveness of consumption. Jevons Paradox thus suggests that as computing efficiency translates into lower operating costs, additional and expanded use of technology becomes more affordable and attractive for more entities, and can increase the overall level of utilization, and therefore, consumption. [42] (See Figure 11)

Figure 11: Jevons Paradox

Compute efficiency and data center power demand

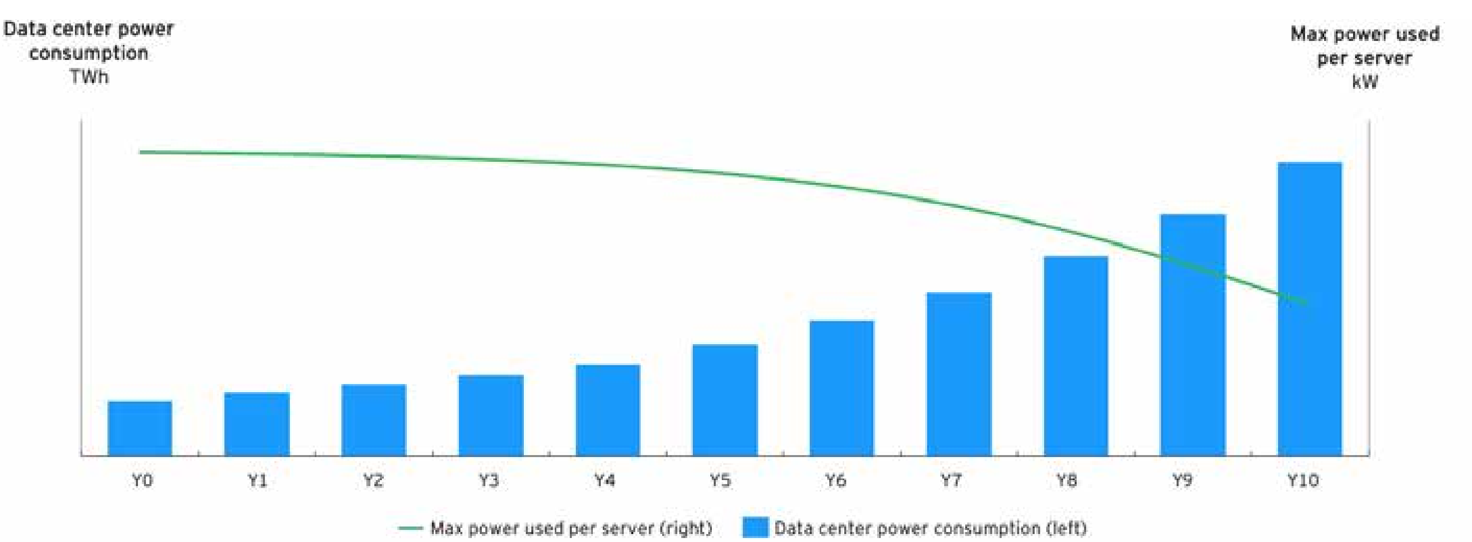

Source: SemiAnalysis; EY-Parthenon analysis

In early 2025, a public disclosure by DeepSeek (a startup based in China) disclosed a design and technology breakthrough that caught the global AI, data center, power and chip sectors by surprise, leading to a market correction in some related US technology, utility and energy investments. DeepSeek's breakthrough reflected cheaper and faster chips and simpler AI training and promised the delivery of less energy intensive and lower costs – Jevons Paradox in action.[43] Many observers were quick to herald the end of the data center boom and a return to less energy intensity, lower capital expenditures and more commoditization of technology.

Within a couple of days, Mag 7 entities, investment analysts and technology observers shared their insights and calmed the stock market, along with technology vendors, equipment manufacturers, power providers and gas suppliers. In essence, the event resulted in what Jevons Paradox points out: technology enhancements will continuously occur, and costs will improve over time, but computing power and energy intensity lead back to the next step-up in demand. It also went largely unnoticed that the chip and AI sectors were already working on similar technology enhancements that would have driven efficiency and costs in a similar manner.

For example, on-site generation with short and simple distribution paths is a potential method to improve efficiency by allowing the capture and use of waste heat. Another developing tool is utilizing a dynamic response system to computing demand, which is characterized by periods of peak demand followed by periods of idle time that can draw more than 50% of full computing power. A dynamic demand response system can realize huge energy consumption savings by powering down idle servers. [44]

Several design, material and equipment solutions are available to further lower PUE values. Direct-to-chip liquid cooling, which circulates a liquid coolant directly over computer chips, is more efficient than traditional air cooling and can reduce rack power consumption by 25%.[45] Microsoft takes this one step further by using a closed-loop liquid cooling system that consumes zero water, saving 125 million liters annually per data center.[46] Another solution is to use higher-voltage power because more powerful chips now can accommodate a higher voltage, which can reduce energy loss by up to 30% [47] Smart designs that simplify power distribution systems can shorten power paths, allow higher voltages closer to racks and reduce the number of failure points by 20%.[48] One Mag 7 entity leverages data and AI to design efficient pathways and rack positioning, minimizing stranded power.[49]

For AI data centers, improving power efficiency for both training and inference workloads is crucial. Limiting GPU power draw can reduce energy consumption by 12% -15%, albeit with a 3% slowdown in training. Another solution is to deploy algorithms that can identify and stop underperforming models early, reducing energy use by up to 80%. For inference, using a mix of GPUs and CPUs based on task complexity can decrease energy use by 10-20% without compromising on response time.[50]

In addition to these actions, an array of other solutions exists, depending on the characteristics and location of a data center, and can bring additional savings. For example, an air economizer, which takes advantage of outside air (e.g., seasonal and nighttime temperature variations) to cool space is a low-cost equipment known to cut data center cooling costs by over 60%.[51] (See Figure 12)

Figure 12: PUE opportunities

Efficiency solutions for AI workloads | % potential energy reduction

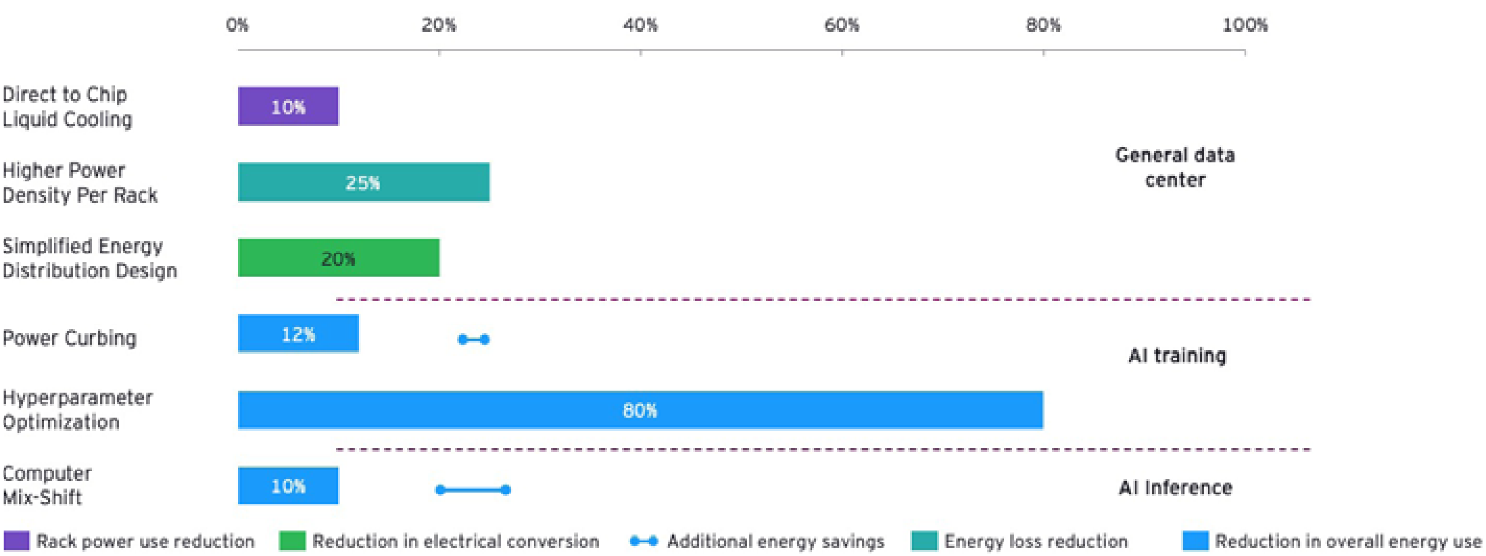

Source: MIT; EY-Parthenon analysis

Advances in power computing and chip technology will continue to lower power consumption per computing unit, reducing unit and data center costs. For example, NVIDIA's DGX B200 server, despite consuming twice the power of the DGX A100, processes 15 times more data per second, enabling data centers to achieve more with less. All these materials, technology, process and computing factors can reduce energy consumption. The question is, can they be sustained?

Consideration of potential gamechangers, for example quantum computing, creates additional uncertainty regarding impacts to data center energy consumption. While little is known to date, reduced demand could follow from more efficient production, but long-term demand could increase as technology proliferates and use cases expand. The question, "Will quantum computing decrease energy consumption because it is so fast and powerful?" is often rephrased to, "Will quantum computing increase energy consumption precisely because it creates more inquiries and demands?" Although some recent power demand requests may filter out as AI hype diminishes, the importance of AI to future economic development suggests that overall energy demand and consumption trends may remain elevated (and could potentially grow).

The jury is out about long-term data center demand as there are many questions where the answers are simply not determinative. Through 2030, the current growth trajectory in projects under contract and those soon to be contracted suggest a very robust level of future expansion, particularly as data centers grow into GW levels as generative AI progresses. Extending this pace of growth beyond 2035 and into the 2040s is far less certain and gives all sponsors and hosts pause about demand sustainability and investment recovery.

Challenges and risks

Each succeeding hyper-and mega-scaler data center proposed has its own uniqueness, with increasing levels of scale, investment and complexity for sponsors and hosts. These hyper and mega-scaler data center challenges start at project conception and extend through actual energization and activation.

For sponsors, challenges start with answering a range of questions important to any facility being proposed:

- What is the purpose of facility stand-up?
- What is the optimal location of the data center site?
- What is the scale expected for the proposed facility?
- What is the ramp-up rate for build-out?
- Where are the future load curve requirements?
- How do these predicates align with available power supply capacity and transmission adequacy?

Speed-to-power is the desired outcome from these front-end considerations, particularly when potential delays may arise, e.g., supply chain delays in the utility sector and potential timing impacts for interconnection to the grid. These insights lead to framing the subsequent electric service agreement, construction agreement and contract commitment (terms and conditions) with the host, including pricing arrangements, operating durations and levels and other formal pricing considerations.

Sponsors then need to consider the range of market and operating risks that could emerge and how to address them with the host, particularly with respect to gaps in facility lifecycle between contract duration for the data center and the anticipated life of newly built operating assets. Finally, sponsors need to understand where fundamental facility value can be enhanced, how to collaborate with hosts to pursue available sources of value and where the potential is to create an expanded cluster configuration. This can enable a broader model that views the data center site and ancillary commercial capacity as value sources. While some of these challenges are simply part of business execution, others, like contract structure, power supply-and-transmission capability, interconnection and project-activation timing, are more critical to project formation and life cycle outcomes.

Hosts deal with the reverse of many of the same challenges that emerge from first contact with potential sponsors and extend over these same end-to-end project stages and can extend over the life of the initial and succeeding contracts. Foundational considerations of hosts start with understanding the completeness of the sponsor's ask, and the financial capacity and legitimacy of the underlying sponsor's objectives. The nature of the sponsor matters as fundamental factors like commitment and capacity are at the heart of any request for electric service process.
(See Figure 13)

Figure 13: Data center stand-up challenges

Front-end considerations

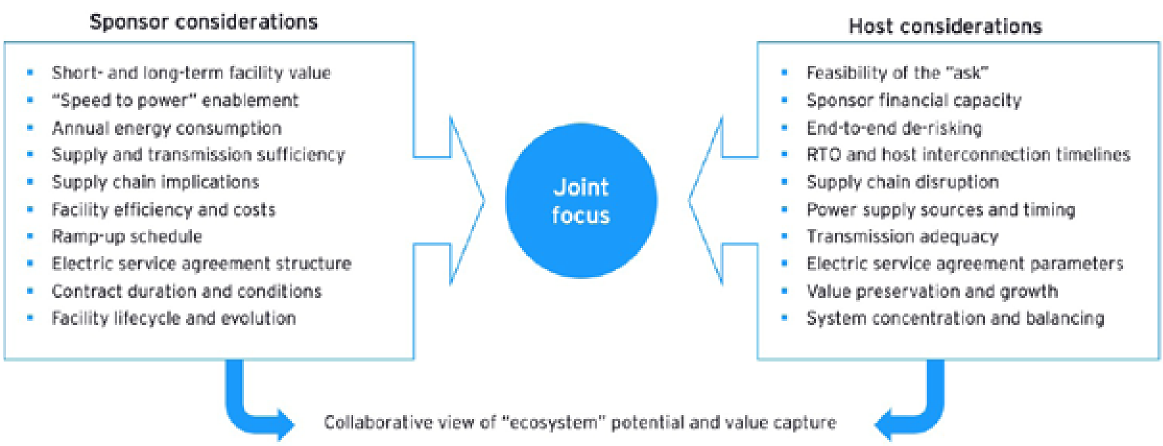

Hosts need to fully understand and thoroughly evaluate the externalities of every request for service received, e.g., RTO interconnection queue time frames, specific supply chain availability or lead times for principal equipment, as well as power supply and transmission build requirements where the host has responsibility.

With these factors articulated, hosts can then focus on refining the electric service request, framing the formal contracts or other agreements and developing appropriate contract pricing, terms and conditions to guide the agreement over a 10-15-year (or longer) period, if available. These activities then provide inputs to refine the nature of project risks and how to de-risk projects to mitigate potential financial exposure to the host or potential affordability impacts to customers. Hosts are particularly focused on understanding if data center siting and scaling have potential unforeseen operating impacts, such as excessive asset concentration or system imbalance within the host's physical footprint that can create longer-term grid challenges.

The challenges hosts face are not limited to interfaces with the sponsor. They also relate to how internal processes are framed and executed to drive effective project management and performance across discrete or parallel projects to enable ecosystems and value creation.

Both participants need to think about how to create joint wins that help the sponsor and host through the duration of the contract and the life of the asset.

Some of these challenges naturally extend beyond simple project planning and execution problems – they signal tangible financial or performance risks that need to be considered and mitigated by and between sponsors and hosts from project inception through activation and until (and perhaps after) facility retirement. It is not in the short- or long-term interests of a sponsor or host to allow these risks to elevate into "showstoppers" that adversely project timing, delivery, or outcomes.

From a host utility perspective, designing and embedding an effective de-risking process is fundamental to avoiding potential exposures and realizing successful data center project outcomes. Hosts are becoming adept with data center project risk identification, assessment, mitigation and management. But risks in project planning, financing and execution are neither dissipating nor fully avoidable – they're simply part of a multi-party process. And as sponsors and hosts are realizing, perspectives and requirements are not always completely aligned.

The list of risk areas that sponsors and hosts typically encounter establishes a working framework for assessment and incorporation into project planning, evaluation, management and contracting. These risks relate to the full development, execution and activation timeframe of the proposed data center. They cover a range of stage-gated risks like feasibility of load curves and ramp-up, adequacy of power supply and transmission, financial protection of customers and performance execution. This list does not include sponsor and host behind-the-meter cost allocation risk from colocation at nuclear plants, where risk (and cost) responsibilities remain unsettled in early 2025.
(See Figure 14)

Figure 14: Risk elements

Data center risk dimensions

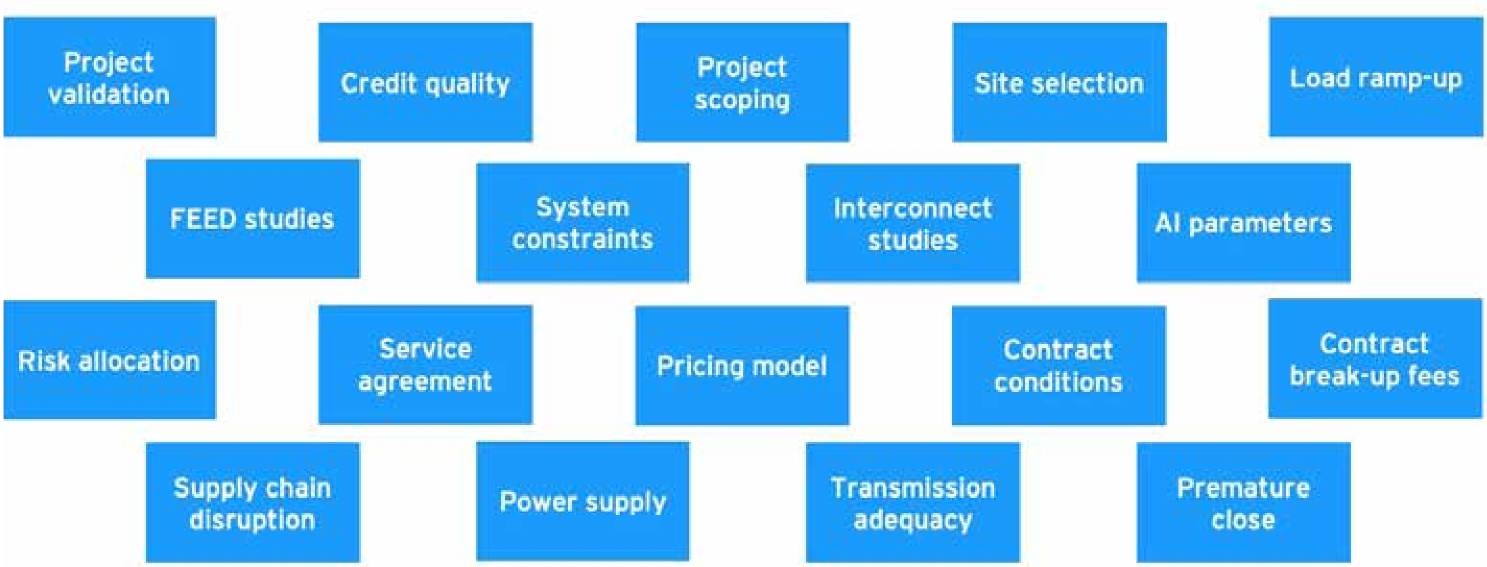

The challenge to sponsors and hosts is how to position to identify potential project-specific risks and impacts, as well as those germane to data center projects regardless of location. First, the way these risks unfold is particularly relevant – a number directly affect project definition and formation and effectively start when the first sponsor contact occurs. Second, as more project specifics are considered, new risks can emerge. Two examples are power-supply sizing and sourcing, and transmission-line access or construction. Sponsors and hosts need to prepare for risks that occur early in decision-making, as well as those continuing through the data center project lifecycle until project sites are retired or abandoned.

To date, hosts have generally adopted both generic (cost responsibility) and tailored (project ramp-up commitments) responses to how risks can be mitigated or avoided. But these measures need to continually evolve from broad fit-for-purpose circumstances identified in preceding data center project evaluation, contracting and execution processes, to tailored solutions that meet future unique project requirements.

Hosts and sponsors have focused on shared areas of potential risk occurring from first contact by sponsor or request for service, continuing through assessment stages, planning and contracting, and extending through the data center lifecycle. These back-end risk outcomes have yet to really materialize to date given the smaller scale of legacy data centers in terms of complexity and the short-elapsed time between larger data center activation and operational run-time. (See Figure 15)

Figure 15: Risk likelihood and impacts

Positioning on the risk grid

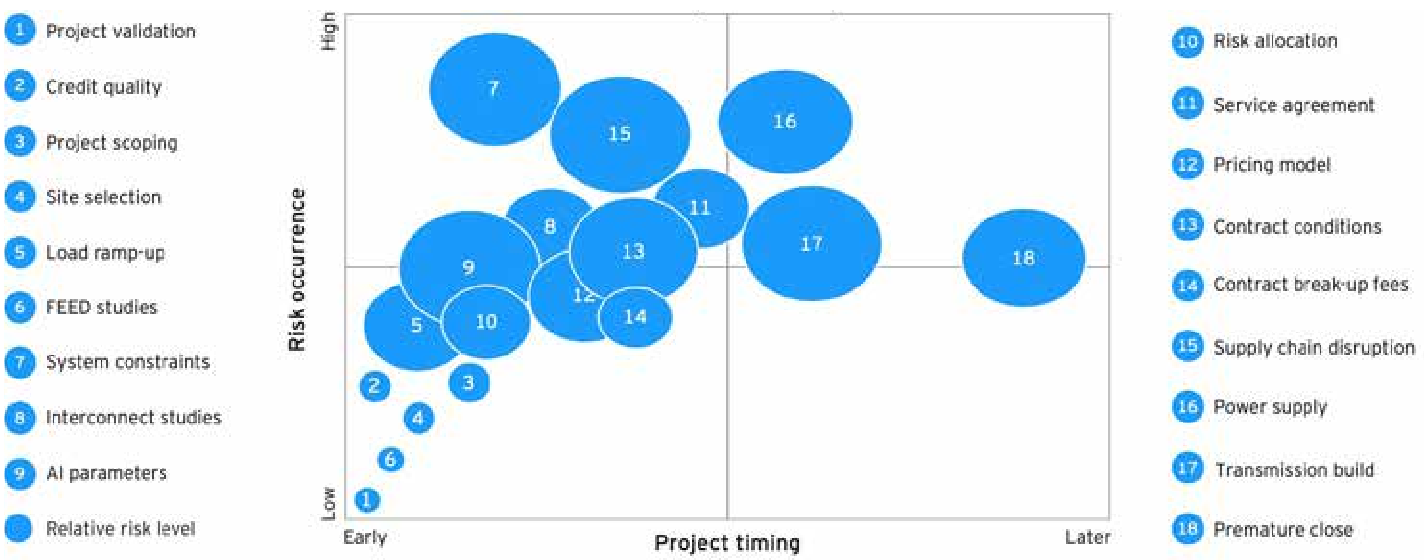

As a matter of philosophy, hosts' focus on avoiding gratuitous project assessment services without payment for related, incurred or out-of-pocket costs for front-end engineering design (FEED) or interconnection studies – which mirrors state regulatory policies governing how non-benefitting customers should pay for benefits derived by another specific beneficiary, like a data center sponsor.

As the project evaluation stage moves to the next level of project planning, new requirements emerge focusing on fundamental challenges related to sourcing power supply. For data centers located in traditional regulated jurisdictions, hosts first assess whether existing generating capacity can serve this new load level, as well as planned service territory demand growth, or whether new supply is required to be built to accommodate a broader scaling of generation for both current and future customers.

If no new power supply is necessary, the challenge level and related risks to generation adequacy is substantially reduced. However, if new power supply were required in a regulated state, the host is expected to provide it and may need to provide back-up supply necessary to support data center needs for constant power. However, sponsors may be able to "bring their own generation" (arrange construction of a supply source, e.g., a reciprocating internal combustion engine (RICE), renewables), particularly for back-up supply or clean energy sourcing, in certain jurisdictions.

The range of new generation construction risks that a host undertakes on behalf of the data center sponsor, and existing and future utility customers, primarily relate to supply chain delivery, construction execution and project delivery performance. These risks include new power supply project planning (including state regulatory approval to build), project management, equipment availability, craft labor adequacy and on-time/at-cost plant delivery. (See Figure 16)

Figure 16: Value chain de-risking

Data center value chain elements

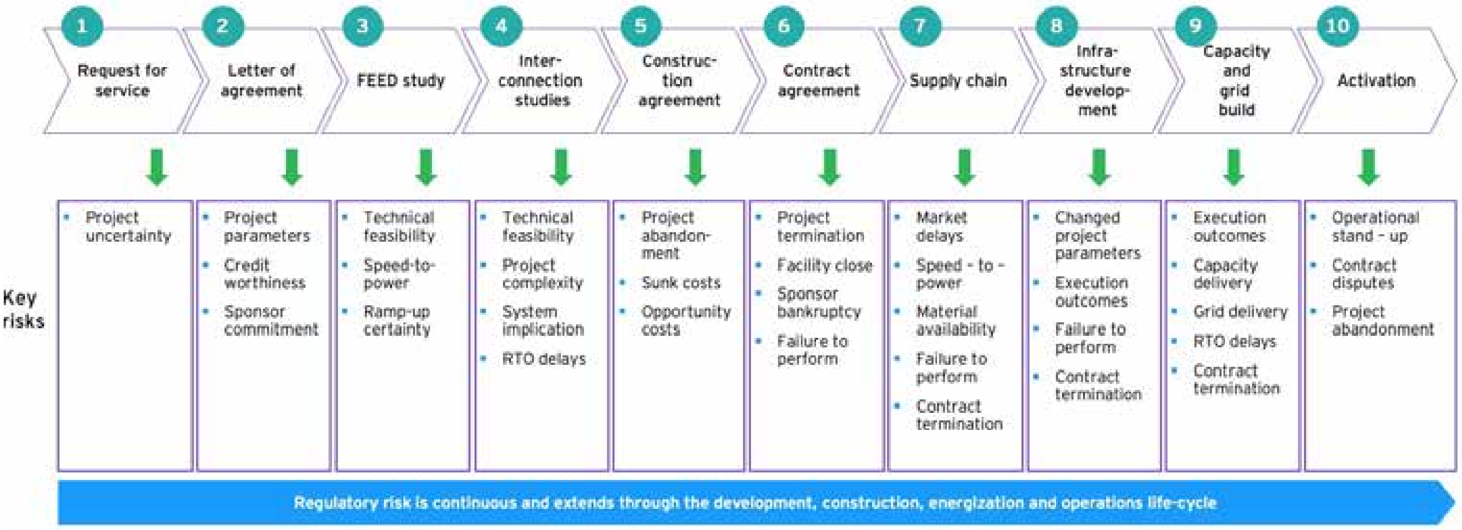

Tomorrow's challenges are being identified every day. Sponsor-and-host experiences, market circumstances, project complexities and contract evolution elevate the range of risks and mitigation approaches both sponsors hosts need to protect their rights during early engagement and formal arrangements.

Mitigation of significant risks involves far more than just delivery of a generation plant itself. Delivery in accordance with contractual commitments is foremost and carries much uncertainty. This requires sufficient contract flexibility, performance allowances and host-related damage limitations.
(See Figure 17)

Figure 17: De-risking the data center experience

Range of protections

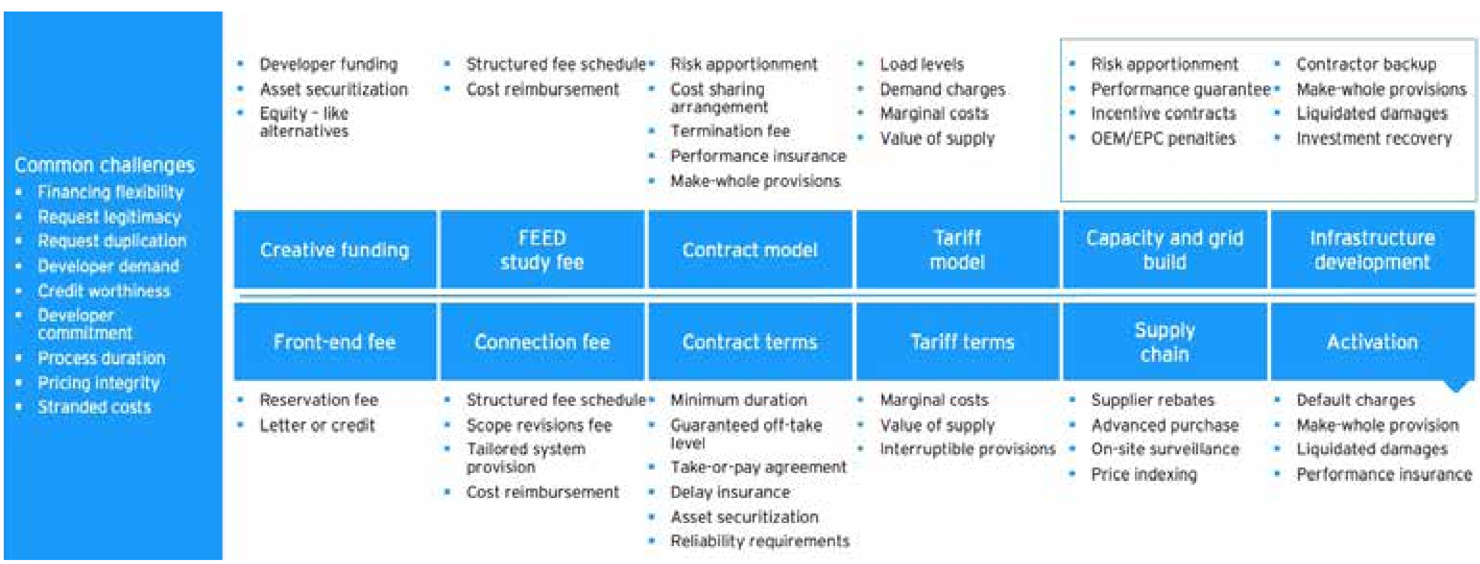

The situation for new transmission capacity in a regulated state is similar to power supply. If sufficient transmission capacity means no new build is required, the need is much simpler for the utility to fulfill. If transmission capacity is not sufficient, the host can seek to plan and construct a new line or expand or upgrade the existing grid. However, just like for new generation, state regulators will have to approve planning and construction of additional transmission capacity, which also carries delay risks as described for power supply. Regulatory review will occur through Integrated Resource Plans (IRPs), which may only occur infrequently (e.g., every three years.) Consequently, hosts need to be prepared to file updates to notify regulators about significant changes to power supply and transmission capital investment for ratcheting data center demand requirements.

This situation will become substantially more complex and elongated if the data center is to be developed, constructed and activated in a competitive market state, i.e., where RTOs exist, which oversee new generation additions and transmission expansion or build. In these situations, formal processes for new generation interconnection are used to screen and review applications for power supply and they are evaluated by RTOs to understand impacts on reliability of the existing transmission system. These RTO processes can involve hundreds of discrete applications related to GWs of new capacity and can take two or more years to move through a large queue of GWs of proposed new generation.

Similarly, new transmission approvals can take 2-3 years (or longer) to obtain necessary RTO approvals and involve dozens of non-data center project types, e.g., reliability, interconnection, upgrades, economic and congestion and require multi-billion-dollar investment levels extending over multiple years between approval and actual energization.

The power supply and transmission addition processes above highlight how sensitive data center approval, permitting, and construction processes and timelines can be to external factors beyond sponsor or host control. When typical supply chain delays (which can be 24-36 months) are combined with new generation and transmission reviews and approvals (which can also take 24-36 months) it is easy to see how data center projects can take two to three times longer than in earlier years to move from initial request for service to finalization. (See Figure 18)

Figure 18: Stand-up timelines
Stages, sequence and timing

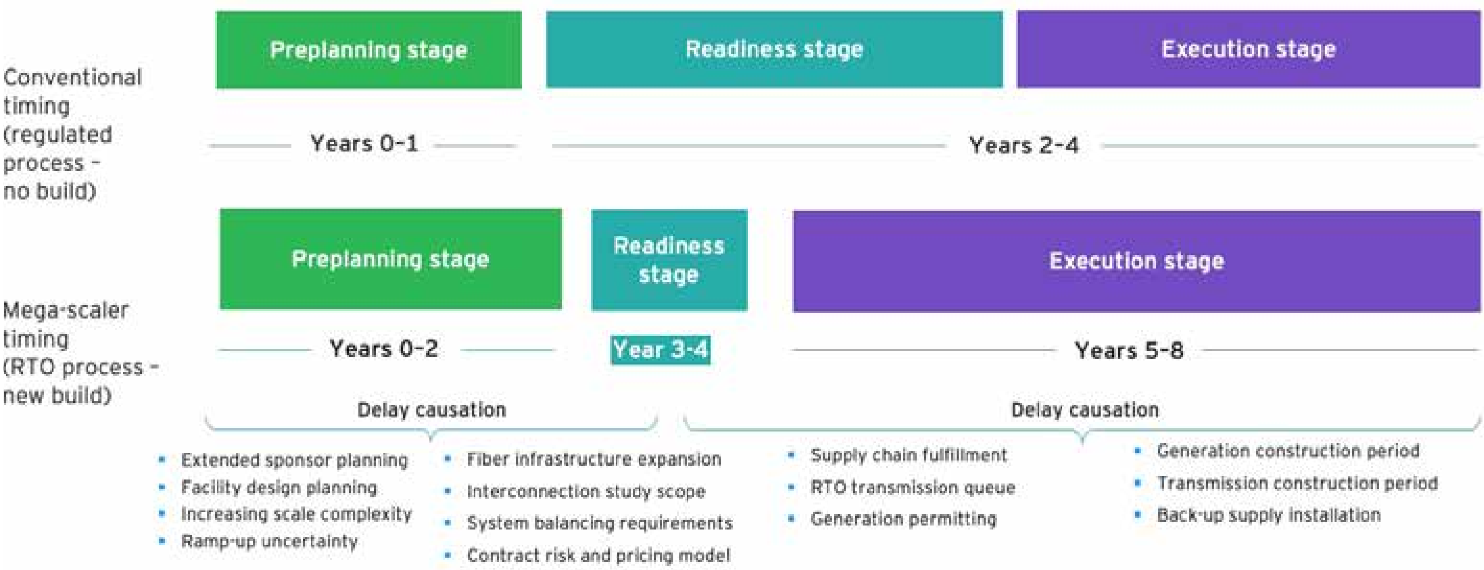

Source: EY-Parthenon analysis

Perhaps the most significant risk area for sponsors and hosts is in overall project contracting: contract structure, terms, conditions and protections, and the pricing models to be applied. This activity takes place in the early years of data center development, e.g., before year three in the mega-scaler example in the above figure. While plenty of critical issues related to data center stand-up exist, the most important factor for sponsors is speed to power, which makes other related issues – like costs and tariffs – secondary to achieving rapid activation. In a sense, the needs of the sponsor to rapidly fulfill its business purpose and meet burgeoning computing power uses is the foundation for success.

Fortunately, in numerous state jurisdictions, formalized large load customer (LLC) tariffs and complementary incentives – like sales and use tax exemptions or extensions, or allowances tied to future investment levels, already exist and provide a working model for data center facility off-take pricing. These frameworks provide a basis that can be acceptable to sponsors and hosts alike and they are consistent with the objectives of state regulators, who want to avoid customers being saddled with costs where they were not responsible for incurrence.

But absolute price levels still matter, as project-specific contract terms and conditions are portable across projects. Differences exist between sponsors and hosts regarding several important issues – project parameters, alignment, assignment and/or apportionment of project risks, build-out and delivery performance, market price elasticity, project lifecycle, take-or-pay conditions, abandonment or retirement implications, facility component costs, site asset requirements, market demand and customer affordability to name a few. All these elements need to be factored into the upfront contracting process and considered by sponsors and hosts over the expected life of the contract.

These issues become particularly acute when the life of the data center sponsor contract is substantially shorter than the life of assets added by the host and its investment may not be fully recovered. State regulators take a dim view of creating stranded costs, and the expectation for recovery by hosts is inconsistent with local customer affordability issues.

This potential outcome needs to be thoroughly considered during contract negotiations. Remedies, like take-or-pay provisions, minimum demand charges, value pricing, sponsor performance conditions, asset transfer and even securitization need to be evaluated for application. Effectively addressing the risks associated with data center development, construction and activation is an area where both sponsors and hosts need to collaborate on drivers, consequences, options and solutions.

To date, state regulators and legislators have been relatively hands-off regarding contract design, pricing models and risk allocation; and content with application of existing LLC tariffs and/or provision of other incentives. However, neither sponsors nor hosts want to create events that cause unnecessary disruption to the current regulatory environment and expansion of the level of state regulator involvement at the contract fulfillment or data center facility lifecycle levels.

Both parties would benefit from thoughtfully and equitably crafting a pricing model that recognizes relative roles, embedded and emergent risks, market shifts, asset life, cost recovery implications, customer impacts and benefits and adaptable risk mitigation options and consequences.

Enhancing execution

The identified challenges and risks above surrounding data centers are many, complex and often long-lived – but they are not insurmountable. Many of these individual challenges and risks can be remediated if considered as potential opportunities, rather than only as problems. And they can be sources of tangible benefit to hosts if approached from the perspective of enhancing "how" proposed data center facilities are addressed from inception through activation and continued sponsor operations.

Hosts need to adopt a different planning and execution mindset from prior circumstances, moving from only emphasizing discrete milestones and sponsor needs to positioning more sophisticated structural, process and engagement solutions across the entire data center development, construction, and activation process. These changes can dramatically shift the role of the utility from a passive provider to an active partner with sponsors over the course of engagement.

Three cornerstones of an evolved project discipline need to be adopted– for firmer execution footing, increased visibility and rigor of performance delivery and an expanded relationship with sponsors:

1. Moving from a project to a portfolio approach
2. Embedding a formal project management model
3. Creating a collaborative mindset with sponsors that transcends data center stand-up interactions

Expanding the thinking on how to approach host stewardship, readiness and execution will serve to enhance project economics, mitigate performance risks, expand relationships and create sources of future value to be captured.

The need to upgrade internal planning and management processes reflects the dramatic advance in data center development, construction and stand-up levels discussed earlier and evidenced among the companies interviewed which create unprecedented load demand in the market. (See Figure 19)

Figure 19: Data center growth curve profile

Incremental GW installed growth to 2030[1]

Note: Compiled or calculated from internal and external data

Source: Host Interviews; Various Analyst reports: Host Q4 Earnings presentations; EY-Parthenon analysis

To emphasize the importance of hosts elevating internal discipline, the panel interview group of hosts indicates they expect to add more than 170 GWs in service territory load from data centers between 2024 and 2030 – in some cases, an increase of two or more times existing peak load within their systems. This heightens the potential risks hosts may encounter in a sustained construction environment at scale.

It's important to recognize the nature of the curve represented in Figure 20 reflects four influencing factors:

- First, LLCs do not receive requests for service or add load in a uniform manner across projects. (It's always open season.)
- Second, how the hosts address FEED and interconnection studies and project future facility scale requirements can vary in service territories and from factors like construction and interconnection schedules and purpose and scale differences.
- Third, for the largest mega-scaler projects, ramp-up rates from sponsors also vary as project complexities emerge at different junctures.
- And fourth, to be conservative, the peer group uses individual probability estimates for converting requests for service into projects until they are confident a project will be initiated.

This last factor causes the most significant impacts to future GWs to be added, as uncertainty over future contracting outcomes affects project commitment. Hosts avoid "over-hyping" potential projects to investment analysts and usually provide a significant "haircut" to the full level of requests for service received when contracting has not been completed.

Obviously sponsors experience similar trepidation in project announcements when requests for service are not yet signed and construction not initiated. The level of uncertainty is magnified as their potential order book is far larger, more diverse and more distributed than that of hosts. Nonetheless, sponsors need to be focused on similar challenges to project development and construction at scale and rethinking how an avalanche of projects can be most efficiently planned and executed.

Projects to portfolio

The fundamental challenges of handling multiple, simultaneous projects can be overwhelming when discrete sponsors of all sizes and experience are vying for host attention during a flood of requests for service. Thinking with clear purpose is necessary to handle the increasing flow of project requests through a tightening funnel – particularly around deployment, focus and execution of host resource capabilities, process design and effectiveness, and ways to simplify and expedite underlying requirements and timelines.

The nature of sponsor project requests extends from enterprise to colocation to crypto to cloud to hyper- to mega-scaler levels, along with a focus maturing from conventional business activities to generative AI. Typically, data center projects are unique undertakings by sponsors and hosts since they can span different purpose, entities, scales, ramp–ups, market profiles, schedules, and geographies. These projects flow into hosts on an uneven basis and make the management of overlapping or duplicative requests for service complicated to assess and process, particularly when inexperienced sponsors require additional attention from local hosts which are already spread thin on resource deployment.

Sponsors and hosts need to evolve from project-by-project orientation to a portfolio model, but each faces a different reality. As hyper and mega-scaler sponsors pursue increasingly larger projects requiring tailored solutions, hosts are facing significant impacts to future peak demand and total commercial load several times their current base.

Hosts can shift their focus from bespoke projects with unique needs (a "once-through" model), to one where the full inventory of in-flight projects is evaluated, processed, prioritized and managed as a portfolio - an "enterprise investment" model with common-purpose frameworks.

The rationale is obvious — there are far more commonalities and dependencies among data center projects than there are differences. With a multi-facility view (e.g., individual Magnificent 7 sponsors, etc.), functional delivery teams can be comprised of common capabilities, (e.g., engineering, legal, etc.) and power supply, transmission and infrastructure assets are integrated at multiple levels (e.g., grid development and system balancing.)

Under the current project model, sponsors and host hosts need to change their thinking from unrelated projects across discrete geographic sites to multiple projects at single sites. The fundamental work supporting each data center facility will still be conducted, but would be scoped, scaled, sequenced and stood up in a uniquely different manner to create a multi-facility ecosystem operated as an integrated site. Some activities like land acquisition, permitting, power supply and transmission planning and financing will need to be entirely rethought to consider new portfolio complexities.

For hosts, a portfolio model can synchronize and optimize the planning, resource deployment, investment analysis, capital optimization and risk management processes. In addition, this model can improve decision speed and quality, sponsor relationships, supply chain positions and third-party coordination.

For sponsors, a portfolio model impacts how they interface with hosts as project numbers and breadth grow. These sponsors have many of the same challenges in data center development, construction and activation as do hosts, e.g., total capital investment management, multiple project alignment and coordination, prioritization trade-off considerations, and need for enhanced foresight of broad market developments.

Some hosts in the peer group indicated they were evolving their internal planning, evaluation, and execution frameworks and processes into a more purposeful portfolio model, but it generally remains a work in process for hosts, with indicated intent by several companies to move beyond a project-by-project planning and execution process. (See Figure 20)

Figure 20: Evolution in management approach

Interviewee positioning

Source: EY-Pathenon analysis

As Figure 20 illustrates, five of the 14 hosts interviewed have adopted a portfolio-driven model, and four others are considering this direction. This suggests the resulting economic and process benefit is worth the additional structure, resources and commitment from adoption. It also indicates that hosts must enhance the nature and level of project perspectives they rely upon today to have a value-centered, unified approach to standing up a portfolio of projects.

The need for purposeful action by sponsors and hosts will become acute as more projects are in flight, more become locationally concentrated, more capital is deployed, more assets need to be interconnected, risk escalates, and financing commitments are required.

Project management

Just as hosts need to evolve their asset management from projects to portfolios, they also need to elevate how they think about existing project management frameworks and approaches to data center interconnection – development, construction and activation need to be collectively managed in an increasingly complex and demanding environment where execution risks continue to become heightened. While project planning and execution are foundational activities and capabilities of hosts, the requirements for success continue to become more demanding and the risks more impactful.

As a rule, hosts generally have developed formal and comprehensive management frameworks to guide them through multiyear, multidimensional and complex capital investment and operational undertakings. These management models have been developed over years of large capital projects or multiyear financial commitments, adding much needed discipline to structuring project resources, planning execution at more discrete activity levels, embedding hands-on management at multiple project levels, and implementing control systems and technology to complement experience and take advantage of prior capital project lessons learned.

Often, large capital projects, i.e., those over $1b, do not have a strong track record of cost and schedule performance, and this can worsen for true mega-projects, i.e., comprising multiple billions of dollars. In most cases, large off-plan projects have not been well conceived, executive management was not sufficiently engaged over the project lifecycle, and work execution performance discipline was below expected levels, whether by owners or third parties.

From the early days of data center development, construction and activation, hosts generally did not have to incur significant costs to enable facility stand-up, as projects were smaller in scale, had limited requirements and were thus cheaper to address. However, as the growth and proliferation of large hyper-scale projects ratchets, more capital commitment for power supply and/or transmission investment is required. And, as current projects consistently extend into the mega-scaler domain (above 1GW and $1b thresholds), delivery periods can stretch five to more than eight years longer, uncertainties about execution increase, and the risks to sponsors and hosts similarly intensify, EY-Parthenon analysis shows.

Consequently, hosts need to reassess their current project or portfolio delivery models and challenge whether their prior experience and capabilities sufficiently extend to a new era of simultaneous, synchronized and sustained investments. This need has become particularly acute as hosts are taking on new power supply and transmission development involving multiyear timelines in a market where critical equipment is in short supply; manufacturing, infrastructure and construction resource supply capabilities are stressed; and sponsor timelines are pegged to ramp up activation and value creation.

For example, sponsors need sufficient power supply and can either contract with OEMs directly or secure power supply development, construction and activation from a host with a track record of delivery. When a hyper- or mega-scaler turns to a host as a power supply source and for transmission access, responsibility is taken on by this entity and circumstances can lead to multiple large MW capacity units, with multiple years of project construction and delivery execution – this can become a high risk undertaking and elevate the need for an enhanced project management discipline.

The need for an elevated project management model is not just directly construction-related, it extends to comprehensive engagement and integration of internal functions across the business to bring the "right" resources to bear on the project. These resources need to not just be loosely aligned, they need to be formally integrated to include areas such as: legal, finance and accounting, tax, regulatory, economic development, real estate, risk management, technology, human resources, engineering, construction, distribution, project controls and project management. Each of these traditional groups are typically involved over the course of a data center project.

Only a few of the larger companies interviewed have evolved to an elevated level of a more integrated project management model, capabilities and cross-enterprise resource dedication, largely because the level of data center activity and increasing scale-driven risks they are identifying indicate there is a need for this formalization.

From our interviews with 14 hosts, six companies have already adopted more permanent, integrated and disciplined project management models, with two presently considering this enhanced approach to data center planning, contracting, execution and stand-up. These organized project teams can include up to 50 internal resources aligned end to end for execution and delivery.

Once again, hosts see the value in leveraging their existing experience to enrich relevant processes – like project management – to preserve or build execution value. In the case of sponsors, the same logic applies. To be effective, this heightened visibility and emphasis on project management needs to work in a close, integrated manner – likely as a seconded, integrated and discrete group. However, it will need no less than a focused, coordinated and collaborative working team consistently engaged that communicates in real-time across the portfolio of projects, with direct line-of-sight into a realm of matters. These include all aspects of the request for service, FEED studies, interconnection studies, fee determinations, risk management, contract framing, negotiation, asset development, and tariff design, application, and pricing, and make-whole arrangements.

Establishment of an enhanced project management model and capabilities set will aid hosts, in preserving shareholder value, protecting customer affordability, enhancing value through reduced risks, enabling new value to be identified and retained through sponsor alignment and collaboration processes and minimizing surprises from planning or execution gaps that could potentially be less recognized.

Hosts need to mature their approach to dealing with sponsors and move beyond a transaction-based model to one where both parties are more aligned and integrated around project planning and delivery and its underlying requirements, individual objectives and desired outcomes.

Moving from a project to a portfolio model implies adoption of a rigorous project management model, while a broadened focus on collaboration suggests that planning and execution processes need to be geared toward considering decisions, investments and activities from a more integrated perspective.

In the right circumstances, this portfolio approach – coupled with expanded sponsor-host cooperation and collaboration – can be readily applied as a fundamental platform for building and embedding a broader ecosystem model where concentrated facility development, construction, stand-up and activation are a natural focus of financial, project and risk management.

Collaborative mindset

As hosts enhance their approaches to fulfillment of data center stand-up, they also need to rethink their long-term positioning with sponsors — particularly the most active Magnificent 7 entities. However, hosts also need to stand back and consider the implications of advancing relationships with companies within the sponsor group and across sponsors collectively.

In many cases, hosts and sponsors have a straightforward transactional relationship, i.e., enabling sponsors to interconnect their data center facilities to the grid as expeditiously as possible. But this is a narrow view of what a sponsor-host relationship could look like. Given the expected sustained growth in number and scale of new data centers into the late 2030s (and potentially well-beyond), an alternative approach — especially long-term positioning with the most active Magnificent 7 sponsors — may better serve the interests of both parties.

The conventional focus of the peer host panel has generally been on addressing four topics: providing general information to sponsors regarding suitable sites for data center stand-up; engaging with sponsors on information requirements necessary to clarify and support a sufficient request for service; providing sponsors with insight into potential local community challenges; and, closely coordinating with sponsors on design, scale and ramp-up to support FEED and interconnection studies. Most of these interactions are intended to be more cooperative than collaborative, as hosts rightfully emphasize arms-length relationships leading to eventual contractual arrangements.

But future facility upscaling, the expected long-term lifecycle and financial and operational implications to hosts and sponsors of the observed data center phenomenon, suggest there is tangible and intangible value to be created by elevating the sponsor-host relationship from purely transactional to one that has mutual-benefit characteristics. Within our 14-host peer group, six entities are still considering whether to move beyond their conventional transactional engagement model, while the other half are acting in a manner to drive mutual benefits.

Many sponsors are serial data center sponsors with a history of parallel facility stand-up. And hyper- and mega-scalers have an even longer history of activated facilities and a robust portfolio of increasingly larger projects serving valuable commercial customers. Similarly, hosts have a more than 100-year presence in their service territories and a natural long-term commitment to their customers by preserving a strong position as a trusted local partner. Neither sponsors nor hosts seek to undermine these positions through intentional or inadvertent actions.

Consequently, it's in the interests of hosts to build and sustain broader business-to-business working relationships to forge a stronger and mutually beneficial market presence, e.g., expanded front-end communication on project siting, tighter early location evaluation and system impact planning, cooperative development of intermediate and longer-term contract frameworks, an embraced "partner of choice" role on relevant power supply and transmission development, open coordination on interconnection execution and collaboration on conceptualization, evaluation, configuration and stand-up of multi-data center cluster or hub environments.

Both parties will necessarily need to preserve their individual interests, but opportunities remain to build a more collaborative relationship and identify actions that can create value potential to each entity. This value can arise from: more integrated asset planning; improved visibility into data center requirements and site configuration; optimizing future capital spend levels and project delivery outcomes; enabling leverage capture of "go-to-market" buying advantages; and reducing the risks large and sustained capital programs can experience from development, construction and activation delays.

An ecosystem model

The advent of the data center phenomenon has spawned multiple GWs of additional load to US hosts between 2022 and 2024 across a range of purposes, sponsors, facilities, tenants, geographies and scales. This growth has occurred from standing up hundreds of data centers over that time period, with no definitive end in sight.

As could be expected from a new market model growing exponentially, the number of transitory data center sponsors can outnumber actual participating entities. It takes little to become a general data center developer – some front-end investment money, good instincts for land value and an ability to remarket developed and undeveloped sites. A deep pool of large-scale parties — like the Magnificent 7 — is focused on securing hundreds of MWs (or multiple GWs) to drive additional power computing.

Data centers are individually unique and represent a "technology in time" purposed for specific use and outcomes. This uniqueness comes from different sponsor perspectives on data center market positioning, which leads to how intended outcomes can be best accomplished, economically, technically and/or operationally. While data center facilities are focused on meeting the needs of the sponsor and related tenants, each differs based on stand-up vintage, facility scale, location and the nature of the entity as the sole or primary sponsor. Sponsors may have a stand-alone investment to drive margin or be a critical enabler of a strategic mission core to the business, like computing power for generative AI.

Similarly, hosts have different views of data center growth – as a "unicorn" that provides a significant new lifecycle revenue and cash flow source, or a way to enhance small and large customer relationships through facility enablement, like seamless energy provision through sustainable power supply.

To stand up a new data center from request for service initiation through actual stand-up and activation can require between eight and ten entities. This number of participants can vary depending on individual data center requirements, e.g., whether a host utility in a typical regulated market currently has sufficient generation and/or transmission, where new power supply is needed in regulated or RTO footprints, and whether the timelines of public processes in an RTO can adequately support the rapid pace of additions to new capacity or transmission paths. It can also vary depending on the number of financiers involved with a project or the parties engaged with sponsors or hosts in enabling data center stand-up.

These unique market structures and conditions can have significant impact on the events that need to occur to secure land parcels and permits, finalize project parameters, conduct front-end load and interconnection studies, order necessary equipment (particularly where supply chain delays persist), upgrade the infrastructure, construct the facility, energize the site and activate the stand-up.

While the scale requirements of any data center development cycle can differ, the range of participants regardless of scale can involve:

- One or more OEMs for power supply or substation equipment.
- A primary contractor for construction.
- One or more transmission OEMs or providers.
- A fiber provider.
- A water access source.
- One or more financiers.
- One or more off takers.
- A sponsor.
- A host.

Generally, these players have operated as individual providers of discrete assets or services and act more separately than as an integrated network. But market processes are shifting and appear ready to evolve to capitalize on tighter participant integration to deliver increasingly larger, more complex, longer duration and more capital-intensive projects.

It is no longer common that each new activated data center project resembles those that came before it — technology components are more powerful and efficient, facility equipment is more compact, processors offer enhanced speed and efficiency, footprints are more sizable, power supply and back-up requirements are growing and site infrastructure is expanding.

A few hosts have migrated to a more formal model for data center hosting — both for configuration and economies. Sponsors are already standing-up ecosystems — sometimes referred to as campuses, clusters or digital hubs — comprising multiple data centers in a designed and managed deployment. In addition, hosts are enabling these facilities through grid, network and distribution infrastructure. As more sponsors move to leveraging this concept, they will find available economies of scale can be realized in asset services (construction) and equipment acquisition (supply chain), fiber configuration and water sourcing, among other areas.

These ecosystems can be built-out in stages, with complete facilities staggered over succeeding years of activation (a serial build) or built-out in parallel fashion with each facility scaled for ramp-up as load grows from either new tenant demand or increased computing power requirements as generative AI capabilities increase (a tailored build). Adopting this type of approach simply captures the natural economies available from facility configuration, i.e., building four 250 MW units is cheaper than standing-up 10 100 MW facilities through more efficient infrastructure (like for fiber), increased power supply build-out (like from natural gas), larger contract letting (like for construction), faster delivery of key equipment (like substations) and higher site productivity, (like from site configuration and contractor continuity).

For sponsors, an ecosystem model enables a fundamental focus on realizing lifecycle economies and benefits from installation or expansion from fiber providers and water access provision or expansion from providers in a greenfield site, OEMs for power supply equipment (where new owned capacity is built) and EPCs and contractors for facility planning, construction and installation of assets and multiple data center buildings within a pre-determined footprint and timeframe. (See Figure 21)

Figure 21: Framing an ecosystem

Conceptual data center ecosystem

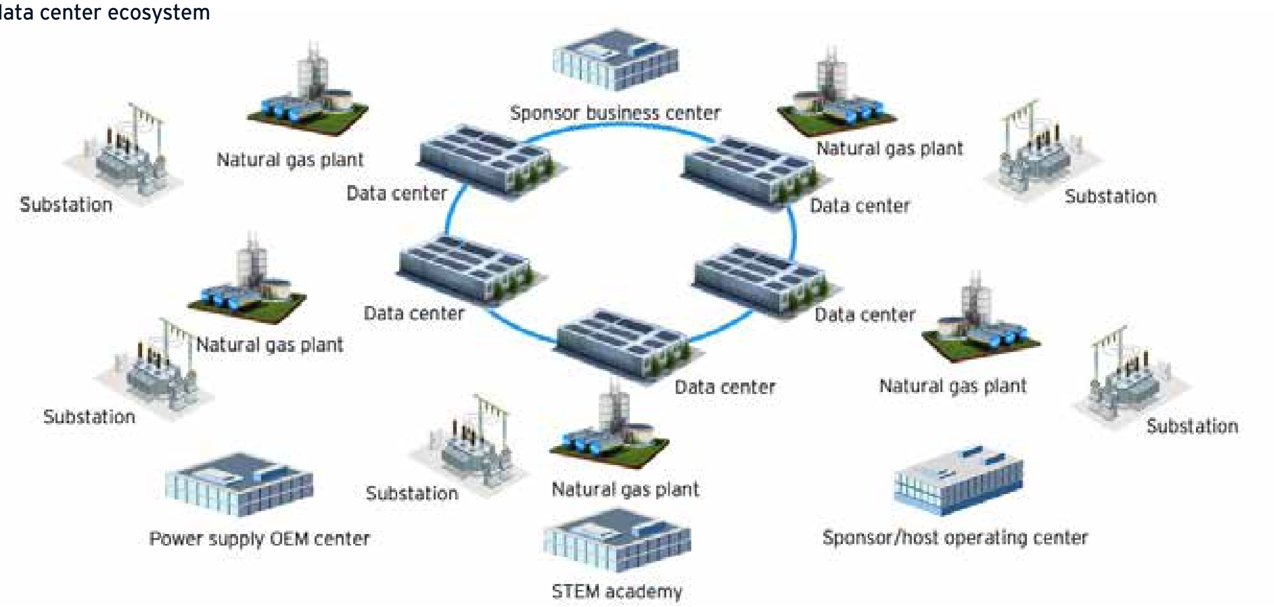

Note: Back up generation may consist of multiple technologies that are not shown in the figure

Source: EY-Pathenon analysis

Similarly, planning for multiple data centers at a single site concentrates the planning window and enables sponsors to align with existing and future order books of OEMs, EPCs and contractors to both ensure availability and secure attractive pricing from increasing order size and certainty.

An ecosystem model also enables the sponsor to add co-located facilities that provide for greater regional presence within the site perimeter for related facilities, e.g., technology support, market and business capabilities and other enabling provider facilities. This model serves to integrate the full range of entities and resources that enable facility infrastructure and support ongoing operation.

From a host perspective, some of the same planning features and outcomes are relevant here as well. When the host needs to build incremental power supply and transmission assets to support new load, economies of scale can be captured through negotiation with third parties across multiple data center asset and facility requirements. Hosts can also seek to secure or realize lower total and/or unit costs of development, construction or installation depending on the nature of the asset and the requirements of the work. And it can establish aligned facilities, as necessary, with the sponsor and other providers to ensure configuration and operations integration.

This is not to say that all ecosystems will look the same in composition and configuration. Practical site considerations, including generation or transmission concentration limitations, need to be individually assessed for both larger and smaller hub locations to avoid excessive financial, operating and asset risk.

An ecosystem enables sponsors and hosts to optimize data center projects from inception through stand-up and activation, and ultimately through retirement. It creates a cohesive, collaborative arrangement to address the types of planning and execution dependencies that exist on large projects, predicated on rapid execution where time to value is important, i.e., the speed-to-power of data center stand-up. Thinking of future data centers from this perspective introduces a dimension that extends beyond the activation of a single facility. It centers on how sponsors, hosts and third parties can capture value in a more purposeful and broader manner. This is particularly important as the scale and pace of data center development and activation continues to increase. This level of scale growth heightens the sensitivity of sponsors to project readiness and stand-up performance and the risks to hosts from future impacts on the adequacy of power supply to support these facilities and internal capacity to manage multiple projects simultaneously.

From both primary and secondary participant perspectives, the sponsor and host can place themselves in advantaged planning, build-out and economic positions that simplify near- and long-term data center development – in continuing single or multiple configurations where they can be leveraged.

But standing up one or more concentrated, multi-facility data center sites is not the entire end game, it is just the first step in a broader evolution and configuration of the ecosystem concept itself. The ultimate stage in its development is a result that enables data center sponsors, providers, partners and hosts to continually be collaborative and effective in standing-up and/or expanding these sites to provide for ongoing economies and effectiveness. (See Figure 22)

Figure 22: Value from an ecosystem

Participant value

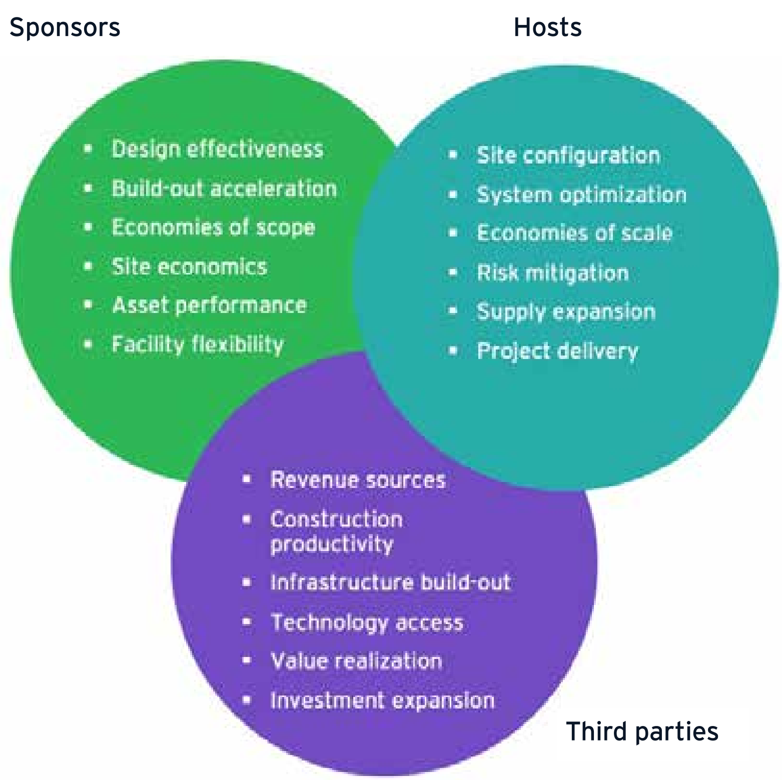

Rather than consider the mission ended with activation of only one or more data centers, the sponsor and host can extend the conceptual and physical perimeter of the relationship in a manner that takes a longer-term view of how each entity can benefit over broader activities and projects. For example, where multiple data center, power supply or transmission assets are constructed, the potential increases for a sponsor and host to co-locate additional facilities, such as a regional technical support organization.

Similarly, an EPC that oversees and drives data center construction can locate a construction planning and delivery support organization that enables services to be provided from within the ecosystem for the specific site, as well as others that may be located nearby, e.g., along a transmission corridor that can support multiple data center sites.

Through the end of 2024, five of the peer hosts interviewed had adopted or initiated some form of an ecosystem or digital hub model where multiple data centers anchored a common site alongside facilities – either sponsors, hosts, OEMs or other project related entities – and were progressing toward a more defined, developed, integrated and sustained relationship.

However, an additional five of the 14 peer hosts similarly indicated they were interested in the potential value from this model or were actively pursuing some element of it.
(See Figure 23)

Figure 23: Future ecosystem focus
Peer company positioning

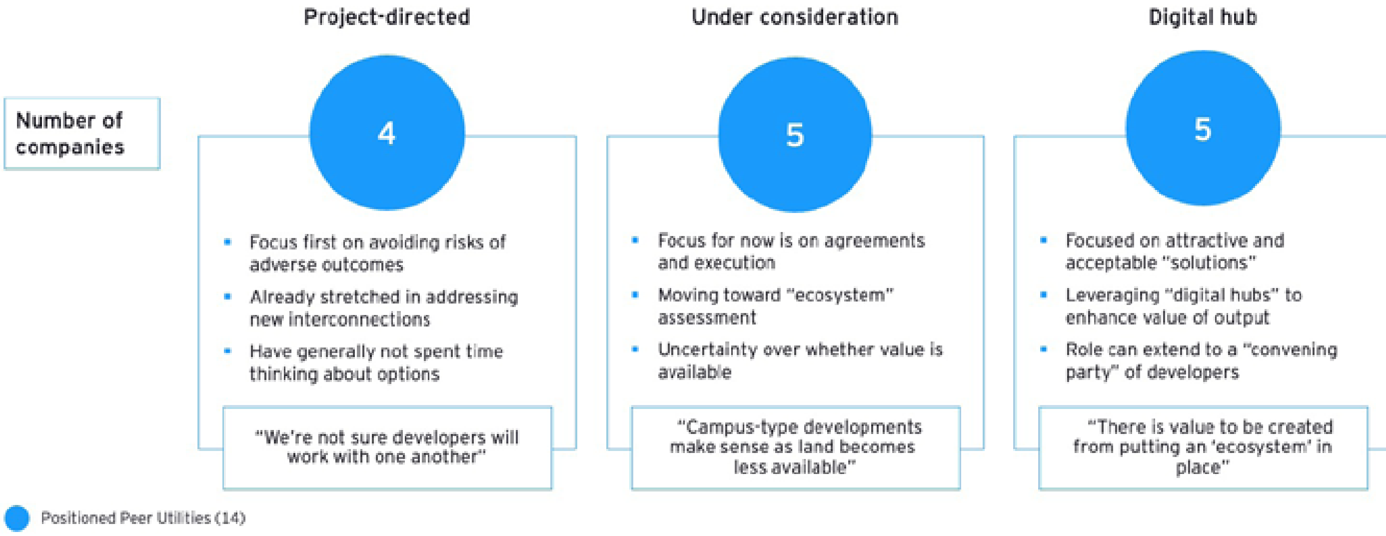

Source: EY-Pathenon analysis

Hosts can think more broadly about potential future data center alignment and integration with an ecosystem model that embraces collaboration and cooperation at its core. This approach further enables the concept of joint value capture to be expanded beyond simple economies. Value will be tangibly captured through production outcomes or asset enablement for sponsors and infrastructure providers, as well as from economies of scale that exist. But this value is further enhanced through thinking about bilateral outcomes for sponsors and hosts alike (and extends to other stakeholders, like communities).

Sponsors obviously receive economic rents available from leveraging and advancing the capabilities of data center computing power utilized within their business model. Hosts capture new revenues from power supply provision and lower infrastructure costs. Both entities collectively create benefits and value for other adjacent stakeholders that participate in site and facility stand-up and ongoing operational support. With the number of requested, planned, contracted, under-construction and ready to activate data centers, the opportunity exists for sponsors and hosts to rethink how they align these facilities to support both near-term operation and longer-term positioning.

As single projects, data centers are often viewed by sponsors as serving specific purposes – to bring necessary computing power to address specific tenant needs, extend computing power availability, accelerate computing power speed and/or increase computing power value. Comparatively, hosts view projects as a welcome source for new revenue growth but typically do not think about them in a more holistic manner that looks across the inventory of projects in-flight to provide the benefits of proximity, integration and risk diversification on an aligned or integrated basis.

While site build-out can occur over several years to coordinate and align supply chain, RTO engagement (if required), power supply, transmission, construction and activation activities, both the sponsor and host will need to think about the build-out as a multi-location, multi-facility, multi-entity project. This will allow both partners and other providers to focus on optimizing configuration, planning, spending, construction and activation activities to enable a larger footprint to be developed and constructed. This potentially drives significant economies of scale and enhances project and risk management across facilities that can be replicated throughout sponsor and host communities.

Not every data center development provides the right circumstances where an ecosystem can provide equivalent value. But in many circumstances, the key players in data center development and activation will find they share similar philosophies of optimizing facility and economic outcomes.

Opportunities are emerging to create value by clustering data centers. To effectively capitalize on scale, configuration and integration, sponsors and hosts need to resolve core execution issues and risks in data center development, interconnection and internal execution.

The case for value

An ecosystem-based model can become a natural source of value when sponsors and hosts evolve from enablers of project outcomes to creators of individual and joint value. Incremental value is created when multiple facilities are on the same site – both physical operating assets and regional or technical support facilities that are developed, constructed and operated as integrated facilities. Creating and capturing economies across a larger and more integrated data center site helps demonstrate benefits for adjacent providers, communities and customers.

If an ecosystem can be created with data center facilities at its center, secondary and tertiary value flows can transfer to financiers, fiber, water and industrial solutions providers, other grid providers, municipalities, regulators and customers: Meaningful jobs can be created; additional property taxes can be collected; power supply, grid and network utilization can improve; electrification can be expanded; financing costs can be lowered and regulatory policies can be advanced. (See Figure 24)

Figure 24: Benefits from data center stand-up
Categories of data center benefits

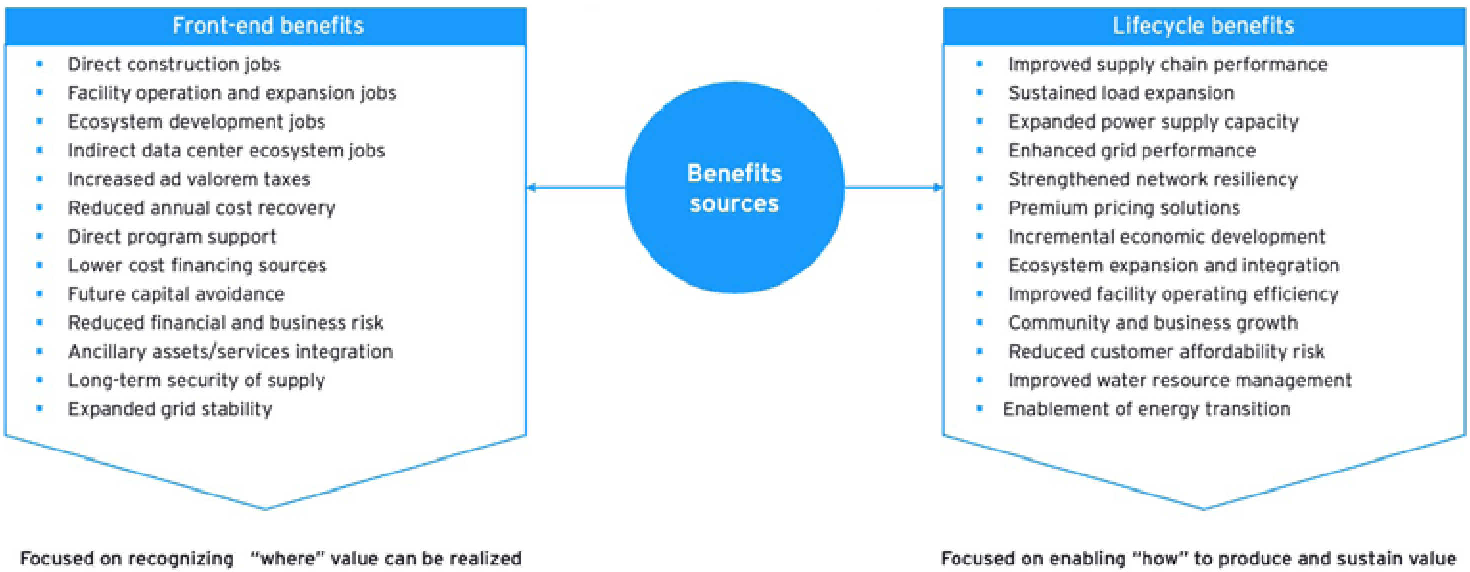

Some data center participants, e.g., fiber and water providers, are fundamental to downstream value delivery and often have already provided a foundation for expanded fundamental services to sponsors over the post-activation lifecycle of a data center. These participants are initial enablers of data center value realization, and their presence typically precedes sponsor or host engagement.

Hosts provide early contribution in the data center assessment process through conduct of the required FEED studies to evaluate feasibility of the site, environmental impacts, facility load, power requirements, timing, building cooling needs, internal technology design, electric infrastructure needs and preliminary costs and sensitivities associated with facility stand-up. This activity is core to understanding high level data center impacts to determine whether project feasibility can be affirmed prior to detailed facility assessment, analysis and design. Value to sponsors at this stage comes from the validation of project purpose and related assumptions and provision of utility perspectives on initial planning parameters and related costs and complexities.

Next stage interconnection studies can define the optimal paths to design, construction and/or activation and include broader and deeper assessment of loads, load impacts and capacity adequacy from related sponsor requests for power supply and transmission, sponsor speed-to-power expectations, current infrastructure sufficiency, system operating impacts, grid interconnection requirements and potential cost estimate modifications and implications. Activities performed by hosts at this stage enhance perspectives on potential planned data center outcomes and allow sponsors to fine-tune their assumptions, work with hosts on facility siting and asset operating parameters and reconsider project scale, ramp-up, timing and complexity of power supply and transmission permitting and interconnection.

Once projects move from development and design to construction and activation, OEMs, EPCs (as needed) and general construction contractors play a direct role in enabling manufacturing, constructing, or activating power supply (including back-up), as well as related activities for transmission upgrade or new design construction. Some of these activities will be supported by the host in conjunction with these third parties, e.g., power supply planning, while others may be directly undertaken by the host itself, e.g., transmission siting and construction.

This stage is critical to the realization of value by the sponsor and host as speed-to-power is highly dependent on project performance — asset design, construction and stand-up — to enable facility activation and system interconnection. Applying a broad concept to definition of front-end and lifecycle value to sponsors, hosts, communities and other stakeholders suggests that data centers can become value engines in ways that extend beyond pure operating outcomes. (See Figure 25)

The nature of benefit to sponsors and hosts is only part of the value equation related to data centers. The communities where data centers are located — including municipalities, taxing authorities, commercial businesses, local suppliers and contractors and customers — can all be beneficiaries of data center project stand-up over the course of a project lifecycle. Similarly, entities, such as financiers, adjacent data center operators and grid operators can also benefit through lower costs and risks.

Figure 25: Business case dimensions

Benefits and value enablers

To date, state regulators have been relatively neutral regarding data center development, execution and stand-up. However, in some cases, they have paused continued data center development. While they are acutely interested in how these new facilities affect customer affordability, they have relied upon hosts to utilize existing pricing structures to avoid customers bearing responsibility for costs they do not benefit from. Tariff design models, along with contract structures, terms, penalties and effective risk apportionment, combine to protect all customers from responsibility for direct (grid-related) or indirect (performance follow-through) issues that emerge later in the data center lifecycle.

In a few locations, however, hosts, regulators or communities have questioned data center project impacts to water access, system imbalances, alternative land use, customer affordability, stranded cost recovery, noise and aesthetic issues and other matters being raised by the public. In response, guardrails to govern handling have been developed to mitigate unforeseen adverse impacts.

As sponsors and hosts elevate the importance of value outcomes from data center stand-up, they need to help regulators and policy makers understand the benefits that will likely emerge from more efficient facility design, enhanced load and power supply planning, faster project execution, reduced risk exposure, acceleration of energy transition goals and lower costs to customers over the long-term.

In their view, if sponsors and hosts work within established principles and boundaries that state regulators have participated in developing or approving, then ground rules are clear and can be effectively administered and applied, as appropriate. Achieving equitable outcomes through each pricing agreement is a cornerstone of every sponsor-host contractual arrangement.

State regulators are also interested in understanding how the role of the clean energy transition is aligned with the power supply sources selected for each facility. As noted earlier, depending on only clean energy sources is a laudable goal, but not realistic and sponsors have moved away from absolute adherence to those objectives in favor of a better balance of "what is practical, what is available and what is economic" — which means it would likely include natural gas where speed to power is shorter and potentially nuclear for larger scale facilities with more allowance for an elongated power availability timeline. Renewables and battery storage will still be a part of these solutions, although they may become ancillary, rather than primary, supply sources as AI inference-focused data centers grow in scale and require 99.99999% performance levels.

It is likely that Magnificent 7 and other sponsors will also begin to frame their data center projects as sources of value that extend beyond the operational benefits. This value may range from continuous power computing for tenants or sponsor self-use in a rich AI inference environment, to the financial contributions from buyers of attendant power supply.

From 2025 through 2030, sponsors will not only drive robust data center stand-up but will also grow into a more economic value capture role beyond computing power and beginning in advance of facility activation — capturing economies of scale in supply, shortening design and construction timeframes to accelerate speed-to-power, driving facility operating economics and enabling further net-zero outcomes.

Capitalizing on opportunity

Hosts are naturally positioned to enable robust, effective and efficient data center expansion. While the future trajectory of data center growth is uncertain, contracted and initiated projects leading to 2030 and slightly beyond have a higher degree of certainty and are reflected in utility IRPs and/or capital expenditure forecasts.

Substantial data center growth opportunity also exists from requests for service or interconnection that have not been committed to or contracted, or where hosts have higher uncertainty about the validity of the request or sponsor capacity to move to a contracting stage in the near-term. Some portion of these potential uncontracted requests for service and interconnection — maybe a significant amount as demand levels soar — will meet the project acceptance conditions of hosts and further increase demand loads necessary to meet sponsor, tenant and off-taker requirements.

Nonetheless, hosts need to be prepared for these potential data center projects progressing and further taxing utility host resources, particularly as greater scale and complexity issues are likely to occur than in prior build-out.

Figure 26: Repositioning sponsors and hosts
Changing the focus

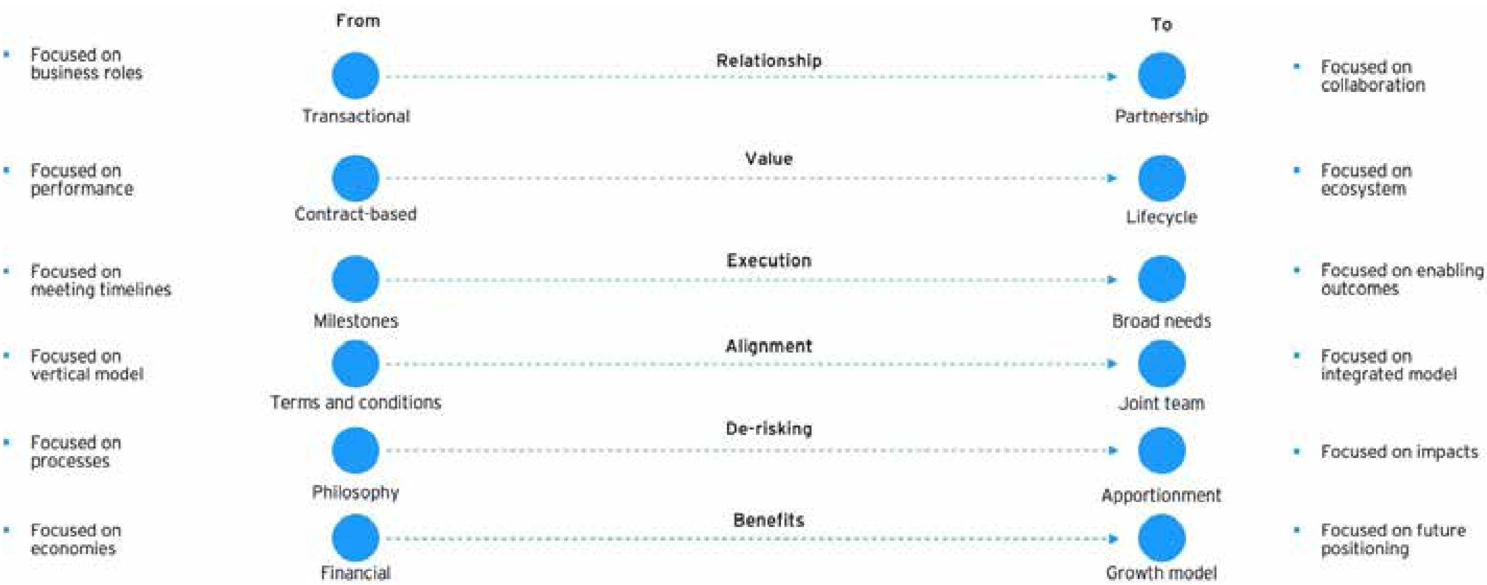

Source: EY-Pathenon analysis

Readiness for the next stage

As potential data center request-for-service backlogs increase and cost and speed-to-power scale up, hosts would be wise to focus on eight key actions to further elevate their market position and capabilities to effectively execute and benefit from a growing backlog of projects in future years:

- Develop a formal program risk evaluation and management model – **think like an ex-ante risk manager v. an ex-post problem solver**
- Stress-test and de-risk each discrete element of the data center execution process – **think like an integrated enterprise vs. a siloed enable**r
- Establish a schedule of common and tailored fees and charges – **think like a recurring service provider vs. an ad hoc performance specialist**
- Create contract structures and tariffs that simplify future contracting and enable risk avoidance – **think like a risk steward vs. a contract negotiator**
- Enhance the ability to drive sustained performance across a large and parallel level of projects **– think like a portfolio manager vs. a serial executor**
- Embed an integrated program management framework across the enterprise – **think like a fully seamless entity vs. an amalgam of related functions**
- Design data center planning, execution and stand-up activities to drive enhanced economic outcomes – **think like a value creator vs. a cost minimizer**
- Position the enterprise as a value-added partner to sponsors and enablers – **think like an active collaborator vs. an arms-length transactional role performer**

If a host can create a more mutually-valuable relationship with mid-sized, hyper and mega-scaler sponsors — particularly the larger Magnificent 7 players — it can position itself as a partner of choice for the long term and identify a range of future growth and value opportunities that would not have been available.

As busy as hosts have been with data center stand-up, their participation has not yet brought a concentrated lens to where and how innovation can be leveraged to enhance internal planning, execution and bi-lateral engagement with sponsors. Innovation can be pursued by both sponsors and hosts across multiple macro- dimensions of a data center project: technical services, materials engineering, equipment efficiency, commercial framing, process (and therefore timelines) simplification, sponsor-host governance and value concepts.

However, it is now time for sponsors and hosts to collaborate on how to drive innovation across discrete activities that can improve outcomes from data center pursuit and stand-up. Identifying and eliminating recurring pitfalls, like equipment shortages, interconnection queues, collaboration models, standardization concepts, value creation, power supply construction durations, pricing structures and financing alternatives, among others, all can improve data center value chain experience and lifecycle outcomes. Focusing more on innovation that leads to intended outcomes is not a one-time activity, it needs to be an embedded capability constantly centered on driving economies, reduced durations, reduced risks and value additionality.

The pace of data center change in just the last five years has increased the MW scale of the average data center facility from 20 MWs in 2019 to approaching 100 MWs in 2025 and new service requests are accelerating in both total levels and scale. A recent EPRI survey of hosts actively involved in data center evaluation, development, construction and stand-up indicates that the average host will see its data center peak load contribution increase from 0% to more than 35% in the next five years.

Even if the lofty growth projections of peak load and total demand growth turn out to be lower than expected, the level of new demand growth will still provide the electric utility industry with a sustained boost to growth not experienced in over six decades. Sponsors and hosts may both enjoy a sustained period of robust growth whether the data center phenomenon lasts for 10 years, 15 years or longer.

With the burgeoning level of economy-wide and AI-driven demand growth, future forecasts remain, notwithstanding the uncertain rate of future growth. The recent news of the smaller, more efficient and cheaper DeepSeek AI model, relative to OpenAI, surprised many and has received close scrutiny. Whether or not Jevons Paradox plays out in light of increased power computing usage effectiveness remains to be seen. But if historical outcomes continue, the era of TWh growth may not be affected as higher efficiency and lower costs combine to produce greater consumption. In addition, continuing technological evolution and an announced national interest in maintaining AI leadership provide an impetus for the future growth trajectory.

To prepare for these growth outcomes, sponsors and hosts need to meet the challenges known today, while preparing for an elevated, accelerated-but-less-certain market where internal planning capabilities and external execution capacities are already stretched. The more attention paid now to identifying and understanding potential future risks and the value that can be captured today and long term, the more likely it is that data center stand-up can meet the requirements of sponsors, hosts, communities, investors and other parties that are affected by data center operations.

Data center projects are high-expectation, high-visibility and high-impact capital undertakings. Yet the future trajectory of data center growth and energy consumption is truly not predictable into the latter stages of the 2030s and beyond. For the short-term — perhaps through the mid-2030s — anticipated outcomes look more predictable, but energy needs and consumption levels, particularly driven by generative AI, are not fully known and are highly dependent on current assumption accuracy.

For the data center phenomenon to continue for several decades, robust AI use cases would need to continue to expand. Industries are planning on continuation of the data center phenomenon for several decades, but this achievement depends on whether robust AI demand continues or plateaus as use case growth peaks at some point in the future. However, NVIDIA's chief executive officer suggested at an AI conference in March 2025 that the world needs 100 times the amount of computation, relative to just one year earlier, due to advanced AI requirements.[52]

Regardless of the ultimate level of realized demand growth that results, both sponsors and hosts will benefit from an extended period of robust growth in technology advancement, capital spend, business expansion and revenue production.

Endnotes

1. EY-Parthenon analysis based on data from EPRI and other market reports. "Powering Intelligence - Analyzing Artificial Intelligence and Data Center Energy Consumption" Electric Power Research Institute (EPRI), Page 4, 5, 13 and 15, 28 May 2024. (https://www.epri.com/research/products/000000003002028905)
2. Wilson, John D., Zimmerman, Zach, "National Load Growth Report 2023," Grid Strategies LLC, Page 3, December 2023. (https://gridstrategiesllc.com/wp-content/uploads/2023/12/National-Load-Growth-Report-2023.pdf)
3. Govekar, Milind, "Gartner Says Cloud Will Be the Centerpiece of New Digital Experiences," Gartner Newsroom, 10 November 2021. (https://www.gartner.com/en/newsroom/press-releases/2021-11-10-gartner-says-cloud-will-be-the-centerpiece-of-new-digital-experiences)
4. Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption" Electric Power Research Institute (EPRI), Page 4, 5 and 15, 28 May 2024. (https://www.epri.com/research/products/000000003002028905)
5. Davenport, Carly, Lee, Brian, et al, "Generational Growth: AI, Data Centers and the Coming US Power Surge," Goldman Sachs Equity Research, Page 14, 28 April 2024, (https://www.goldmansachs.com/pdfs/insights/pages/generational-growth-ai-data-centers-and-the-coming-us-power-surge/report.pdf).
6. Moss, Sebastian, "Exclusive: After Meta cancels Odense data center expansion, projects are being rescoped globally," Data Centre Dynamics.com, 14 December 2022, (https://www.datacenterdynamics.com/en/news/exclusive-after-meta-cancels-odense-data-center-expansion-other-projects-are-being-rescoped).
7. "Meta halts construction of two data centres in Denmark," Reuters, 15 December 2022, (https://www.reuters.com/technology/meta-halts-construction-two-data-centres-denmark-2022-12-15).
8. Shi, Madeline, "Meet the 10 Most Active Data Center Investors," PitchBook News, 19 September 2024, (https://pitchbook.com/news/articles/meet-the-10-most-active-data-center-investors).
9. H.R.434 – "CHIPS and Science Act," Congress.gov, 9 August 2022, (https://www.congress.gov/bill/117th-congress/house-bill/4346).
10. EY-Parthenon analysis based on data from Grid Strategies and other market reports. Wilson, John D., Zimmerman, Zach, "National Load Growth Report 2024," Grid Strategies LLC, Page 3, December 2024, (https://gridstrategiesllc.com/wp-content/uploads/National-Load-Growth-Report-2024.pdf).
11. PJM Resource Adequacy Planning Department, "2025 Load Forecast Report," PJM, Page 6, January 2025. (https://www.pjm.com/-/media/DotCom/library/reports-notices/load-forecast/2025-load-report.pdf)
12. PJM Resource Adequacy Planning Department, "2021 Load Forecast Report," PJM, Page 3, January 2021. (https://www.pjm.com/-/media/DotCom/library/reports-notices/load-forecast/2021-load-report.pdf)
13. Neel, Evan, "Large Load Interconnection Status," Electric Reliability Council of Texas, 25 January 2024. (https://www.ercot.com/files/docs/2024/02/06/LLI-Queue-Status-Update-2024-1-25.pdf)
14. Neel, Evan, "Large Load Interconnection Status," Electric Reliability Council of Texas, 5 August 2024. (https://www.ercot.com/files/docs/2024/08/07/04-lli-queue-status-update-2024-8-5.pdf)
15. Srivathsan, Bhargs, Sachdeva, Pankaj, et al., "AI Power: Expanding Data Center Capacity to Meet Growing Demand," McKinsey Insights, 29 October 2024. (https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/ai-power-expanding-data-center-capacity-to-meet-growing-demand)
16. Berns, Maurice, Lee, Vivian, "How CEOs Can Achieve Both AI and Climate Goals," Boston Consulting Group, 8 November 2024. (https://web-assets-pdf.bcg.com/prod/ceos-achieving-ai-and-climate-goals.pdf)
17. EY-Parthenon analysis based on data from EPRI and other market reports. "Powering Intelligence - Analyzing Artificial Intelligence and Data Center Energy Consumption" Electric Power Research Institute (EPRI), Page 18, 28 May 2024. (https://www.epri.com/research/products/000000003002028905) and Choe, Richard, Nguye, Braden, "2025 Communications Infrastructure Outlook" JP Morgan, 6 Dec 2024.
18. EY-Parthenon interview, EY-Parthenon analysis, https://www.datacenter.com, "USA Data Centers," Data Center Map, 1 Dec 2024. (https://www.datacentermap.com/usa/)
19. Anderson, Jared, "Utilities Face Challenges, Opportunities from AI-Driven Data Center Power Demand Growth," S&P Global Commodity Insights, 1 April 2024. (https://www.spglobal.com/commodity-insights/en/news-research/latest-news/electric-power/040124-utilities-face-challenges-opportunities-from-ai-driven-data-center-power-demand-growth-report)
20. Sickinger, David, Regimbal, Kevin, "High-Performance Computing Data Center Water Usage Efficiency," National Renewable Energy Laboratory, August 2016. (https://www.nrel.gov/computational-science/reducing-water-usage.html)
21. Moore, Susan, "The Data Center Is (Almost) Dead," Gartner, 5 August 2019, (https://www.gartner.com/smarterwithgartner/the-data-center-is-almost-dead).
22. Dinsdale, John, "On-Premise Data Center Capacity Being Increasingly Dwarfed by Hyperscalers and Colocation Companies," Synergy Research Group, 12 July 2023, (https://www.srgresearch.com/articles/on-premise-data-center-capacity-being-increasingly-dwarfed-by-hyperscalers-and-colocation-companies).
23. Powell, Phill; Smalley, Ian, "What is a Hyperscale Data Center?" IBM Think, 21 March 2024, (https://www.ibm.com/think/topics/hyperscale-data-center).
24. Dutt, Arjun; Renno, Paul; et al, "AI Changes Big and Small Computing," Bain & Company Technology Report 2024, (https://www.bain.com/insights/ai-changes-big-and-small-computing-tech-report-2024).
25. Ungles, Brian; McWilliams, John; et al, "U.S. Data Center Development Cost Guide," Cushman & Wakefield, 4 March 2025, (https://cushwake.cld.bz/Data-Center-Development-Cost-Guide-2025/2-3).
26. "Microsoft announces $3.3 billion investment in Wisconsin to spur artificial intelligence innovation and economic growth," Microsoft website, 8 May 2024, (https://news.microsoft.com/2024/05/08/microsoft-announces-3-3-billion-investment-in-wisconsin-to-spur-artificial-intelligence-innovation-and-economic-growth).
27. Sriram, Akash, "Alphabet to Invest $3.3 Billion for Two Data Centers in South Carolina," Reuters, 26 September 2024, (https://www.reuters.com/technology/alphabet-invest-33-bln-two-data-centers-south-carolina-2024-09-26).
28. Landry, Jeff, "Meta Selects Northeast Louisiana as Site of $10 Billion Artificial Intelligence Optimized Data Center," Louisiana Economic Development, 4 December 2024, (https://www.opportunitylouisiana.gov/news/meta-selects-northeast-louisiana-as-site-of-10-billion-artificial-intelligence-optimized-data-center-governor-jeff-landry-calls-investment-

a-new-chapter-for-state).
29. Smith, Brad, "The Golden Opportunity for American AI," Microsoft On the Issues, 3 January 2025, (https://blogs.microsoft.com/on-the-issues/2025/01/03/the-golden-opportunity-for-american-ai).
30. "USA Data Centers," DataCenterMap.com, (https://www.datacentermap.com/usa).
31. EY-Parthenon company interviews and SEC filings
32. Berkeley Lab, "PJM Data Show Substantial Increases in Interconnection Costs," Energy Markets & Policy, 19 January 2023. (https://emp.lbl.gov/news/pjm-data-show-substantial-increases)
33. Seiler, Ken, "New Interconnection Process Reaches Next Milestone," PJM Inside Lines, 21 December 2023. (https://insidelines.pjm.com/new-interconnection-process-reaches-next-milestone/)
34. Glick, Richard, "Explainer on the Interconnection Final Rule," Federal Energy Regulatory Commission, 28 July 2023. (https://www.ferc.gov/explainer-interconnection-final-rule)
35. Smith, John, "AI for Interconnection (AI4IX)," U.S. Department of Energy, 25 November 2024. (https://www.energy.gov/gdo/ai-interconnection-ai4ix)
36. Berns, Maurice, Lee, Vivian, "The Semiconductor Industry in the AI Era," Capgemini, February 2025. (https://www.capgemini.com/wp-content/uploads/2025/02/CRI_Semiconductors_Final_WEB.pdf)
37. Weinschenk, Rose, Donnellan, Douglas, "2023 Data Center Staffing Survey," Uptime Institute, 15 December 2023. (https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2023StaffingSurvey.Report.12152023.pdf)
38. Freed, Peter, Clements, Allison, "How to Reduce Large Load Speculation? Standardize the Interconnection Process," Utility Dive, 19 February 2025. (https://www.utilitydive.com/news/data-center-large-load-interconnection-process-clements/740272/
39. Mallory, Ryan, "PUE: Power Usage Effectiveness," Flexential, March 2024. (https://www.flexential.com/system/files/file/2024-03/PUE-Power%20Usage%20Effectiveness-hvc.pdf?t=1740386980)
40. Bizo, Daniel, "Global PUEs – Are They Going Anywhere?" Uptime Institute Blog, 4 December 2023. (https://journal.uptimeinstitute.com/global-pues-are-they-going-anywhere/)
41. Smith, John, "Efficiency in Google Data Centers," Google Data Centers, 2025. (https://datacenters.google/efficiency/)
42. Jevons, William Stanley, "The Coal Question; An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal Mines," Energy History at Yale, 1865 (https://energyhistory.yale.edu/w-stanley-jevons-the-coal-question-1865/)
43. Rich, Brandon. "DeepSeek Explained: What is it and is it safe to use?" University of Notre Dame, Information Technology Center, 31 January 2025. https://ai.nd.edu/news/deepseek-explained-what-is-it-and-is-it-safe-to-use/
44. Smith, John, "Data Centers Best Practices," Pacific Gas and Electric Company, Page 64 and 76, 2024. (https://www.pge.com/assets/pge/docs/save-energy-and-money/rebate-and-incentives/DataCenters_BestPractices.pdf)
45. Perez, Kristen, Porter, Chris, et al, "Strategies for Maximizing Data Center Energy Efficiency," NVIDIA Technical Blog, 23 May 2023. (https://developer.nvidia.com/blog/strategies-for-maximizing-data-center-energy-efficiency/).
46. Solomon, Steve, "Sustainable by Design: Next-Generation Datacenters Consume Zero Water for Cooling," Microsoft Cloud Blog, 9 December 2024. (https://www.microsoft.com/en-us/microsoft-cloud/blog/2024/12/09/sustainable-by-design-next-generation-datacenters-consume-zero-water-for-cooling/?mscm=1)
47. Smith, John, "Data Centers Feel the Power Density Pinch," Electronic Design, 2024. (https://www.electronicdesign.com/technologies/power/whitepaper/21170904/electronic-design-data-centers-feel-the-power-density-pinch)
48. Kalyanaraman, Prasad, "AWS Announces New Data Center Components to Support AI Innovation and Further Improve Energy Efficiency," Amazon Press Center, 2 December 2024. (https://press.aboutamazon.com/2024/12/aws-announces-new-data-center-components-to-support-ai-innovation-and-further-improve-energy-efficiency)
49. Kalyanaraman, Prasad, "AWS Announces New Data Center Components to Support AI Innovation and Further Improve Energy Efficiency," Amazon Press Center, 2 December 2024. (https://press.aboutamazon.com/2024/12/aws-announces-new-data-center-components-to-support-ai-innovation-and-further-improve-energy-efficiency)
50. Foy, Kylie, "AI Models Are Devouring Energy. Tools to Reduce Consumption Are Here, If Data Centers Will Adopt," MIT Lincoln Laboratory, 22 September 2023. (https://www.ll.mit.edu/news/ai-models-are-devouring-energy-tools-reduce-consumption-are-here-if-data-centers-will-adopt)
51. Smith, John, "Data Centers Best Practices," Pacific Gas and Electric Company, Page 10, 2024. (https://www.pge.com/assets/pge/docs/save-energy-and-money/rebate-and-incentives/DataCenters_BestPractices.pdf)
52. Nellis, Stephen, Cherney, Max A., "Nvidia CEO Huang says chipmaker well positioned for shift in AI," Reuters, 18 March 2025.

EY | Shape the future with confidence

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multidisciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

Ernst & Young LLP is a client-serving member firm of Ernst & Young Global Limited operating in the US.

About EY-Parthenon

Our unique combination of transformative strategy, transactions and corporate finance delivers real-world value — solutions that work in practice, not just on paper.

Benefiting from EY's full spectrum of services, we've reimagined strategic consulting to work in a world of increasing complexity. With deep functional and sector expertise, paired with innovative AI-powered technology and an investor mindset, we partner with CEOs, boards, private equity and governments every step of the way — enabling you to shape your future with confidence.

EY-Parthenon is a brand under which a number of EY member firms across the globe provide strategy consulting services. For more information, please visit ey.com/parthenon.

US SCORE no. 27205-251US_2
ED None

ey.com

EY-Parthenon energy and resources strategy consulting teams help energy companies enhance their value today while successfully reshaping their businesses to meet tomorrow's energy needs. Our energy strategists apply the latest technology, leading business practices, state-of-the-art digitization, mergers and acquisitions (M&A) experience and more to help your business transform and thrive.

Contacts

Tom Flaherty
Senior Advisor
EY-Parthenon, Ernst & Young LLP
tom.flaherty@parthenon.ey.com
+1 214 208 7129

Gordon Bell
Principal
EY-Parthenon, Ernst & Young LLP
gordon.bell@parthenon.ey.com
+1 415 894 8727

Bill Bradley
Managing Director
EY-Parthenon, Ernst & Young LLP
bill.bradley@parthenon.ey.com
+1 512 542 7618

Jag Mukherjee
Managing Director
EY-Parthenon, Ernst & Young LLP
jag.mukherjee@parthenon.ey.com
+1 571 278 9017

Jeffrey Miller
Power & Utilities Sector Leader
EY-Parthenon Americas
jeffrey.miller2@parthenon.ey.com
+1 404 817 5879

Stephanie Chesnick
Power & Utilities Sector Leader
EY Americas
stephanie.chesnick@ey.com
+1 713 750 8192