orrick

Megawatts to Megabytes:

Orrick's Business and Legal Guide to Developing, Financing, and Powering Tomorrow's Data Centers

JULY 2025

Table of Contents

Introduction

Data centers sit at the intersection of three megatrends shaping the future: exponential AI growth, rapid scaling of low carbon power solutions, and massive capital deployment needed to develop the supporting infrastructure.

The rapid expansion of data centers across the globe is reshaping the landscape of digital infrastructure, driven by the surge in digital demands, particularly from AI and cloud computing services. This transformation presents both **opportunities and challenges** for data center developers, power producers, and investors. Understanding the interplay between energy solutions, infrastructure development, and investment strategies is critical for stakeholders aiming to capitalize on the data center revolution.

- **Data center developers** must manage the risks that accompany rapid scaling — navigating complex regulatory frameworks, securing increasingly scarce interconnection capacity to the electric and digital grids, and responding to heightened scrutiny over energy use, water consumption, and land impacts. Developers are being called upon to innovate at the intersection of operational resilience and sustainability, balancing proximity to customers and network infrastructure with the need for firm, scalable, and lower-carbon power.

- **Energy developers** are confronting the challenge of meeting soaring data center power demands, which, per the International Energy Agency, are projected to increase by over 900 terawatt hours (TWh) around the world by 2030. Hyperscale data centers—requiring 100+ megawatts (MW) of power — are at the forefront of this demand, driving the need for innovative solutions such as onsite generation technologies and new utility tariffs. This moment presents a unique opportunity to deliver transformative, low-carbon energy systems that support 24/7 reliability while meeting climate and regulatory goals.

- **Investors** in the data center space are witnessing a diversification of financing strategies that reflect the sector's maturity and appeal as a robust asset class. Traditional financing methods are evolving, with securitizations, project financings, and joint ventures becoming increasingly prevalent. These approaches offer flexibility in managing capital-intensive projects and leveraging stable revenue streams from long-term leases with high-credit tenants. Understanding and aligning these financial mechanisms with risk, credit, and sustainability considerations is essential for long-term success.

Orrick's global team of lawyers are deeply engaged in helping clients develop, power, and finance data centers and related infrastructure. This report provides a comprehensive, practice-oriented analysis of the legal, commercial, and regulatory issues shaping the future of the industry — and the solutions emerging to meet them.

1 Reputational Considerations and Social License to Operate

The digital demands of AI deployment — requiring land, water, and energy resources — have brought renewed scrutiny over the environmental and community impacts of data center developments. But concern about the rapid growth and scale of data centers, their energy-intensity and potential economic and environmental impact is not new. In 2012, Greenpeace published the report "How Clean is Your Cloud?", assigning data center operators "Company Scorecards" that applauded companies with renewable energy goals and challenged others to action. The result? Data center operators became early adopters of sustainability targets, setting bold carbon commitments and goals. Google, Microsoft, Amazon, Equinix, Digital Realty, Iron Mountain, Switch and many others publicly stated goals to achieve 100% renewable energy supply (some already meeting those goals) and have been among the first movers helping to innovate deal structures to achieve these goals.

However, as one data center client commented, **"the challenges for data center development are simultaneously global and intensely local."** In addition to company-wide, global commitments, data center companies must also engage local communities to ensure they are addressing local impacts, needs, and priorities. In other words, they must obtain and maintain a "**social license to operate**." Unlike formal regulatory approvals, a social license is an informal agreement that reflects the community's acceptance and support of a project. Achieving this license requires **early, proactive engagement with local stakeholders** to address concerns, share community benefits, and demonstrate the company's commitment to being a good neighbor.

Failure to adequately address community concerns can **delay or prevent project development**. In Ireland, concerns over the environmental impact and energy consumption have led to increased requirements to qualify for grid interconnection and a de facto moratorium on new data center connections in Dublin. Similarly, Amsterdam imposed a temporary ban on new data centers in 2019 due to worries about their strain on the power grid and urban space. Data center operators have an opportunity to export lessons learned from these restrictions to ensure they are employing best practices that will foster positive relationships, ensure the long-term viability of their projects, and the positive contribution of their developments to local communities.

Some of the key developmental challenges addressed in detail below are not surprisingly the same issues scrutinized by local communities:

- **Energy Use:** Developers may face community opposition due to data centers' massive energy demands, both because communities fear new data centers will impact their own access to electricity and because increased energy demands may impact climate goals. To address these concerns, data center operators may procure energy from renewable or carbon-free energy sources directly or through the local utility while also implementing energy-efficient technologies. [See Section 2.5 on Contracting for Power.]

- **Utility Cost-Shifting:** Some communities are concerned that people and businesses unrelated to data centers will be required to pay for new energy generation resources to support data centers. Data center developers can ease these concerns by contracting for power directly, or by using utility-tariffs to pay for new energy projects. [See Section 2.5 on Contracting for Power.]
- **Water Use:** Data centers have historically relied on water-intensive cooling systems, creating concerns in regions facing water scarcity or drought conditions. To address these concerns, operators are increasingly adopting innovative cooling technologies, such as air cooling, liquid immersion cooling, closed-loop direct-to-chip liquid cooling, or use of recycled water systems. By minimizing water waste and optimizing cooling efficiency, data centers can reduce waste, reduce their environmental footprint, and address the needs of the communities they operate in. Collaborating with local water authorities and investing in water conservation initiatives can further alleviate community concerns. [See Section 3.1 on Lease Considerations.]
- **Land Use:** Some communities are worried about data center developers buying large parcels of real estate. Developers should proactively engage with local stakeholders to understand whether the community will mount future zoning and permitting challenges. [See Section 3.1 on Lease Considerations.]

While some of the concerns about data center impacts — energy and water intensity, for example — are universal, each community is unique and any initiatives, agreements or partnerships should be done in collaboration with the community to ensure alignment with local priorities and needs. By establishing trust early and maintaining transparency and accountability in operations, data center operators can build a robust social license to operate, reducing the risk of opposition and ensuring benefits to both the data center and the community.

2 Powering Data Centers

According to McKinsey's analysis, by 2030, U.S. data center power demand alone is expected to rise by 400 TWh, growing at 23% annually, and could represent 30%–40% of new net demand. This demand would require $500 billion in infrastructure investment. The size and scale of data centers are constantly expanding, with "hyperscale" data centers that require at least 100 MW or more of capacity accounting for around 41% of worldwide data center capacity. Hyperscale data center capacity is projected to exceed 60% by 2029 (Synergy Research Group). The rapid expansion of data centers is driving unprecedented energy demands, necessitating very significant strategic investments and innovative utility tariff designs to help streamline regulatory processes and ensure that the rapid growth can be met by dedicated clean resources.

To meet their customers' needs, data center operators and users require 24/7 energy deliverability, fast ramping capabilities for peak demand periods, and high redundancy. With their unique power demands, the rapid growth of data centers is raising concerns among utilities, particularly in states lacking customer choice, as new resources are needed to meet this demand. Specifically, utilities cite concerns about integrating these new loads into their systems while maintaining reliability, affordability, and sustainability. They fear price increases for non-benefiting customers and the potential of stranded assets due to technological changes or a drop in data center demand. Moreover, with many data center users pursuing low-carbon or carbon-free goals, there is an added challenge in meeting data center growth with qualifying resources. To achieve these multifaceted goals, data center-focused clean utility tariffs may facilitate the expedited development of dedicated clean energy resources, equitable financing mechanisms, and appropriate cost allocation structures.

2.1 Key Considerations for Data Center Developers and Grid Capacity

Data centers are among the most energy-intensive infrastructure assets — and demand not just large volumes of power, but near-perfect reliability. The industry standard of **"Five 9s" (99.999% availability)** permits only about five minutes of downtime per year, making fully firm, uninterrupted grid supply essential. Most developers pursue grid interconnection capacity sufficient to always meet peak demand, typically through utility upgrades or direct transmission investment. Depending on region and scale, this process can take months or years for hyperscale projects. But that model is rapidly becoming unworkable.

Data centers are among the most energy-intensive infrastructure assets — and demand not just large volumes of power, but near-perfect reliability.

A System Under Strain

The explosive growth of data center demand, combined with the broader electrification of transport and industry, is pushing transmission infrastructure to its limits. In many jurisdictions, delays for new connections now stretch into the **decade-long range**. The UK, for example, has seen 10+-year wait times in some regions. In other parts of Europe, developers must secure new dispatchable generation or storage to qualify for interconnection capacity.

Many governments recognize the urgency. Yet policy reform and grid expansion are slow-moving by nature. Many developers simply cannot wait — especially given the speed at which AI and digital services are expanding.

Data Center Developer Dilemmas: What Are the Options?

Faced with grid constraints, data center developers are increasingly forced to choose between:

- **Waiting** for long-delayed transmission upgrades.
- **Building onsite generation** to compensate for non-firm grid supply.
- **Constructing fully islanded sites** with 24/7 self-generation.
- **Relocating** to regions where capacity is still available.
- **Acquiring grid access on the secondary market**, such as purchasing decommissioned industrial land with legacy connections.

For developers committed to proximity — whether for latency, fiber infrastructure, or customer requirements — onsite generation often becomes the most attractive option. But it brings trade-offs and complexity.

2.2 Grid Interconnection Strategies: Co-Located Generation vs. Network Load Service

Grid capacity constraints are redefining data center development. While firm grid interconnections remain the ideal, constrained transmission infrastructure is forcing developers to consider complex alternatives. Similar to co-generation facilities previously popular for industrial manufacturing sites, onsite gas generation may be a necessary first step in some cases, but it must be **carefully structured and future-proofed** to align with both commercial objectives and carbon reduction imperatives.

The winners in this space will be those who can navigate **technical, regulatory, and reputational complexity** — and bring resilient, scalable power solutions to the front lines of the digital economy.

Connecting hyperscale data center load to the interstate transmission grid is an increasingly complex, costly, and time-consuming process. Developers must carefully weigh the trade-offs between direct interconnection as a "network load" and pairing with behind-the-meter generation. Each path offers unique benefits and risks across timing, reliability, regulatory treatment, and cost recovery.

The winners in this space will be those who can navigate technical, regulatory, and reputational complexity — and bring resilient, scalable power solutions to the front lines of the digital economy.

Network Load Service: Direct Interconnection Without Co-Located Generation

In the traditional model, a data center connects directly to the transmission grid and is designated by the utility as a **network load**. The transmission provider studies the service request and, if upgrades are needed, finances them upfront and recovers costs from data centers through its transmission tariff. The data center developers may be required to post security for the upgrades, which is typically refunded upon energization.

This approach often has **shorter study timelines** than those involving new generation, but also presents notable downsides:

- **No dedicated power generation**, leaving the data center fully reliant on grid conditions.
- **Curtailment risk** during transmission outages or congestion.
- **Potentially higher wholesale energy costs**, especially if located far from generation centers.

Co-Located Generation with Grid Interconnection

Co-located generation involves pairing a data center with a new or existing power plant. Co-located generation provides enhanced reliability and may accelerate energization by enabling the data center to interconnect using the generator's existing or pending interconnection agreement. This structure can **shield the data center from curtailment** and offer greater control over energy sourcing. However, it brings significant regulatory and practical complexity.

- Interconnection of new generation is **significantly delayed** across many regional transmission organizations (RTOs) due to study backlogs.
- The **standard timeline** for completing an interconnection agreement is approximately three years from the initial request to the final agreement, with delays stretching to more than **six years**.

RECENT REFORMS

In 2023, the Federal Energy Regulatory Commission (FERC) introduced a sweeping overhaul of the interconnection process to address these delays:

- Transition from "first-come, first-served" to a **"first-ready, first-served" model**.
- Projects must demonstrate **90% site control** and post a **commercial readiness deposit** equal to 10% of expected network upgrade costs.
- FERC began imposing **penalties on transmission providers** for missed study deadlines.

While implementation of these reforms is underway, the **benefits may take years** to fully materialize.

Accelerated Paths: Surplus and Replacement Rights

To bypass long interconnection queues, data center developers are exploring alternative interconnection mechanisms:

- **Surplus interconnection rights:** This approach involves leveraging underutilized capacity from existing generators. For example, where an aging coal plant operates at only 20% capacity, a developer may co-locate a new solar plant to supply the remaining interconnection capacity. By attaching to an existing interconnection agreement, the solar plant skips the interconnection queue. Although this approach may offer lower costs and faster approval, the surplus interconnection is **subordinate to the primary interconnection** — if the host loses its rights, so does the surplus project.
- **Replacement interconnection rights:** Under this approach, a proposed generator can **inherit interconnection rights from a retiring facility**, provided it does not exceed the prior injection level. This approach may avoid the need for full studies and allow for fast-track development. However, it may require **joint ventures** or transition agreements with the retiring asset owner to prevent queue jumping.
- **Energy parks:** This approach involves integrating multiple generation assets, storage solutions, and co-located loads behind one point of interconnection as a form of large-scale microgrid. Energy parks offer data centers cost savings and faster energization; however, integrating energy parks into the grid may require changes to current market rules, and they may be difficult to finance given their size and the large number of parties and technologies involved.

2.3 Selection of a Power Technology

The energy sources and technologies available to power data centers present several important considerations, including reliability, scalability, cost, regulatory compliance, and environmental impact, among others. Selecting a suitable power technology requires an assessment of the relative benefits and drawbacks of each option.

We note that, while not specifically addressed in this report, power facilities constructed to serve data centers will require their own financing. Project financings and, to the extent applicable, tax credit monetization transactions are often the predominant financing structure for such projects. Data center developers will need to consider this as they source their power supply.

ONSITE GAS: A VIABLE BUT COMPLICATED OPTION

Natural gas-fired generation is often the best onsite solution for firm, dispatchable power. However, while cleaner than coal or oil, it still produces emissions and may conflict with publicly stated carbon goals, exposing developers to reputational and regulatory risks.

How gas is deployed significantly affects its classification:

- Direct supply to the data center typically results in scope 2 GHG emissions, indirect emissions as the result of purchased energy, which can be offset using renewable energy certificates (RECs or Guarantees of Origin), particularly if the generator is not owned or operated by the data center itself.
- However, direct linkage increases public visibility and reputational exposure, especially for sustainability-minded operators.

Balancing interconnection needs, carbon reporting and public commitments becomes a strategic challenge. Solutions will depend on a developer's risk appetite, sustainability goals, and investor expectations.

1. Natural Gas

Natural gas-fired power plants are a **critical tool for ensuring reliable, dispatchable, large-scale energy delivery** — especially for hyperscale data centers that require hundreds of megawatts of capacity with near-perfect uptime. For decades, natural gas has served as a cornerstone of industrial power systems, and it remains one of the few technologies capable of balancing scalability, geographic flexibility, and dispatchability.

For developers and investors looking to meet immediate data center power needs, natural gas offers a distinct combination of advantages:

- **Reliability and Dispatchability:** Natural gas plants provide stable, on-demand baseload and peaking capacity, essential for mission-critical data operations.
- **Siting Flexibility:** Plants can be developed near data center campuses in industrial zones, avoiding lengthy transmission buildouts or congested grid interconnections.
- **Scalability:** Natural gas facilities can be built to match the 50–500+ MW scale now typical of AI and hyperscale campuses.
- **Lower Emissions Compared to Legacy Fuels:** Compared to diesel or coal, natural gas emits significantly less CO_2, NO_X, SO_2, and particulates — making it a cleaner option for near-term deployments.
- **Combined Heat and Power (CHP):** Recovered thermal energy can reduce a data center's cooling energy demand, increase total system efficiency, and improve project economics.

Developers and operators are also leveraging **next-generation gas technologies and efficiency upgrades** to reduce emissions and increase performance. Today's plants can be equipped with:

- **Combined-Cycle Gas Turbines (CCGT)**, which improve efficiencies up to 60% by capturing and reusing waste heat.
- **Selective Catalytic Reduction and low-NOx** burners to reduce air pollutants.
- **Carbon Capture and Sequestration (CCS)** systems to mitigate CO_2 emissions at the point of combustion — an increasingly important (but expensive) feature for regulatory compliance and ESG performance.

These innovations allow developers to deploy natural gas solutions that meet current emissions requirements, and they're also better positioned for a future energy mix that includes hydrogen blending, renewable fuels and low-carbon operational mandates.

Despite its strengths, natural gas power presents several challenges that must be addressed through proactive planning and risk mitigation:

- **Equipment Procurement Lead Times:** Global supply constraints and high demand have led to multiyear lead times for major components like gas turbines and heat recovery steam generators (HRSGs). Developers may need to secure procurement contracts early in project cycles.
- **Fuel Supply and Infrastructure:** Viable projects require access to high-pressure gas lines, compression infrastructure, and contingency fuel strategies. Siting must be coordinated with pipeline networks and local permitting authorities.
- **Commodity Price Risk Management:** Data center developers looking to utilize natural gas power solutions may need to bear (at least partially) commodity price risk of natural gas and may consider hedging and other contractual solutions to mitigate the short- and long-term risks.
- **Environmental and Regulatory Risks:** While natural gas-fired generators emit fewer CO_2 emissions (especially as compared with diesel backup generators that are often used in data center power systems), they are not considered a long-term clean energy solution unless paired with CCS or renewable fuels. Additionally, incidental releases or "leaks" of natural gas in transportation emit methane (CH4), which is a more potent greenhouse gas compared to CO2. Projects may face regulatory scrutiny or declining policy support in jurisdictions with aggressive decarbonization targets.

FUTURE-PROOFING: FROM GAS TO ZERO-CARBON

To mitigate long-term risks, developers considering natural gas must look ahead. Projects should incorporate clear pathways for transitioning to lower-carbon fuels like green hydrogen or implementing CCS as soon as technology and economics allow. Structuring these projects to evolve with policy, technology, and investor requirements may become critical to preserving optionality and the social license to operate.

2. Nuclear

Nuclear power has provided stable baseload power to the grid for decades and presents a compelling option for data centers seeking stable carbon-free energy solutions. Nuclear power, with its ability to provide consistent baseload power, offers significant advantages in meeting demand for clean, reliable baseload power.

Large-Scale Nuclear Power

Traditional nuclear power plants have long been recognized for their capacity to generate substantial amounts of electricity with minimal carbon emissions. A single traditional nuclear reactor typically generates approximately **one gigawatt of electricity** with availability of over 90%. These facilities are a natural fit for hyperscale data centers, which often require hundreds of megawatts to operate efficiently. However, the substantial capital expenditure and construction time associated with building and permitting new traditional nuclear power plants is an obstacle to development.

In the last 20 years, only three new traditional nuclear power plants have been built and commissioned in the United States: Watts Bar Unit 2 (2016), Vogtle Unit 3 (2023), and Vogtle Unit 4 (2024). As a result, data centers primarily focus on entering power purchase agreements that support **recommissioning** nuclear power plants previously decommissioned or **extending** the life of currently operating nuclear power plants. Given the limited number of options available for recommissioning and extensions, many data centers are instead focusing on small modular reactors for future development plans.

Constellation recently announced significant nuclear-powered data center agreements with leading technology companies:

- **Microsoft:** Constellation and Microsoft entered into a groundbreaking agreement under which Microsoft will procure power matched on an hourly basis from Constellation's nuclear fleet. The deal involves delivering carbon-free electricity to Microsoft's Virginia data center operations and is one of the largest 24/7 carbon-free nuclear power agreements in the U.S. to date.
- **Meta:** Meta signed a major agreement with Constellation to supply nuclear power for its data centers in the PJM region. This transaction is designed to match Meta's hourly data center load with nuclear energy, supporting the company's goal of achieving net-zero emissions across its value chain.

These transactions signal a growing trend among hyperscalers toward leveraging nuclear energy — particularly existing assets — for meeting real-time carbon-free energy targets.

Small Modular Reactors (SMRs)

SMRs represent an innovative approach to nuclear energy, offering a more flexible and scalable solution for data centers. These reactors are smaller in scale, ranging from 1 to 10 MW (also known as microreactors) to approximately 350 MW of output. Their smaller size results in lower capital expenditures than traditional nuclear power reactors. They may also be deployed in a modular fashion, standardizing construction and operation across multiple projects and allowing a single site to host multiple units built together or sequentially. The proposed designs of SMRs incorporate advanced safety features including passive cooling systems and the use of advanced nuclear fuels designed to avoid reactor malfunctions and core compromises. These advanced safety features mean that SMRs can have smaller physical footprints compared to traditional nuclear power plants, which require large exclusion zones under current regulatory standards. This allows SMRs to be located closer to data centers and supports future behind-the-meter deployment.

SMRs deploy existing fission reaction technology but are generally considered first-of-a-kind from a regulatory, financing, and market perspective. In recent years, developers have been forced to navigate challenging and complex regulatory landscapes. However, nuclear power development currently benefits from bipartisan political support as a carbon-free resource with substantial economic benefits. In fact, Congress recently enacted the ADVANCE Act, which directs the U.S. Nuclear Regulatory Commission to reduce licensing application fees and expedite the licensing process. The current U.S. administration also enacted several executive orders to streamline the review and approval process for new nuclear reactors and provide other forms of government support. While the present regulatory scheme is a challenge to SMR development, the evolving political landscape is expected to reduce regulatory burdens and support SMR innovation.

3. Geothermal

Geothermal power — unlike many other renewable energy sources — can serve as a **baseload resource** with **24/7 firm availability**. This makes it uniquely suited for powering critical infrastructure such as data centers, which require constant, reliable electricity regardless of weather or time of day.

According to the U.S. Department of Energy, geothermal energy could supply up to **120 GW** of generation capacity in the U.S. by 2050 — enough to meet over **16% of projected national electricity demand**. This long-term vision reflects not only the environmental benefits of geothermal energy, but also its role in stabilizing grids with a growing share of variable wind and solar generation.

Geothermal is gaining traction in regions where traditional renewables face siting, land use, or intermittency challenges — particularly in the **Asia-Pacific region**, where population density and grid reliability constraints complicate solar and wind deployment. Unlike solar and wind farms, geothermal plants have a compact footprint and minimal visual impact, making them more compatible with urban or industrial zones.

However, geothermal energy has historically been underutilized due to high upfront capital costs and geological limitations. Traditional geothermal technologies are economically viable only where high-temperature resources are easily accessible near the surface. That landscape is now shifting rapidly due to technological innovations. **Advances in drilling and subsurface engineering** — adapted from the oil and gas industry — are unlocking deeper and more complex geothermal reservoirs. Recent breakthroughs include:

- **Enhanced Geothermal Systems (EGS):** Using techniques such as hydraulic stimulation to create artificial reservoirs in otherwise dry hot rock.
- **Horizontal Drilling:** Increasing reservoir contact and thermal output.
- **Closed-Loop and Modular Geothermal:** Innovative systems that do not require water-intensive open reservoirs.

These advancements should extend the geographic viability of geothermal power into areas previously considered uneconomic or geologically unsuitable.

Another accelerating trend is the co-location of data centers with geothermal resources, particularly in regions with supportive regulatory frameworks and high energy demand. Tech companies are leading the way:

- **Google and Baseload Capital (Taiwan)**: Google signed a corporate Purchase Power Agreement (PPA) to support a new 10 MW geothermal facility to power its Taiwanese data center and chip manufacturing operations — marking one of the region's most advanced clean energy integrations.
- **Google and NV Energy + Fervo Energy (Nevada)**: In the U.S., Google is pursuing an innovative arrangement where up to 115 MW of geothermal energy from Fervo Energy will be used to power its growing data center load in Nevada. Fervo's use of fiber-optic sensing and horizontal drilling is setting new standards for geothermal monitoring and efficiency.

Beyond electricity generation, geothermal energy can also be used to support **direct cooling of data centers** using **geothermal heat pumps** or absorption chillers. These systems can significantly reduce electricity demand for cooling (often 30%–40% of total data center load) as well as **water consumption**, a growing sustainability and regulatory concern in many regions.

4. Solar and Wind Resources

Renewable energy resources, such as wind and solar, are commonly used to power data centers. However, data centers may not be able to rely solely on these generation sources for 24/7 uptime due to the inherent intermittency of wind and solar generation.

To achieve reliability, developers may need to supplement wind and solar generation with:

- Deliveries of conventional power from the market;
- Diesel backup generators;
- Battery storage that stores renewable energy to use at times when renewable energy generation is unavailable; or
- In some cases, hydropower, nuclear or geothermal.

There is also a business case for utilizing existing wind or solar projects as sources of power for data centers. Typically, during periods of negative pricing in ISO regions, renewable energy producers may choose to curtail generation instead of paying the grid operator to generate in the negative price environment. However, by co-locating a data center at the wind or solar site, this excess power can be directed to power the data center load rather than being curtailed due to negative pricing. In this type of co-located arrangement, the data center remains connected to the grid to ensure a continuous power supply when the renewable source is not generating. This approach provides a win-win situation, both optimizing the renewable energy resource and maximizing revenue for the renewable energy producers who otherwise may have curtailed the supply.

PHYSICAL DELIVERY – 24/7 CARBON-FREE DELIVERY STRUCTURE

Because wind and solar power are intermittent resources, data centers requiring 24/7 power leverage a mix of resources to meet their load requirements on a 24/7 basis.

One example is **Microsoft's** 24/7 carbon-free delivery structure to power its data center in the State of Washington. The agreement involved several underlying renewable PPAs in which Microsoft purchased physical renewable power from developers and an innovative energy management agreement in which an independent power producer combined its own hydropower resources with the renewable resources to meet the data center load. The parties had to balance the imperative that the data center never went without power with the intermittency of the carbon-free resources. To achieve this, the parties specified a hierarchy of delivery from the carbon-free resources and carefully negotiated scheduling provisions to account for all contingencies.

5. Fuel Cells

Fuel cells are a clean energy technology that generates electricity through an electrochemical reaction — typically using hydrogen and oxygen — without combustion. This process results in significantly lower emissions compared to traditional fossil-fuel-based generation, with water and heat as the primary byproducts. As the global demand for digital infrastructure surges, fuel cells are emerging as a **scalable, dispatchable power solution for data centers**.

In some cases, fuel cells provide continuous, behind-the-meter power directly to data centers, reducing reliance on the grid and avoiding the risks associated with transmission congestion or interconnection delays. They can also serve as an alternative to diesel generators for providing backup power during grid outages, offering faster start-up, lower emissions, and compliance with increasingly strict air quality regulations.

Fuel cells offer **high reliability** — often with availability factors above 99.9% — making them well-suited to meet the always-on power requirements of hyperscale and co-location data centers. Additionally, data center developers often prioritize fuel cell systems for their relatively fast deployment timelines and easier permitting compared to traditional generation. This makes them **strategically valuable as a "bridge" resource**, providing interim power for several months or years while larger permanent grid infrastructure is built or upgraded.

From an investment standpoint, fuel cells can help mitigate the power delivery bottlenecks that have delayed or constrained new data center developments in key U.S. and European markets. In regions where grid interconnections are projected to take 3–5 years or longer, fuel cell installations can enable earlier revenue generation from data center assets — improving project IRRs and unlocking portfolio value.

Major technology firms and data center operators have deployed fuel cells at scale:

- **Equinix** has partnered with **Bloom Energy** to install fuel cell systems at 19 of its U.S. data centers, with a total capacity exceeding 100 MW. These systems operate continuously, offsetting grid power and contributing to Equinix's long-term sustainability goals.
- **Microsoft** has announced a partnership with the **Electricity Supply Board** in Ireland to power a forthcoming data center with renewable energy powered hydrogen fuel cells, positioning the project as one of the first large-scale deployments of zero-carbon fuel cell technology in Europe.
- Additional pilot projects by **Amazon**, **Google**, and **Meta** are exploring both natural gas-powered and green hydrogen fuel cell technologies as part of broader efforts to decarbonize data center operations and enhance grid independence.

Key considerations for data centers looking to deploy fuel cells include:

- **Technology Maturity:** While natural gas-based solid oxide fuel cells are commercially mature, hydrogen fuel cells are still gaining traction as green hydrogen becomes more available.
- **Policy and Incentives:** Federal and state-level incentives enhance the economics of fuel cell projects, particularly those utilizing low-carbon or renewable fuels.
- **Siting and Permitting:** Modular design and lower emissions profiles make fuel cells easier to permit than diesel generator sets or gas turbines — often enabling deployment in urban or constrained areas near data hubs.
- **Revenue Models:** Opportunities exist in long-term energy-as-a-service (EaaS) agreements with data center operators, as well as potential grid services revenues through demand response or capacity markets.
- **Environmental Considerations:** While hydrogen fuel cells create electricity without emitting any CO2, overall emissions depend on whether the hydrogen is produced using renewable energy (green hydrogen) or natural gas (blue hydrogen).

6. Energy Storage

The case for utilizing energy storage — including battery energy storage systems (BESS) and other storage technologies — to help manage the power supply equation for data centers is compelling. BESS can provide the following functionalities, depending on ISO/RTO status and state regulatory requirements, and subject to the constraints and challenges described below:

- **Load-Matching:** If co-located with solar, gas, or other power generation resources, BESS can potentially be used to help **address data center power demand**. A BESS with a four-hour duration, however, is not likely to satisfy the 24/7 load-matching goals of hyperscale data centers. Long-duration energy storage (LDES) solutions, although costly, are considered viable alternatives — and multiple LDES technologies and vendors have commercially viable projects already in operation.
- **Time-Shifting:** A stand-alone BESS, which is a battery interconnected to the grid without a co-located generation resource, can potentially be used to **decrease a data center's energy costs** by arbitraging market prices (charging when local marginal prices (LMPs) are low and discharging when LMPs are high). The viability of this approach depends on whether the data center is in an ISO, whether the BESS is behind or in front of the data center's revenue meter, and other regulatory considerations.
- **Reliability:** A BESS placed behind the data center's revenue meter can serve a reliability function, including during power failures. However, installing a BESS system is costly, and the relatively short duration of most BESS reduces its attractiveness.

Although the potential for utilizing BESS to support data center requirements exists, not many BESS-data center transactions have occurred to date. The relative costliness of BESS (combined with recent uncertainty in U.S. tax credit — now resolved with President Trump's signing of the "One, Big, Beautiful Bill" on July 4, 2025 — and tariff policies), short BESS durations, and regulatory considerations are all contributing factors. In addition, the inconsistency across the U.S. of mature capacity markets, and changing capacity accreditations for BESS have resulted in similarly inconsistent incentives for BESS deployment. As data center power demands increase and U.S. capacity markets develop in the coming years, however, we believe that mature BESS technologies are poised to play a critical role in addressing these demands.

7. Other Technology Solutions

As demand pressures on the grid increase — due in part to the proliferation of data centers — and obstacles to expanding the transmission network remain significant, grid-enhancing technologies have become critical tools to increase the existing network's capacity, reliability, and efficiency. These technologies encompass a variety of innovations, such as:

- **Sensor Technologies:** By deploying sensing and monitoring devices throughout transmission and distribution systems, utilities can gather real-time data on condition and performance. These devices enable early detection of operational issues, leading to improved reliability and efficiency and reduced maintenance costs. As extreme weather events become more frequent due to climate change, these technologies are increasingly vital for monitoring the stability of transmission and distribution infrastructure.
- **Dynamic Line Rating (DLR):** DLR systems are installed along transmission lines and monitor real-time weather conditions and other environmental

factors. In response to these conditions, the DLR technology modifies the rating assigned to the line. As a result, DLR technology provides data on a transmission line's real-time power carrying capacity, as well as forecasted carrying capacity. This granular data helps to maximize existing transmission infrastructure by ensuring energy generators fill remaining capacity.

- **Flexible Alternating Current Transmission Systems (FACTS):** FACTS technology improves network efficiency by increasing control over power flow, reducing power loss, and maintaining power quality. Distributed FACTS (D-FACTS) is a scaled-down version of FACTS with the benefits of lower cost and easier installation.
- **System-Wide Load Flexibility:** System-wide load flexibility refers to the power grid's ability to integrate large, controllable energy demands — such as data centers or electric vehicle charging — by allowing these loads to temporarily reduce or shift their electricity usage during periods of grid stress. This approach could enable the addition of substantial new data center loads without necessitating significant infrastructure upgrades, thereby maintaining grid reliability and affordability; however, it requires that data center operators and hyperscalers are willing to lose power supply for around 40 hours each year (0.5% of the time) when power supply is most constrained, which is well above the 99.999% reliability that the industry strives for.

2.4 Contracting for Power

To meet the unique energy demands of data centers, there are a variety of contractual arrangements and mechanisms that provide strategic opportunities for securing power while achieving environmental goals.

The Evolution of Utility Tariffs to Serve Data Center Load

Behind-the-meter strategies cannot work for all locations as a result of geographical, technical, and regulatory impediments. As a result, data centers connecting directly onto the grid require solutions to integrate large demand loads into utility tariff frameworks.

Where the existing grid does not have sufficient generation or transmission capacity to support the interconnection of an additional data center, the data center developer will often be required to absorb resulting financial and contractual risk. The utility is often not in a position to provide the tens or hundreds of megawatts of power required to power the data center without the build-out of additional dedicated generation and transmission capacity.

In these circumstances, the **data center developer may be required to source and deliver new generation projects to the utility** and/or in parallel, to pay the utility to construct expensive and lengthy transmission upgrades — all at the data center developer's cost and risk, before any tenant is signed up, and before any sleeved or dedicated PPA is entered into between the generator and the utility. To do so, the data center developer will require significant balance sheet support or third-party financing to absorb these risks, creating a number of commercial and legal issues to resolve.

Many of these arrangements between data center load and utilities are negotiated on a bilateral, one-off basis. However, as data centers proliferate, we expect that more utilities will adopt uniform tariffs to provide a more consistent approach to incorporating data center load. Ideal tariff structures balance data center priorities for reliable, affordable, and clean electricity, with the utility's priority to mitigate the costs and risks related to building large, new energy resources to power data centers.

Green Tariffs

Many data center developers and customers have ambitious clean energy pledges that require they generate onsite carbon-free electricity, or else purchase clean energy attributes — i.e., Renewable Energy Credits (RECs) — which tie the data center's electricity usage to a renewable energy project's generation.

Green Tariffs allow data centers to offset their overall electricity consumption with RECs or buy bundled renewable electricity (electricity paired with RECs) from a specific project. Green Tariffs allow data centers to match their overall electricity usage with renewable energy resources, but they do not necessarily match real-time energy consumption with renewable energy generation.

24/7 Clean Utility Tariffs

Some markets are moving toward a more sophisticated clean energy procurement strategy called "24/7 carbon-free energy" or "hourly matching." This strategy **requires that buyers match their electricity usage with carbon-free electricity generation on an hourly basis**, typically using resources located on the same grid where the electricity is consumed.

24/7 clean utility tariffs can facilitate hourly matching by aligning hourly data center load with dedicated, carbon-free energy generation. Unlike traditional green tariffs, 24/7 clean utility tariffs match actual consumption with real-time clean supply, improving both sustainability and reliability. Crucially, these tariffs are designed to reflect the true cost of service without imposing rate distortions that deter regional investment. However, from the utility perspective, it can be a challenge to contract for sufficient renewable generation to deliver 24/7 clean energy to customers, and so the costs of these tariff products can be significantly higher than other tariff offerings.

Clean Transition Tariffs

Clean Transition Tariffs (CTTs), which go by other various names such as Accelerating Clean Energy tariffs, go one step further by offering a **structured, transparent financial mechanism** for delivering **new, dedicated clean resources** in direct partnership with data center operators and energy developers.

CTTs are gaining traction in **non-Independent System Operator (ISO) utility regions** where utilities function as monopolies, and do not permit data centers to purchase electricity from outside suppliers. CTTs offer a breakthrough path to develop dedicated clean resources without shifting costs to nonparticipating customer classes or triggering prolonged regulatory disputes. This approach bypasses many of the bottlenecks in traditional regulatory pathways by:

- **Isolating cost responsibility** to beneficiaries, avoiding cost shifts to other customer classes.
- **Reducing regulatory friction** by clarifying financial commitments up front.
- **Securing long-term certainty** through binding agreements and defined risk mitigation mechanisms.

Well-crafted CTTs emphasize collaboration between stakeholders to ensure that the cost of service aligns with actual demand, to create predictable cost structures, and to establish financial mechanisms to allow for the expedited development of dedicated clean energy resources.

In conjunction with CTTs, innovative contractual and operational mechanisms can expedite data center deployment. These mechanisms may include:

- Financing structures that share initial investment costs and guarantees between utilities and large customers for new generators developed to serve data centers within the utility's territory.
- Encouraging data center operators to program behind-the-meter generation and storage assets, where available, to be dispatchable to interconnecting utilities under limited and defined periods, thereby contributing positively to the grid during certain system events.
- Creating incentives or requirements for demand-side management programs at data centers, such as selective load shedding or prioritized load management during defined system events (e.g., prioritizing training modules over latency-sensitive applications at a given data center).
- Establishing a low-cost standby backup service for data centers primarily served by behind-the-meter generation.

NV ENERGY'S CLEAN TRANSITION TARIFF (CTT)

NV Energy's new CTT illustrates this model in practice:

1. Available to customers with **≥ 5 MW** average hourly load.
2. Rates based on the utility's **Base Tariff General Rate**, with credits for generation because power is supplied by a dedicated resource.
3. Long-term agreements matching the **asset's lifespan**.
4. **Risk protections** via liquidated damages and security requirements.

This type of structure helps provide developers with bankable certainty, utilities with cost containment and data center customers with reliable, clean energy tailored to their needs.

"Take-or-Pay" Tariff Structures

Unlike the versatile tariff structures discussed above, "Take-or-Pay" tariff structures prioritize mitigating utility risk above all else. This structure requires that data centers pay all or a percentage of their contracted capacity regardless of actual energy use. This structure ensures data centers cover the cost of building new generation resources; however, significant fixed costs deter investment and complicate development.

Behind-The-Meter (BTM) PPAs

"Behind-the-meter" refers to co-located power plants that deliver electricity directly to an energy load without using regulated transmission lines as an intermediary. BTM powering of data centers is an attractive option for corporates and data center developers lacking consistent, affordable, or readily available grid energy. These transactions are typically negotiated in a PPA similar to the thousands of megawatts of PPAs already in existence for rooftop and commercial & industrial solar projects. But BTM data centers that do not have any utility grid connection, even for backup power, are rare, because data center developers must over-build generation to ensure that they have reliable supply at all times, resulting in significant incremental development costs.

BTM PPAs offer data centers many **advantages**:

- On-site electricity generation can help ensure reliability by providing a stable source of power disconnected from the wider grid.
- BTM PPAs may enable data centers to directly source renewable energy, helping them meet sustainability targets and reduce their carbon footprint.
- BTM generation can be customized to satisfy the specific needs and operational requirements of a data center.

However, certain unique considerations exist:

1. Because data center loads are significant, the project may require a large percentage of the buyer's land; **site arrangements and allocation of site responsibilities are key**, and the contractual interface must be evaluated.
2. Parties will need to consider the most efficient and desired outcome of the project following **a default of either party or termination of the PPA** and, in particular, the future power sales arrangement for the generator if the data center shuts down, no direct grid interconnection is available, and the generator becomes "stranded" without a load to serve.
3. **Storage can be co-located with generation** to shift energy supply when the renewables facility is generating excess output or grid energy is not sufficient to satisfy data center demand.

Build-Transfer Arrangements

A data center developer may facilitate sourcing of new generation through a build-transfer agreement (BTA).

A BTA is a **hybrid acquisition and construction contract** in which the BTA counterparty, typically a project developer, secures and initially owns the land rights, permits, interconnection rights, and all other assets necessary to construct and operate the generation facilities. They're also responsible for constructing (either directly or via a third party) the facilities. This activity all occurs for a fixed price. Afterwards, on the "closing date" under the BTA, the data center developer, as "buyer," takes ownership of the project assets, and the seller thereafter remains responsible for achieving final completion of the facility.

For its part, **the data center developer is responsible for paying the purchase price under the BTA**, typically in installments which may be structured such that the buyer is essentially providing project financing for the late-stage development and construction of the facility. If the seller or its parent is financing the development and construction, installments may include a relatively modest pre-closing closing deposit, a closing payment, and one or more post-closing installments conditioned on achievement of substantial and/or final completion payment, depending on the closing conditions.

BTA Advantages

- The data center developer can **shorten the time required to secure generation capacity** by entering into a BTA with respect to a generation facility in an advanced stage of development.
- A BTA provides an opportunity (within limits) to **customize** the generation facility to better address the data center's specific requirements, including reliability, ramping capabilities, regulatory compliance, and environmental impact, among others.
- The BTA structure is **versatile** and may be used to purchase a behind-the-meter project, a network load project, a project utilizing replacement interconnection rights, or a project interconnected to the same utility network and utilizing CTTs (where available).
- Through a BTA structure, the data center developer would eventually **own the generation facility's tax credits** (investment or production) and could monetize them via self-use, a third-party tax equity arrangement, or a third-party sale.

BTA Disadvantages

- BTAs require a **significant commitment of capital** and a willingness on the part of the data center developer to take on some development and construction risk, which may be considerable depending on the BTA terms.
- BTAs are **complex and bespoke, requiring the parties to identify and allocate various risks**, including development, equipment procurement, construction, financing, regulatory and tax risks, force majeure and change-in-law events, and the parties' rights and obligations if the BTA is terminated before closing occurs. The complexity can be compounded by introducing a joint venture structure, which may be needed to facilitate third-party tax equity monetization or to enable two or more parties to utilize the generation facility's capacity.
- The data center developer will **own and operate** (or engage a third party to operate) the project, which activities have their own costs and risks.
- Depending on the nature of the project and the applicable regulatory environment, the **project could become an expensive stranded asset** if the data center fails or the revenue-generating contracts are breached or terminated prematurely.

Virtual PPAs

As noted above in [Chapter 1: Social License to Operate], many data center companies have pledged 100% renewable energy use for their operations and have innovated the structures to obtain this goal.

One such innovation is the Virtual Power Purchase Agreement (VPPA), which offers a flexible and impactful way to achieve these objectives by allowing operators to make renewable energy claims by financially supporting renewable energy projects without requiring physical delivery of electricity. Through VPPAs, data centers procure Renewable Energy Certificates (RECs), but unlike a REC-only contract, the VPPA allows the purchaser to facilitate the development of new clean energy resources (sometimes called "additionality") and potentially benefit from the financial settlement of energy price differences.

The adoption of VPPAs presents several **advantages** for data center operators:

- Provides **price stability** by locking in a fixed rate for renewable energy over the contract term, typically 10 to 20 years.
- Enhances the operator's reputation by **demonstrating a commitment to sustainability** and supporting the transition to a low-carbon economy. This can be particularly beneficial in regions where community and regulatory pressures demand greater environmental responsibility from data centers.
- Can be **tailored to meet specific corporate goals**, such as sourcing energy from particular types of renewable projects or geographic locations, allowing data centers to strategically align their energy procurement with broader business objectives.
- Because VPPAs are financial transactions, they allow data centers to meet their renewable energy goals in regions where resource availability or regulatory or cost hurdles make **renewable procurement** difficult.

As with all power contracts, Virtual PPAs have some **drawbacks**:

- VPPAs typically involve **long-term purchase commitments**.
- Since VPPAs do not involve the physical delivery of electricity, renewable energy generated from the contract **may not directly reduce the carbon footprint** of the buyer's operations.
- Because VPPAs often involve setting a fixed price for RECs, the buyer is **exposed to market price fluctuations** and may end up paying more than the market rate.

In 2024, S&P Global estimated that data centers procured over 17 GW of clean energy through direct third-party power purchase agreements. The VPPA trend will surely continue as one of the most effective tools to navigate the complexities of renewable energy procurement while contributing to the global effort to combat climate change.

REC Transactions

Because of the physical delivery challenges associated with carbon-free or renewable energy and, in turn, the challenges to satisfy data center load, data center providers may choose to purchase RECs via one or **more long-term agreements from a renewable energy resource**. The owner of such resource can then match the RECs to the data center's load on a 24/7 basis. It should be noted that it is difficult to achieve actual 24/7 matching from renewable energy alone; however, certain protections can be built into the agreements to ensure that the matching is done at the highest possible rate.

Procuring RECs from multiple renewable resources in the vicinity of a data center and arranging for 24/7 matching is an efficient way to make unique green claims without the hindrance of physical delivery and local market constraints (i.e., not enough carbon-free sources to delivery on a 24/7 basis).

Hedging Transactions

Data centers can benefit from active energy hedging strategies, particularly in ISO-markets. Energy hedging is a financial strategy that allows companies to lock in energy costs using tools like futures, swaps, or options. These strategies enable data centers to ensure 24/7 power available to the data center facility, mitigate pricing risk caused by spikes in power pricing in given intervals (e.g., daytime summer intervals), and much like VPPAs, match utility-delivered power supply with available clean power from renewable energy resources.

Unlike bilateral PPAs (whether physical or virtual) or utility tariffs discussed elsewhere in this paper, hedges are defined by their flexibility. Specifically:

- **Tenor:** Can be structured on a short-term basis (including on a day-ahead basis) or for longer terms of 7 to 10 years (although hedges rarely approach the true long-term nature of PPAs).
- **Form:** Entered into pursuant to bilateral agreements, through commodities exchanges, or through ISOs. They can be contracted through industry-standard instruments such as the ISDA, the EEI, the WSPP, or through bespoke instruments.
- **Contracting Parties:** Providers can include not only utilities, but also independent power producers, energy storage companies, commodity trading desks, and insurance providers.
- **Products:** May include energy, capacity attributes, ancillary services, and environmental attributes/ renewable energy credits.
- **Terms:** Can include fixed-for-floating swaps (e.g., TB2/4 hedges for BESS projects), call options (e.g., heat rate call options for thermal generation) and put options (whether based on price or on revenue).

While allowing more flexibility than traditional instruments, hedges can come with their own set of **unique risks**. These may include:

- **Creditworthiness and Liquidity:** Unlike public utility-approved instruments that may be included in a utility's rate base, **hedges are ultimately backed by both parties' balance sheets combined with other forms of credit support** (e.g., letters of credit). General problems with a hedge provider's balance sheet, including as a result of issues with other projects, may affect any given hedge in its portfolio (including one with a data center provider).
- **Market Events:** Market- or weather-related events that affect physical assets may have knock-on effects under hedges. The most poignant example was during Winter Storm Uri when a large number of fixed-volume hedges in the Electric Reliability Council of Texas (ERCOT) market ultimately failed.

The sophistication of a hedging strategy is really in the mind of the conceiver. Risks to the strategy are creditworthiness, liquidity, and increased financial exposure.

2.5 Regulatory Considerations

Network Load Classification for Co-Located Projects

A growing regulatory issue concerns whether data centers with behind-the-meter generation must be designated as "network loads," requiring them to pay for transmission upgrades even if they rarely draw power from the grid.

In a 2024 decision, the Federal Energy Regulatory Commission (FERC) rejected amendments to an interconnection agreement between PJM and **Susquehanna Nuclear, LLC**, aimed at supplying an **Amazon** data center behind-the-meter. Protesters argued that the data center should be classified as a network load and share in transmission upgrade costs. FERC agreed, citing unresolved reliability concerns.

FERC has since opened a **"show cause" proceeding** to evaluate PJM's treatment of large co-located loads. FERC's decision could have **major cost and timing implications** for developers seeking to co-locate generation with data centers.

As transmission bottlenecks and policy reforms reshape the grid interconnection landscape, data center developers must adopt **informed, flexible strategies**. Whether pursuing direct interconnection or co-locating with generation, success will depend on navigating evolving FERC rules, interconnection queue mechanics, and the emerging regulatory treatment of hybrid generation-load facilities. Choosing the right path requires a careful balance between **cost, control, risk, and speed to market** — with increasing need for creative approaches with respect to power supply and utility partnerships.

Limitations on Supply Arrangements

Because electric utilities maintain regulated monopolies in most states, state law determines how easy or difficult it is for data centers to contract dedicated energy resources. Currently, **only 13 states have comprehensive retail electric choice programs**, allowing generators to sell or transfer electric energy directly to co-located loads, including data centers. Another eight states have limited programs for such arrangements.

If a state does not have a retail choice program and the data center owner is not itself the owner of the on-site generation, the generation owner must "sleeve" sales of electric energy through an intermediary. Typically, these sleeved transactions are accomplished through **back-to-back PPAs** where, under one agreement, the generation owner makes wholesale sales of electric energy to the intermediary (*e.g.*, the interconnecting utility) and, under the other agreement, the intermediary makes retail sales of electric energy to the data center load.

Exemptions under the Public Utility Holding Company Act

Sleeved transactions are common even in states with robust retail choice programs because they preserve exemptions from federal regulation arising under the Public Utility Holding Company Act (PUHCA). Originally enacted in 1935 and partially repealed in 2005, **PUHCA establishes onerous accounting, recordkeeping, and reporting requirements** for holding companies of entities that own in-service generation and transmission facilities. These requirements were designed to ensure that holding companies were not subsidizing their utility-related businesses with nonutility-related revenues and vice versa.

There are multiple options for pursuing exemptions and waivers from the PUHCA requirements, but the most common and favored exemption by developers and financing parties alike is to have **the generation-owning project company self-certify with FERC as an exempt wholesale generator (EWG)**. To do this, the project company must affirm that it is engaged exclusively in the business of owning or operating, or both owning and operating, facilities used to sell electric energy at wholesale. Although EWGs can engage in certain incidental activities, they cannot engage in any retail sales or transfers of electric energy unrelated to wholesale electricity sales. An exception to this rule would be the provision of station power to other EWGs that might be located behind the same point of interconnection.

Using a sleeved transaction accomplishes two goals: it facilitates **service to co-located loads** even in states that do not have retail choice programs, and it **preserves EWG status**.

In states where it is possible for a generator-owner to sell or transfer electric energy at retail behind the meter to co-located data centers, a generator-owner can seek alternative exemptions from PUHCA but still might be restricted by their financing arrangements. **Financing parties typically require that generator-owners secure EWG status to ensure that they will not become subject to the accounting, recordkeeping, and reporting requirements under PUHCA**, particularly because the financing party may need to foreclose on their interests on the underlying generating facility. Accordingly, even in states with retail choice, many independent power producers use sleeve transactions to maintain EWG status in accordance with their financing obligations.

Need for Market-Based Rate (MBR) Authority

Any entity that makes wholesale sales of electric energy within the contiguous United States outside of the ERCOT region of Texas must file the rates, terms, and conditions of such sales with FERC. To make **wholesale sales of energy at negotiated or market-based rates, including pursuant to a PPA, a seller must obtain MBR authority** from FERC. For example, independent power producers that enter into sleeved transactions or that wish to make wholesale sales of any excess generation produced by their projects must obtain MBR authority before making any such sales, including test sales from their projects.

Data center owners do not require MBR authority to purchase energy, but they could require it to the extent they intend to resell any energy purchased under their supply arrangement or if they wish to enter into power marketing activities in wholesale markets.

Obtaining MBR authority from FERC is a relatively simple process for most wholesale sellers. A seller must demonstrate to FERC that it and its affiliates lack, or have adequately mitigated, both horizontal (generation) and vertical (transmission) market power in the relevant wholesale market. In addition, a seller must file with FERC a standard MBR tariff that provides terms, conditions, and any applicable limitations relating to its sales of wholesale energy, capacity, and ancillary services. Once a seller obtains MBR authority, it must file with FERC quarterly reports with summary information about its wholesale sales. Additional filings would be required to reflect increases in generating capacity controlled by the seller and its affiliates. As "public utilities" under the Federal Power Act, entities with MBR authority also must obtain prior FERC authorization for financing arrangements and changes in direct or indirect ownership of the seller.

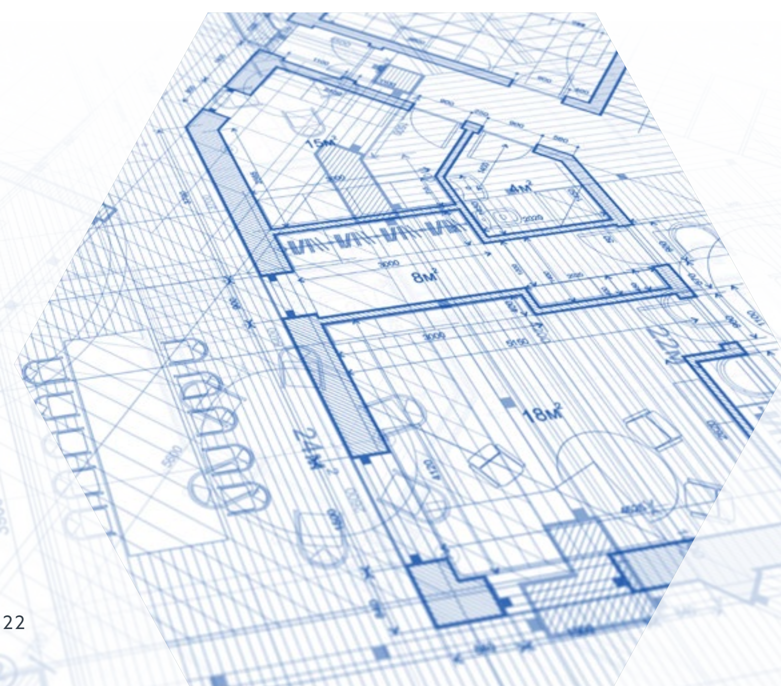

3 Leasing Considerations

Leasing a hyperscale data center, whether a fully operational site or through a build-to-suit arrangement, presents many considerations for prospective data center tenants and developers. For businesses looking to lease capacity, hyperscale data centers can offer strategic advantages by providing flexibility, scalability, and cost efficiency over owning a data center outright. Choosing the right facility will provide a foundation for future growth while minimizing risks.

For developers, the leasing process requires an assessment into a number of factors such as local regulations and political environment, environmental conditions and constraints, availability of incentives, network infrastructure, and proximity to potential users. In addition to navigating where to lease, choosing the right facility requires ensuring that the terms of the lease align with the tenant's needs and adequately address challenges such as scalability, workforce availability, and cybersecurity risks.

3.1 Pre-Lease Considerations

Prospective tenants must conduct thorough due diligence to ensure the proposed data center site and facility will meet their current requirements while also supporting future growth. The evaluation process should examine critical factors including location, provider capabilities, infrastructure scalability, and resource availability — all of which significantly impact long-term operational success and total cost of operation. Data center operators — referred to in this section as the "landlord" — should consider tenant priorities when choosing a future data center's location and infrastructure.

The location of a hyperscale data center significantly impacts **latency, regulatory compliance, and disaster recovery** capabilities. Further, putting aside power availability issues, any new data center development should consider the burden of environmental and zoning issues as well as the benefits of any tax or other economic incentives.

Zoning

Local zoning and land use regulations have emerged as critical determinants of data center project viability, as states and municipalities have taken differing positions on hyperscale data center development. Many jurisdictions have welcomed data center development (e.g., Central Ohio) by adapting land use regulations specific to data centers. Data centers, however, typically have low employment density, high energy and water usage, and intense demand for connectivity. These characteristics often place them in tension with zoning codes that were designed for traditional commercial and industrial uses.

Most data centers fall under **industrial or light industrial use** categories. Zoning codes typically impose floor area ratio (FAR) restrictions that limit the ratio of total floor area to the size of the parcel. Because data centers often prioritize equipment density and internal cooling over vertical expansion, FAR restrictions can artificially constrain capacity. Similarly, height restrictions may prevent the construction of large vertical mechanical systems such as cooling towers or air handling units. Site coverage requirements, setback mandates, and parking minimums (which are often unnecessary given low staffing levels) can also increase costs or reduce usable area.

To accommodate data center development, jurisdictions that aim to welcome data center development may **offer by-right use in utility-rich industrial zones**, increase FAR and height limits in designated corridors, and remove outdated parking requirements. Jurisdictions that are restricting further development — particularly those in saturated markets — are unwilling to adapt existing land use regulations. If a party desires to acquire and then expand or redevelop an existing data center, it should ensure that such expansion or redevelopment will not conflict with existing zoning code or receive significant public opposition.

Water and Cooling

Water availability has become an **increasingly critical constraint** in global data center development, rivaling power availability as a primary factor in site selection and long-term operational planning. Data centers typically use a large amount of water in their cooling systems.

Water consumption in data centers has reached unprecedented levels. Virginia's data centers alone consumed nearly two billion gallons in 2023, representing a 64% increase from 2019. Advances in technology can help manage water constraints and reduce the correlation between computing power and water use. Traditional systems — based on cooling towers and chillers — consume thousands of gallons per MW per day. In contrast, **closed-loop systems target heat at the source** and may require minimal water. However, advanced cooling technologies can increase capital costs, require technical expertise to manage performance risk, and sometimes involve higher power usage effectiveness.

Where data centers use large volumes of water, data center projects are under increasing scrutiny from water authorities and the public, especially in arid or semiarid regions like Arizona, Utah, and California. Conversely, there are jurisdictions, such as the Ashburn corridor in Virginia, that are investing in increased water utility capacity to accommodate further data center expansion.

Developers now face **more complex environmental impact assessments** and may try to acquire or lease water rights in advance to ensure adequate supply throughout a facility's lifespan. For developers, particularly in regions with reliable and inexpensive water supplies, this represents a strategic opportunity. Parcels with secure water access now command premium valuations, especially among hyperscalers seeking to de-risk expansion pipelines.

Other Considerations

Regulatory Compliance: In addition to local zoning ordinances, many states have regulatory schemes that may impact the development or operation of data centers including **consumer data privacy laws**. Landlords and tenants should consider these requirements in determining the cost and/or burden of operating a data center in a particular jurisdiction.

Incentives: To attract data center development, many states, counties, and municipalities have offered **tax or other incentives** in connection with the initial construction or operation of the data center. Whether these ultimately benefit the landlord or tenant, the parties should explore all potential incentives in selecting a site for a new data center development.

Proximity to Users: Considering how quickly we receive our Google search results, it's easy to forget that there is a physical aspect to the transfer of data. Parties should select locations for data centers that both **minimize latency and optimize performance** for end users.

Disaster Recovery: Physical disaster risk will vary by location. The risks inherent with a site in California are different from those with a site in Iowa. Parties should confirm that the proposed data center has **sufficient redundancy** and failover systems in place to account for potential disasters.

Infrastructure: Data center tenants should evaluate the quality and capacity of telecommunications infrastructure at the data center. Such infrastructure must ensure that the data center can deliver the computing power and network connectivity needed for the tenant's business.

Landlord: The landlord plays a crucial role in maintaining the data center and ensuring tenant satisfaction. Given the number of recent entrants into the hyperscale data center market, tenants should adequately vet the landlord entity as well as its personnel. It's important to confirm that the landlord has adequate experience in the operation of hyperscale data centers of similar size and geography. With respect to build-to-suit projects, tenants should prioritize landlords with significant experience in managing hyperscale data center construction and navigating supply chain issues.

Service Quality: If available, tenants should discuss other customer's past experiences with the landlord. Unlike in a typical triple net lease, the landlord will provide significant services to the tenant throughout the term of the lease. Each tenant will have unique requirements as it relates to its intended use of a data center site, and those should be communicated at the outset of the lease discussions. Tenants should ensure that the landlord is willing to accommodate and adapt to the tenant's specific needs and requirements.

Financial Stability: Tenants should review the corporate structure and upper tier ownership of the landlord to assess whether any equity investors are competitors of the tenant and evaluate the landlord's financial capacity to complete construction.

Growth Potential: The landlord and tenant should consider the potential growth of the data center project or campus. This may be structured through multiple leases, options, or rights of first refusal/offer. Such expansion may require significant upgrades to the data center campus, including building or expanding substations to meet any increased power demands. Modular infrastructure of the initial hyperscale development will help to accommodate future growth of the overall project.

3.2 Lease Terms

Hyperscale data centers are designed to support robust, scalable applications and services, often serving as the backbone for cloud providers and large enterprises. Leasing capacity in these facilities presents unique challenges and opportunities.

Security and Access: Security is paramount in hyperscale data center operations to protect sensitive data and infrastructure.

Physical Security: Considering the sensitivity of information housed within any data center, the parties should ensure that the premises have robust physical security measures in place, including gated perimeter access, surveillance, and biometric access controls.

Cybersecurity: In addition to physical security, the tenant should employ appropriate software security tools to prevent unauthorized access to any data contained in the data center facility.

Assignment: Assignment clauses in data center leasing agreements allow tenants and landlords to transfer their lease obligations under certain conditions and subject to specific requirements.

Flexibility: In data center leases, assignment restrictions are typically imposed on both the landlord and tenant. The landlord should seek to prohibit direct assignments by tenants but should be prepared to permit certain indirect assignments (i.e., changes of control) of tenants. Due to the scope of services being provided by a data center landlord, data center leases often restrict the landlord's ability to sell or otherwise transfer the data center (directly or indirectly) without the tenant's consent.

Approval Process: The lease should clearly outline the procedures for either party's proposed assignment. Such procedures should include notice requirements, qualified transferees, prohibited assignees (e.g., competitors), financial requirements of a proposed assignee, etc.

Acceptance Testing: Acceptance testing ensures that the leased space meets specified requirements before occupancy.

Criteria: The data center lease should clearly define the criteria for acceptance testing to ensure all data center systems meet specified operational standards. The lease should describe the various commissioning stages and the requirements at each stage.

Timeline: The lease should specify a timeline for acceptance testing and a detailed procedure for the resolution of any issues identified.

Documentation: Throughout the term of the lease, the parties should maintain thorough documentation of any commissioning, acceptance testing procedures, and the results thereof.

Service Level Agreements (SLAs): SLAs outline the expected performance and reliability standards for data center services to be provided by (or on behalf of) the Landlord.

Metrics: The lease should clearly detail the various metrics for each of the data center components, including requirements for response times, temperature and cooling requirements, and acceptable data transmission delays.

Penalties: The parties should establish clear, objective penalties for noncompliance with any SLA terms. These are often in the form of rent or service credits for the tenant's benefit; however, repeated SLA violations may eventually result in an option to terminate all or a portion of the data center lease.

Review: The parties should regularly review and update SLAs to address any recurring issues or to reflect any changing requirements of the tenant.

Building Management Systems (BMS): Building management systems are essential for monitoring and controlling data center operations. They support proactive management of the data center while providing tenants with transparency regarding service delivery.

Monitoring: Prior to entering into any data center lease, the tenant should review and approve the landlord's proposed form of management systems reports to confirm they contain the information needed by the tenant. These reports should provide sufficient detail regarding the management of the data center's electrical, cooling, mechanical, and electromechanical systems.

Upgrades: The monitoring plan should also provide for regular upgrades to the data center's systems. The lease should clearly outline the scope of any required upgrades and specify which party is responsible for the cost of such upgrades.

Planned Maintenance: Planned maintenance is crucial for ensuring the reliability and longevity of data center infrastructure.

Schedule: The lease should include a detailed regular maintenance schedule to minimize downtime or other disruptions at the data center.

Documentation: Whether or not provided to tenant, the landlord should maintain detailed records of maintenance activities throughout the term of the lease.

Audits: Hyperscale data center audits are comprehensive evaluations of a data center's infrastructure, operations, and security practices, designed to ensure compliance with the standards specified in the lease.

Frequency: The parties should determine the frequency of any internal or external audits based on the size and type of data center facility together with the particular needs of the tenant.

Scope: The lease should clearly define the scope of any required audits (such as evaluating physical equipment (including any power or cooling equipment), disaster recovery plans, and network connectivity).

Reporting: The parties should ensure that the results of any audit are documented and delivered to the landlord, tenant, and any other relevant parties (e.g., lenders).

End-of-Lease Requirements: End-of-lease requirements outline the obligations of both parties upon the expiration or termination of the data center lease.

Notice Period: Given the long lead time that may be necessary to locate a new tenant for a data center facility, landlords should require significant advance notice from tenants at the expiration of the lease term (e.g., vacate premises, extension, purchase, in each case to the extent applicable).

Purchase Rights: In a single tenant data center facility, the tenant should consider negotiating for an option to purchase the data center facility at the expiration of the lease term. Alternatively, the tenant should consider seeking term extension options or rights of first refusal or rights of first offer in connection with a sale of the data center to a third party.

Restoration: The lease should specify in detail the requirements of the tenant at the expiration or termination of the lease, including requirements for removal of any equipment within the data center and restoring the premises to its original or a specified condition.

Final Inspection: To ensure compliance with any requirements at lease expiration or termination, the landlord and tenant should coordinate a joint inspection of the data center and all equipment thereon prior to the expiration of the term or promptly following lease termination.

4 Financing

In 2024, there was an unprecedented level of capital expenditure in data center infrastructure. Global capital expenditure on data centers **surged 51%, reaching a total of $455 billion**, according to Dell'Oro Group. This increase was primarily driven by extensive investments made by hyperscale cloud providers in infrastructure optimized for AI workloads. Worldwide spending on data center hardware and software achieved an all-time high of $282 billion in 2024, which represented a 34% growth compared to the previous year, highlighting the rapid expansion of the digital infrastructure landscape.

The massive amount of capital expenditure required to develop, power, equip, and operate modern data centers has created a voracious demand for financing. Bank loans continue to be a vital source of capital for developers and operators. However, financing sources are diversifying beyond traditional bank lending, reflecting the sector's maturity and appeal as an asset class. Key alternative financing strategies include private debt, securitizations, and specialty finance. Tax-exempt public financing of data center-adjacent infrastructure may have its own role to play as the capital demands of this asset class continue to boom and seek out increasingly innovative financing options.

The choice of financing strategy for data center assets depends on a number of factors, such as:

- The balance sheet strength and creditworthiness of the data center developer and operator.
- Tolerance of developers and sponsors for full-recourse versus limited or nonrecourse financing.
- The sponsor's financial objectives, including managing leverage toward the desired rate of return.
- Development stage of the data center asset, including leased-up status.
- Capital needs of the project.

For large-scale, greenfield construction projects, particularly hyperscale facilities, where the loan is secured by the project's anticipated future cash flows rather than the developer's overall balance sheet, it allows for risk isolation but involves complexity, cost, and time — and requires a granular understanding of the project risks being absorbed by lenders and how to mitigate those risks.

The remainder of this chapter delves into how a variety of financing sources cater to the needs of data center developers, operators, and sponsors.

4.1 Bank Loans

Bank loans remain a foundational source of financing for data center businesses and projects. Large financial institutions have the balance sheets to provide financing in the quantum required by today's data centers, either on a bilateral basis or — in the case of large financings in the hundreds of millions or several billions of dollars — with a multibank lending syndicate. These large financial institutions can also utilize different underwriting approaches to structure, arrange, and execute the type of financing with the characteristics that most closely align with the objectives of data center developers, operators and sponsors. Broadly speaking, bank loans to data center businesses come in at least three varieties:

- Full-recourse corporate loans, which banks underwrite based on the creditworthiness of an operating business.
- Limited or nonrecourse loans, which banks underwrite on the basis of a specific asset's particular attributes.
- Portfolio loans, which banks underwrite on the basis of a pool of data center assets that, collectively, cross-collateralize the financing.

Corporate Loans

As with other cash-flow positive businesses, data center companies that own and operate a mature portfolio of data centers borrow from banks for a variety of purposes — funding working capital requirements, financing strategic acquisitions, building land banks for future development, and bridging early-stage projects until they become financeable on a standalone basis. These **corporate, parent-level loans** to data center companies look similar to what other operating companies may have on their balance sheets — variable interest rate obligations with a margin that can fluctuate based on how levered the business is relative to its earnings, meaningful principal amortization, and one or more financial maintenance covenants. These loans can either be secured by substantially all of the business assets (for non-investment grade companies) or unsecured (for investment grade companies). The documentation for bank loans includes various affirmative and negative covenants that are binding on the parent borrower and, with some exceptions, its subsidiaries, including restrictions on additional debt incurrence, acquisitions and other investments, dividends and distributions, and asset sales. In addition, a number of large data center operators are organized as real estate investment trusts, or REITs, with credit facilities tailored to reflect their unique tax and other operating requirements.

ADVANTAGES

- Lower cost of capital across the full credit spectrum of borrowers.
- Ability to arrange large financings by syndicating and spreading credit risk across multiple banks.
- Potential to capitalize on relationship banking for better terms and access to a full suite of credit products.
- Often no cost to refinance or replace bank financings.

CHALLENGES

- Repayment risk not isolated to data center assets but rather attaches to the entire enterprise.
- Increased leverage can create downward pressure on corporate credit ratings and other credit metrics.
- Financial maintenance covenants exacerbate difficulties arising from underperformance of business plan.
- Enterprise-wide negative covenants can hinder operational flexibility.

Limited / Nonrecourse Financing

There are two different models for limited or nonrecourse financing of data centers — the project finance model and the real estate finance model. **Project finance** involves limited or nonrecourse funding, where lenders and investors finance a specific data center project based on its expected future cash flows, rather than relying solely on the parent company's overall financial health. By comparison, **real estate finance** is a form of limited or nonrecourse financing where lenders underwrite based on the value (typically on an "as-completed basis") of the data center building and underlying land. Both project finance and real estate finance are often used for construction debt, especially for large-scale greenfield hyperscale facilities still in development, and for bridge loans for projects still in the process of being leased up. Lenders in a project financing typically require **contracted stable cash flows** pursuant to long-term agreements with tenants, while real estate lenders will size loans based on a **maximum loan-to-cost and loan-to-value (LTV) formula**. In both cases, lenders may require limited recourse to project owners or sponsors in the form of nonrecourse carveout guarantees (so-called "bad-boy" guarantees), completion guarantees, interest guarantees, and carry cost guarantees.

ADVANTAGES

- Protects the parent company's assets by isolating financial risks within the specific data center project, subject to delivery of required guarantees from the parent company (or another deep pocket, such as a financial sponsor) which will vary on a deal-by-deal basis.
- Allows companies to fund capital-intensive projects that might otherwise be too costly to undertake based solely on their balance sheet.
- Can bring together multiple capital sources (lenders, investors, equity partners) structured specifically for the project's needs and maximize the lowest-cost available mix of capital sources.

CHALLENGES

- Generally involves more complex structuring compared to traditional corporate loans, with intensive due diligence of the project (including the real estate, leases, entitlements, power supply, construction costs, and other operating expenses) and mitigation of any material risks identified. As such, it may result in higher transaction costs and longer timelines than corporate financing.
- Requires comprehensive and accurate financial modeling to predict cash flows based on in-place leases and projected lease-up and to ensure the project can meet its debt obligations.

- In the case of real estate financings, lenders require third-party appraisals as a condition to providing the financing and on a periodic basis during the term of the loan.
- Borrowers are typically subject to greater oversight, tighter covenants, and stricter borrowing conditions from lenders, leading to higher compliance costs. A deep pocket guarantor will be required to deliver guarantees as well and may be subject to reporting and net worth and liquidity requirements.

Project finance is seeing increased activity, particularly for funding the construction of new large-scale hyperscale data centers.

- **Vantage Data Centers** secured $5 billion in incremental green loan financings to support demand for its North America platform, including a $2.25 billion construction loan for its Ohio campus and a $2.75 billion corporate facility upsize. Led by Mitsubishi UFJ Financial Group (MUFG) and Societe Generale, the transaction is one of the first construction loans in the data center industry to have achieved private investment-grade ratings.
- **Switch Bighorn** secured $4.5 billion in debt financing for a 300 MW capacity turnkey project and a 360 MW capacity project in Reno, Nevada, with both projects anchored by an investment-grade global cloud services provider. The cash flows for these projects derive from fixed lease payments under a triple net lease, with the tenant responsible for operating expenses.
- **Stack Infrastructure** secured multiple green financing deals for construction and expansion of data centers in Virginia, including $350 million for a 150 MW campus in Manassas and $301 million for a 200 MW campus in Loudoun County.
- **DataBank** also secured construction debt financing last year from Siemens Financial Services for another hyperscale data center in Virginia.

Portfolio Loans

Portfolio loans — where a group of data center assets collectively secure and support a single financing structure — have some elements of corporate loans and limited / nonrecourse financings but are distinguishable from them in key ways. On the one hand, unlike single-asset project or real estate financings, lenders are not underwriting based on the attributes of a single data center. On the other hand, portfolio loans are not backed by the financial strength of a diversified operating company or its sponsors, which means lenders must still understand and underwrite to specific asset-level attributes. Portfolio loans are typically secured by a **pledge of the equity in the asset companies** that own the underlying data center assets and, in many cases, by asset-level security (mortgages on the land and improvements and security over the personalty, fixtures and material contract rights).

Portfolio loans can be used for a variety of purposes, including financing under a single debt facility multiple data centers that may be at different development stages. Portfolio loans can also be a way of monetizing assets — for example, by funding a dividend to the portfolio's owners — when bid/ask spreads for an M&A-style exit may be too wide to transact upon and when capital markets are not accessible to the portfolio due to its asset composition or challenging market conditions.

ADVANTAGES

- Diversified asset base can lead to lower cost of capital than single-asset financings.
- Flexible use of proceeds, including development costs and capital expenditures, operating expenses and return of or recycling capital invested by portfolio owners.
- Often include an "accordion" feature that allows portfolio owner to obtain incremental debt financing by bringing additional assets into the portfolio over time.

CHALLENGES

- Cross-collateralization of financing puts multiple assets at risk if one or more data centers in the portfolio experience stress.
- Lenders may still require extensive due diligence as a condition to providing the financing and ongoing asset-level reporting over the term of the financing.
- Complex structuring and documentation, particularly for portfolios where data centers are located in multiple jurisdictions or in transactions where lenders require asset-level security.

4.2 The Role of Private Debt

Private debt has played an increasingly prominent role in financing operating companies, financial sponsors, and real estate and infrastructure assets across a variety of industries and market segments. Moody's estimates that total private credit assets under management **totaled nearly $2 trillion at the end of 2024, with that total projected to rise to $3 trillion by the end of this decade**. The popularity of private debt as an asset class and the influx of capital from a growing universe of investors mean that the amount of "dry powder" managers need to draw on and deploy will continue to grow at a rapid pace. As a result, private debt is staking out new frontiers in the quest to identify quality investment opportunities.

Private debt has played an increasingly prominent role in financing operating companies, financial sponsors, and real estate and infrastructure assets across a variety of industries and market segments.

The data center industry is one area where private debt has sought out investment opportunities and by all indications will continue to do so for the foreseeable future. Data centers also appeal to several different strategies within the private debt universe. Funds whose primary investment strategies include real estate, infrastructure, senior secured loans, mezzanine debt, or structured equity all have found opportunities to deploy capital in data center assets and businesses.

ADVANTAGES

- Versatility of private debt enables it to be utilized in many of the same bank loan structures described in the previous section.
- Private debt has a higher risk tolerance than bank lending, which allows developers, operators and sponsors to push leverage levels beyond what banks may be willing to underwrite. Utilization of PIK (payment-in-kind) interest features allow sponsors to reduce or defer capital contributions to the project, potentially increasing internal rate of return.
- Because private debt markets are not regulated to the same degree as banks, private debt can in some cases be nimbler and more flexible in certain situations (for example, in real estate-based lending where banks must comply with various requirements under the federal Financial Institutions Reform, Recovery, and Enforcement Act, or FIRREA, that do not apply to private debt).

CHALLENGES

- The flexibility of private debt does come with a cost, as the all-in-yield for private debt, all else being equal, is typically higher than it is for bank loans. That premium, however, has trended lower recently, as private debt increasingly competes directly with bank lenders for deal flow.
- Because private debt portfolios must manage returns and maintain capital invested over longer time horizons, refinancing private debt can come with a cost — such as prepayment premiums, minimum investment returns or exit fees — that is not as prevalent in the bank lending market.
- Private capital may not provide a "one-stop shop" for all needs of data center developers and operators, compared to the full menu of products that bank lenders offer (cash management, hedging, working capital facilities and letters of credit / bank guarantees).

4.3 Securitizations

Securitizations offer data center operators an additional source of capital markets financing. They provide operators with access to a broad range of debt investors with varied appetites for risk and often result in a lower cost of financing when compared to other financing options.

So far, data center securitizations have been in the form of either:

- **Asset-Backed Security (ABS)** offerings, in which secured **notes** sold to investors **are repaid from cash flows generated by the data centers** themselves, usually from unencumbered customer leases and/or fees payable to the data center owner and operator under service-related contracts; or
- **Commercial Mortgage-Backed Security (CMBS)** offerings, in which pass-through certificates are repaid from the debt-service payments on, and/or the refinancing of, a **mortgage loan secured by the data center's real property**, with the scheduled mortgage loan payments, in turn, being supported by the rent payable by tenants under the data center's customer contracts or leases.

Each of these structures presents distinct advantages and disadvantages.

ABS Structures

Because data centers are structures that sit on real property often owned by the data center operator, **they are easy to finance through mortgage loans paired with a related CMBS offering** that relies primarily on the credit quality of the center's tenants and the terms of their leases. However, if a data center is not encumbered by a mortgage and/or has large revenue streams generated by customer leases or service contracts, the **receivables from those leases and contracts allow the operator to undertake traditional, receivables-backed ABS offerings**.

Data center ABS offerings are structured and rated according to a **hybrid of CMBS, project finance, and ABS rating methodologies**. While the rating and thus the viability of a data center CMBS offering will focus on the ratio of the underlying securitized loan to the value of data center property itself, the rating and viability of a receivables-backed, data center ABS offering will focus on a comparison of (i) the revenues expected to be received by the operator under its leases and service contracts over the term of the ABS and (ii) the cost of operating and maintaining the data center. Other factors that are considered in rating a data center ABS transaction include:

1. The predictability of revenues.
2. The creditworthiness of the facilities' customers.
3. Whether there are multiple customers (as would be typical for co-location facilities).
4. The diversification of cash flows across customers and customer industries.
5. The terms and provisions, including termination provisions, of customer leases and service contracts.
6. The age and physical condition of the facilities and their technology (including the potential for the technology's obsolescence).
7. An operational assessment of the data center operators and servicers and their respective management teams.
8. The prospective demand for the data center in the geographic location in which it is situated.
9. Available credit enhancements, usually in the form of debt-tranching, overcollateralization, or debt-service reserves.
10. To some extent, whether the LTV ratio of the anticipated debt-to-site values of the ABS are secured by real property.

Master trust structures, traditionally used in credit card receivables and auto loan ABS offerings, can provide the flexibility and efficiency required to navigate the dynamic landscape of data center ABS offerings. These structures allow ABS issuers to add new data center assets to the trust as needed, helping to maintain or enhance the overall value of the asset pool. By adjusting the collateral composition over the course of the transaction, data center ABS issuers can better manage risks associated with fluctuations in data center performance or economic changes, ensuring the trust remains appealing to investors. In addition, master trusts enable ABS issuers to issue additional debt over time, allowing them to capitalize on favorable market conditions, respond to shifts in investor demands or interest rates, and manage their financing strategies without needing to establish new securitization structures for each issuance.

ABS structures also lend themselves to financing in the short-term asset-backed commercial paper (ABCP) market, which, historically, has been a reliable funding source for a wide array of trade receivables.

CMBS Structures

The securitization of data centers in CMBS offerings is a specialized but growing area. Some transactions feature data centers as the sole or primary collateral, while others include them as part of diversified pools. The two primary types of CMBS transactions are **single asset single borrower (SASB) structures**, which generally involve the securitization of a large-balance mortgage loan by one or more originators or co-originators and **pooled mortgage loan structures, often referred to as a "Conduit,"** which generally involve the securitization of a diverse pool of mortgage loans of varying sizes and characteristics by multiple originators.

More recently, SASB securitization structures have offered borrower sponsors, such as data center owners or operators, with **an agented process** to raise capital from various investors by seeking origination, from a commercial lender, of a single mortgage loan to be securitized. This structure provides borrowers and investors with flexible financing options and allows securities to be pre-marketed, ensuring investor preferences are considered during the structuring process. In a SASB structure, a mortgage loan can have a fixed or floating interest rate, and the certificates issued can match this rate type. These certificates can also have different features to attract investors, while complying with real estate mortgage investment conduit (REMIC) tax rules and other legal and regulatory limits. In recent SASB CMBS transactions, the issuance of the securities is generally contingent upon the closing of the loan, with pre-marketing of the securities to be issued conducted as part of the pre-closing process.

By contrast, a **Conduit structure is characterized by multiple mortgage loan sellers** transferring a pool of mortgage loans (made by the various loan sellers or their affiliated originators to various borrowers and which are backed by various property types) to a common law trust which issues the CMBS. Conduit transactions are generally structured with the more senior classes of certificates being registered under the Securities Act of 1933 and the various mezzanine and/or junior classes of certificates being sold as part of a private placement. In Conduit transactions, to produce advantageous rating criteria and robust investor interest, and to reduce concentration risk, the pools are structured to produce diversity in borrower, sponsor and tenant concentration, loan size, and geographic location, among other characteristics.

In the event that all certificates backed by a mortgage loan (which may be large) included in an SASB CMBS transaction are not placed, or to reduce such mortgage loan's concentration in any one SASB or Conduit transaction, the mortgage loan may be severed into multiple notes which can be included in one or more subsequent Conduit CMBS transactions. Accordingly, it may be included in both a SASB and one or more Conduit transactions, or in multiple Conduit transactions, all serviced under one of such securitizations (called the lead securitization), and with one such note being the **"controlling note."**

For both SASB CMBS transactions and Conduit CMBS transactions, the investor purchasing the most subordinate class or classes of certificates can act as or appoint a **controlling class representative** that has certain rights, including the right to advise and direct the applicable special servicer with respect to certain actions regarding specially serviced mortgage loans.

Comparing CMBS and ABS Securitizations

Securitization of data centers as part of CMBS and ABS securitization transactions has certain advantages and challenges, which may not be totally aligned. However, in general, many of the advantages and challenges of general ABS structures are shared by CMBS transactions, as listed below.

SHARED ADVANTAGES OF CMBS AND ABS SECURITIZATION TRANSACTIONS

- Provide access to potentially less costly capital and a broader investor base.
- Allow operators to monetize future income streams.
- Can extend leverage and lower borrowing spreads compared to other methods of financing.
- Appeal to investors due to their generally stable cash flows, long-term leases, and the essential nature of their services.

SHARED CHALLENGES OF CMBS AND ABS SECURITIZATION TRANSACTIONS

- Sensitivity to interest-rate fluctuations.
- As a relatively new asset class (with the first deals closing in 2018), rating methodologies are still evolving.
- Require stable and predictable revenue streams, from long-term leases or other receivables-generating contacts, whereas some data center leases are short-term.
- Can involve complex structuring, documentation and multidisciplinary expertise.
- Tenant contracts do not permit disclosure of the name of the customer/tenant and address, creating challenges during the marketing process.

Private Placement Securitization Transactions

Private placement securitization transactions involve selling shares or bonds to preselected investors and institutions rather than on a public exchange. While the majority of data center securitization activity to date has been in the form of Rule 144A offerings, private placements offer several potential benefits to borrowers utilizing securitization for their funding needs.

These include the potential for **lower transaction costs** compared to broadly syndicated institutional offerings and a more streamlined execution process. Private placements typically involve **fewer intermediaries** and more direct communication between the lender(s) and the borrower, facilitating faster deal execution and reducing transaction expenses. These arrangements can also foster deeper, more collaborative relationships and better alignment of interests between the lender(s) and borrower, especially during the diligence process, when lenders tend to gain deeper insight into the borrower's business. Private placements of data center securitizations have the additional benefit of **limiting dissemination to the broader market** of the rates, duration, and applicable covenants and conditions on the debt and descriptions of the underlying collateral that are included in Rule 144A offering materials.

The **terms of the debt issuance** in a private placement transaction are typically set forth in an indenture executed between an indenture trustee and the special purpose entity that is established by the borrower for the purpose of holding the collateral and offering the debt that is secured by the collateral. In both Rule 144A and private placement data center securitization transactions, the repayment of the notes is secured by the cash flow generated by the data center assets or real estate interests contributed to the securitization, while being de-linked from the corporate credit quality of the data center operator.

In addition, private placement data center securitizations are sold **directly to lenders through a note purchase agreement**. This means that lenders have a direct contractual relationship with the borrower and can negotiate specific terms or protections tailored to the transaction. This makes private placement ABS transactions well-suited to circumstances where a prospective data center borrower/issuer has unusual attributes or history, or where the underlying data center properties, leases or service contracts have attributes that may be more effectively addressed with prospective lenders via direct, bilaterial discussions and negotiations typical for a private placement transaction rather than via traditional Rule 144A debt offering disclosures.

While the number of participants in the private placement esoteric securitization market continues to increase, and those participants generally have the ability to take large allocations of debt, the Rule 144A market has become relatively well-established for offering data center borrower/issuers very competitive pricing and a strong investor appetite for the offered debt. As a result, Rule 144A offerings for data center securitization notes are likely to continue to be the predominant path for data center securitization issuers to access the market, but private placements offer an alternative means of accessing credit that may be a good fit for certain data center operators.

4.4 Public Finance Considerations

Public finance, or municipal finance, generally involves the **issuance of debt by state or local governmental agencies** to finance costs of public infrastructure or certain other eligible costs. In many jurisdictions, the developer of property is responsible for the financing and construction of public infrastructure improvements, such as roads or public utilities, that are necessary for the development of the data center. However, it may be possible to coordinate with local governmental agencies to provide for the issuance of **lower-cost tax-exempt municipal bonds**, payable from assessments or taxes imposed on the data center property, where the proceeds of the bonds are used to finance or reimburse costs of such public infrastructure.

ADVANTAGES

- Lower cost of capital for required public infrastructure due to local governmental agency's ability to access tax-exempt financing.
- Potential for turnkey financing, where infrastructure improvements are constructed by the developer and acquired by local agencies with proceeds of tax-exempt bonds upon completion.
- Debt is not on the developer's balance sheet but secured by assessments or taxes imposed on the property (or other arrangements for payments in lieu of taxes), and generally payable on a level debt service basis over 20 years or more.
- Supporting local agencies in the issuance of bonds may help develop relationships and support the acceleration of construction of any other public improvements necessary for or supportive of development.

CHALLENGES

- Local governmental agencies may have competing interests regarding development.
- Governmental entities may impose additional procedural requirements and/or have a different pace for completing transactions.

5 Mergers and Acquisitions

The data center industry has experienced significant Mergers & Acquisitions (M&A) activity in recent years, driven by the **increasing demand for cloud services and the explosion of AI learning intelligence and large language models**.

Key Drivers of M&A Activity

- **Demand for Scalability:** Data center acquisitions help companies scale their operations quickly and efficiently. This is particularly important for hyperscalers and other cloud service providers who need to expand their infrastructure to meet growing customer demands.
- **Geographic Expansion:** Acquisitions allow companies to enter new markets and regions, providing better service to local customers and reducing latency.
- **Technological Advancements:** The need to stay competitive with the latest technology drives companies to acquire data centers that offer advanced capabilities, such as AI and machine learning support.
- **Cost Reduction:** M&A can help companies achieve economies of scale, reducing operational costs and increasing profitability.
- **Strategic Partnerships and Consolidation:** The data center industry continues to see consolidation as companies merge to strengthen their market position and create synergies.
- **Efficiency:** Greenfield projects often require significant lead time to become operational due to several obstacles, including obtaining access to sufficient real estate for the data center (many of which encounter resistance from local communities), acquiring grid interconnection, navigating permitting requirements, and completing construction, among others. M&A offers a fast track for buyers in need of current or near-term operational capabilities.

Important Legal Considerations

- **Resource Requirements:** Data centers require significant and reliable power, cooling, and network connectivity. Ensuring that the data center has contracts in place to address such requirements is an important aspect of the diligence process.
- **Regulatory Compliance:** Data centers must comply with a variety of increasingly onerous regulations, including data protection laws, cybersecurity requirements, and environmental regulations. Ensuring compliance is critical during the due diligence process.
- **Data Privacy and Security:** Given the sensitive nature of the data stored in data centers, ensuring robust data privacy and security measures is a priority, and it is important to consider the legal agreements in place addressing liability and responsibility for system downtime and data breaches.
- **Intellectual Property:** Transactions often involve the transfer of proprietary technology and software. It is important to clearly define the ownership and rights to use intellectual property post-acquisition.
- **Real Estate and Zoning Laws:** Data centers require specific real estate considerations, including land rights, permitting, and zoning laws. [See Chapter 3: Leasing Considerations.] Legal due diligence must assess these factors to ensure no future operational issues.
- **Contractual Obligations:** Existing contracts with customers, suppliers, and partners need to be reviewed to understand any obligations or liabilities that may transfer with the acquisition.
- **Employee Transition:** M&A activity often involves the transfer of employees, which requires compliance with labor laws and consideration of employment contracts, benefits, and potential redundancies.

- **Antitrust and Foreign Direct Investment (FDI):** Data center acquisitions may attract scrutiny from both antitrust authorities and foreign direct investment regulators, so it is important to conduct transaction-related regulatory approvals early in the transaction process.

Transaction Structuring in Data Center M&A Deals

- **Asset vs. Stock Purchase:** Data center transactions can be structured as asset purchases or stock purchases. Asset purchases allow buyers to select specific assets and liabilities, **potentially avoiding unwanted obligations**. Stock purchases involve acquiring the entire company, including all assets and liabilities, which can be simpler. Stock purchases generally enjoy better tax treatment — by avoiding double taxation — **but may carry more risk** because in a stock sale, all liabilities, even unknown liabilities of the target, typically remain with the acquired entity (and may be addressed via indemnities or, increasingly, W&I insurance), whereas an acquiror in an asset deal may exclude all liabilities other than those specifically agreed to be assumed in the transaction.
- **Joint Ventures and Partnerships:** In some cases, companies may opt for joint ventures or partnerships instead of full acquisitions, especially when the parties to the M&A transaction have different capabilities (such as partnerships between data center developers and independent power producers). This structure **allows for shared ownership and risk**, providing flexibility in investment and operational control.
- **Leaseback Arrangements:** Given the significant real estate component of data centers, another option is a sale-leaseback arrangement. This allows the seller to monetize the asset while continuing to operate the data center under a lease agreement.
- **Financing Considerations:** M&A deals often involve complex financing arrangements, including debt financing, equity financing, or a combination of both. Structuring the transaction to optimize tax benefits and minimize financial risk is crucial.
- **Earn-outs and Contingent Payments:** To bridge valuation gaps, earn-outs, or contingent payments may be used. These structures tie a portion of the purchase price to future performance metrics, aligning incentives between buyers and sellers.

6 Tax Considerations

Data centers, while not currently benefiting from unique federal income tax incentives, can leverage various **state and local incentives designed to attract investment**. These incentives often include property tax abatements and favorable sales tax regimes. Additionally, data centers may incorporate renewable energy sources, such as wind and solar, and energy storage technologies to manage their energy needs effectively.

Under current law, renewable energy facilities like wind and solar may qualify for either a Production Tax Credit (PTC) or an Investment Tax Credit (ITC) under the Internal Revenue Code. Energy storage technologies may qualify for an ITC, though not a PTC.

The **PTC** is calculated based on the power produced by the facility and is governed by Sections 45 and 45Y of the Internal Revenue Code. Section 45 applies to projects starting construction before the end of 2024, while Section 45Y introduces a "tech-neutral" PTC for zero-emissions technologies. The credit is determined annually based on the energy produced and sold to an unrelated party over a 10-year period, with a statutory base rate of 0.3 cents per kilowatt hour (1992 dollars), adjusted annually for inflation. As of 2024, the inflation-adjusted base rate is 0.6 cents per kilowatt hour (KWh). Facilities meeting prevailing wage and apprenticeship (PWA) requirements benefit from a five-times multiplier, bringing the current PTC rate to 3.0 cents per KWh.

The **ITC** is based on the cost of ITC-eligible energy property within a facility or energy storage technology, as outlined in Sections 48 and 48E of the Internal Revenue Code. Section 48 applies to projects starting construction before the end of 2024, while Section 48E focuses on zero-emissions technologies. The ITC is calculated as a percentage of the cost of eligible property, with a base rate of 6% and an increased rate of 30% if PWA requirements are met. The ITC is subject to a five-year recapture period if the facility is destroyed, abandoned, or no longer qualifies.

Both the ITC and PTC offer additional credit bonus adders for facilities located in designated "energy communities" or meeting domestic content requirements (which require the purchase of components from domestic manufacturers). Each adder increases the PTC by 10% and the ITC by 10 percentage points. For example, a facility claiming the ITC that meets PWA and qualifies for both adders would receive an ITC equal to 50% of the cost of its eligible property.

On July 4, 2025, President Trump signed into law House Bill 1, more commonly known as the "One, Big, Beautiful Bill" ("OBBB"). OBBB significantly revised Sections 48E and 45Y of the Internal Revenue Code, most notably with respect to wind and solar technologies. While these and other technologies remain eligible for both PTC and ITC, new hurdles present additional qualification requirements. For example, for wind and solar technologies, to qualify for PTC or ITC, such projects must either start construction prior to July 4, 2026 or be placed in service by December 31, 2027.

7 International Trade and Compliance

The U.S. government's national security-related legal requirements as well as policy concerns associated with international trade and investment have continued to intensify, particularly as they relate to China. These requirements continue to have important effects in the technology and infrastructure sectors, and a number of compliance concerns should be considered by stakeholders in the data center industry. Data center developments must confront these **key compliance areas**:

- Import tariffs
- Export controls
- U.S. government focus on foreign investments in the U.S. energy, infrastructure, and technology sectors

Various National Security Import Tariffs

The Trump administration has utilized various legal authorities to impose new and, in several cases, unprecedented tariffs that may significantly impact the costs of developing U.S. data centers.

Imports of steel, aluminum, steel derivative, and aluminum derivative products are subject to 50% tariffs imposed by the Trump administration under Section 232 of the Trade Expansion Act of 1962. These tariffs can raise costs associated with, among other things, **steel and aluminum products** used in the construction of data centers.

Further, the Trump administration is pursuing **a national security investigation** under the authority of Section 232 of the Trade Expansion Act of 1962 that will likely lead to tariffs on imports of semiconductors and potentially certain electronics items containing semiconductors from many, if not most, countries worldwide. If imposed, these tariffs could substantially raise the costs of procuring foreign-made semiconductors and other electronics needed to power data centers in the U.S. Subject to ongoing court challenges, the Trump administration has imposed wide-ranging so-called "reciprocal" tariffs on most countries, ranging from 10% to over 100%, and has also imposed ostensibly fentanyl-related tariffs on Canada, Mexico, and China. The fate and duration of these tariffs is unclear.

Export Control Considerations

The United States, Canada, the United Kingdom, the European Union and its member states, and other countries administer export controls. Export control requirements prohibit or restrict **exports, reexports, and in-country transfers** of certain goods, services, software, and technology. An authorization of one or more government agencies may be required before transferring items across international borders or, in many circumstances, within foreign countries. Applicable export controls **depend on the assigned sensitivity** of the items being transferred.

A key concern in data center development relates to export controls specific to **chips essential to powering data centers**. For example, procurement of U.S.-origin chips and certain foreign-made chips otherwise subject to U.S. jurisdiction for international data centers often would require an evaluation of the export control status of the relevant chip (based on the chip's production and capabilities), and of the intended destination (i.e., the location of the data center). In some cases, the transaction may require a U.S. government and/or other government license that would be difficult to obtain.

CFIUS Considerations

For national security reasons, the U.S. President has the authority to block foreign investment in the U.S. that is subject to screening by the Committee on Foreign Investment in the United States (CFIUS). CFIUS is authorized to screen not only transactions that could result in a foreign person having, directly or indirectly, control over a U.S. business, but also certain noncontrolling investments and real estate transactions. Parties are required to notify CFIUS of certain types of transactions within CFIUS's jurisdiction.

In some circumstances, a CFIUS filing may not be legally required but, especially from the perspective of a foreign investor, may be prudent. These include investments in projects that involve a security sensitive aspect, such as proximity to a sensitive U.S. government site. Data center projects may be located near such sites. Further, CFIUS is likely to deem many data center projects to otherwise have national security sensitivities, so foreign parties engaged in development or operation of data center projects should seriously consider whether notifying CFIUS of such engagement is required or prudent.

President Biden issued Executive Order 14083 on September 15, 2022 (the "CFIUS Executive Order"), which provides formal direction on the risks that CFIUS should consider when reviewing transactions within its jurisdiction. The CFIUS Executive Order specifically directs CFIUS to focus on, among other things, **a transaction's effect on U.S. technological leadership and supply chain resilience and security** in areas affecting U.S. national security, including artificial intelligence.

A DIVESTMENT OF CRYPTOCURRENCY MINING REAL ESTATE

Demonstrating its focus on geographic proximity, in May 2024, CFIUS required the divestment of real estate located near a U.S. Air Force base acquired by **MineOne Partners Limited** which is ultimately owned by Chinese nationals, for national security reasons. MineOne purchased the land in 2022 for the purpose of establishing a cryptocurrency mining operation. In addition to proximity to the Airforce base, CFIUS assessed the risk associated with the presence of specialized equipment on the property used to conduct cryptocurrency mining operations, some of which was foreign-made and was found to present national security concerns, including given its potential capability of facilitating surveillance and espionage activities.

It is important that data center developers analyze CFIUS risk, including potential mandatory filing requirements, if they **plan to source financing from foreign persons**, including entities directly or indirectly controlled by foreign individuals, entities, or governments.

CFIUS ordinarily clears transactions involving investors or acquirers headquartered in countries allied to the U.S. However, CFIUS may condition clearance of sensitive transactions on mitigation arrangements, such as modifications of transaction structure, contractual commitments from the transaction parties to do (or not do) certain things, or both. Transactions involving acquisition of, or a significant investment in, a U.S.-based data center by a party headquartered in, or with ties to, China are likely to be intensely scrutinized by CFIUS and often will not be cleared by CFIUS.

Developers of U.S.-based data centers should **conduct due diligence on potential foreign investors**, including their ownership and control, as well as their commercial relationships and other connections to China and Russia. In May 2023, CFIUS issued guidance confirming its position that it has the authority to request information "with respect to all foreign investors that are involved, directly or indirectly, in a transaction, including limited partners in an investment fund" and about "governance rights and other contractual rights that investors collectively or individually may have in an indirect or direct acquirer or the U.S. business."

CFIUS continues to expand its efforts to identify transactions within its jurisdiction that parties have not notified to CFIUS and that may present national security risks. CFIUS may contact transaction parties post-closing, even years later, to request information and potentially a filing.

8. Conclusion

The convergence of hyperscale data center growth, clean energy demand, and evolving capital markets is reshaping infrastructure development at an unprecedented pace and scale. Parties must navigate a complex, high-stakes environment characterized by grid constraints, interconnection delays, and increasing regulatory scrutiny — while responding to a global push for sustainable, always-on digital infrastructure. Delivering 24/7 reliable and carbon-free power to these facilities requires innovative energy contracting mechanisms, including 24/7 clean utility tariffs, build-transfer agreements, behind-the-meter solutions, and evolving utility partnership models. Equally critical is the need for agile financing tools — ranging from project finance to securitization — that can accommodate both the front-loaded capital intensity of development and the stable, long-term cash flows of operational data centers.

Meanwhile, developers must contend with not only the technical and economic challenges of large-scale infrastructure deployment but also reputational and political considerations that can shape community support and regulatory outcomes. Securing a social license to operate now requires credible, location-specific strategies around land, energy, and water use. The developers and investors who succeed in this space will be those who bring an integrated view — one that aligns power procurement, infrastructure financing, and community engagement with the evolving demands of the digital economy. As this report highlights, a forward-thinking, legally sophisticated, and operationally flexible approach is essential to powering the next generation of data infrastructure.

An Integrated Platform for Data Centers

We offer comprehensive legal guidance at every stage of the data center lifecycle, from power procurement, offtake, capital raising, leasing and M&A to deploying new technologies.

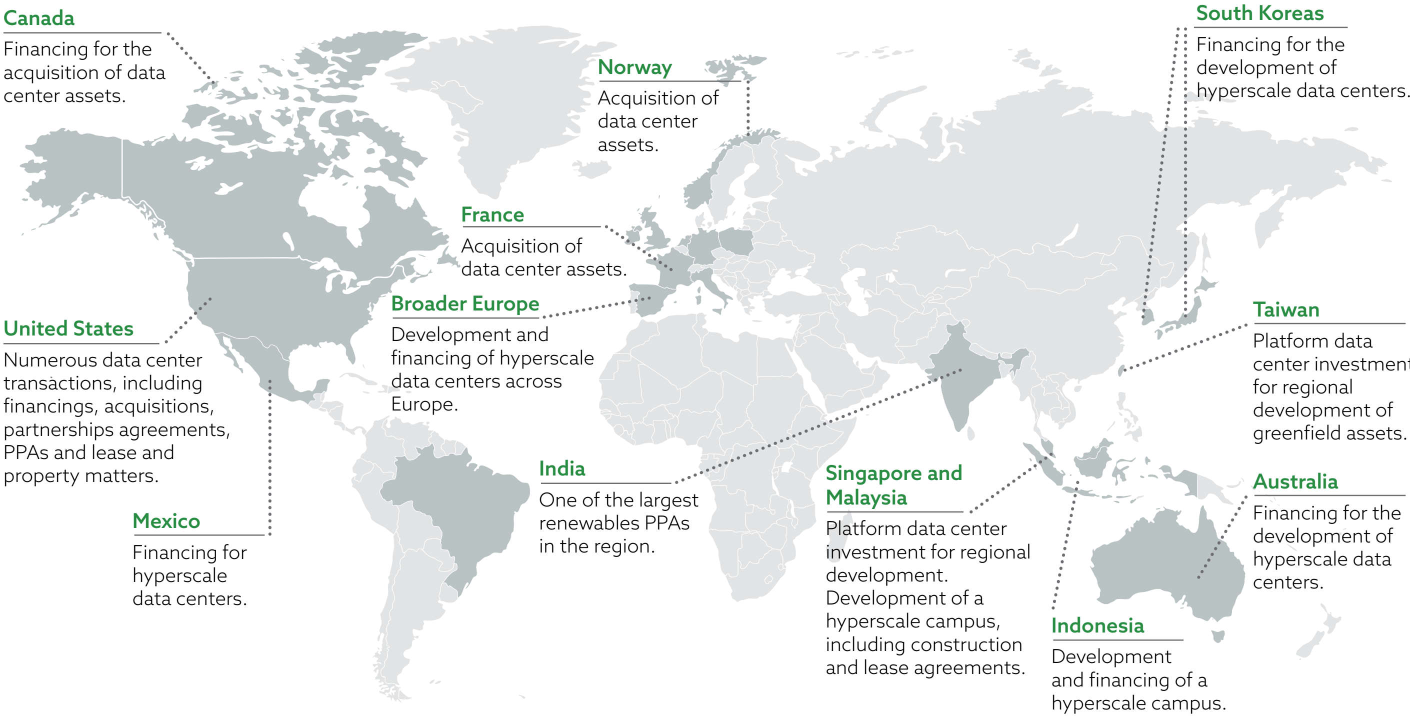

PLAN | DEVELOP | CONSTRUCT | OPERATE

Power Solutions

- Behind-the-meter
- Clean Transition Tariffs
- Utility Sleeves
- Solar, Wind, Nuclear/SMR, Geothermal, Storage
- Gas-fired
- PPAs and VPPAs

Corporate Advisory

- Finance
- Leasing
- Construction
- M&A and Joint Ventures
- International Trade/CFIUS
- Disputes
- Cybersecurity and Data Privacy

Orrick is a leading global law firm with offices in 25 global markets servicing the Energy & Infrastructure, Technology & Innovation, Finance and Life Sciences sectors. Our cross-disciplinary teams provide comprehensive support for data center operators, developers, investors and power developers. Over the past 15 months, we have supported more than 25 innovative data center transactions — including six landmark PPAs — across 20 countries.

inspiratia
The Energy Transition Specialist

#1 Global, North America, Europe, Asia-Pacific PPA Advisor

#2 Global Energy Transition Firm

FT FINANCIAL TIMES

19 Financial Times Innovation Awards
For Energy Sector Projects *2010 – 2024*

Chambers AND PARTNERS

Band 1 – US 2025
Public Private Partnerships
Energy Transition
Renewables & Alternative Energy
Power & Renewables: Transactional
Corporate/M&A – Northern California
Public Finance

The Authors and Editors of This Report

Powering Data Centers

Jonathan Ayre
Partner,
Energy & Infrastructure
Houston
jayre@orrick.com

Lowell Chandler
Associate,
Energy & Infrastructure
Washington, D.C.
lchandler@orrick.com

Kyle Drefke
Partner,
Energy & Infrastructure
Washington, D.C.
kdrefke@orrick.com

Matthew Gluschankoff
Law Clerk,
Energy & Infrastructure
Seattle
mgluschankoff@orrick.com

Kayla Grant
Senior Associate,
Energy & Infrastructure
Washington, D.C.
kgrant@orrick.com

Teresa Hill
Partner,
Energy & Infrastructure
Seattle
teresa.hill@orrick.com

Giji John
Partner,
Energy & Infrastructure
Singapore/Houston
gjohn@orrick.com

Cory Lankford
Partner,
Energy & Infrastructure
Washington, D.C.
clankford@orrick.com

Lana Le Hir
Partner,
Energy & Infrastructure
Orange County
llehir@orrick.com

Rohit Sachdev
Partner,
Energy & Infrastructure
San Francisco
rsachdev@orrick.com

Hannah Schwartz
Associate,
Energy & Infrastructure
Los Angeles
hschwartz@orrick.com

Dahl Thompson
Partner,
Energy & Infrastructure
Houston
dahl.thompson@orrick.com

Alfred Weightman
Partner,
Energy & Infrastructure
London
aweightman@orrick.com

Paul Zarnowiecki
Partner,
Energy & Infrastructure
Washington, D.C.
pzarnowiecki@orrick.com

Leasing Considerations

Peter Bergan
Partner,
Real Estate
Chicago
pbergan@orrick.com

Ari Bessendorf
Partner,
Energy & Infrastructure
Singapore
ari.bessendorf@orrick.com

Financing

Walter Alarkon
Senior Associate,
Energy & Infrastructure
New York
walarkon@orrick.com

Janet Barbiere
Partner,
Structured Finance
New York
jbarbiere@orrick.com

Devin Brennan
Partner,
Public Finance
San Francisco
dbrennan@orrick.com

Elizabeth Elias
Partner,
Structured Finance
New York
eelias@orrick.com

Zachary Finley
Partner,
Banking & Finance
San Francisco
zfinley@orrick.com

Adam Griffin
Partner,
Energy & Infrastructure
New York
agriffin@orrick.com

Andrew Hughes
Partner,
Energy & Infrastructure
Tokyo
ahughes@orrick.com

Romina Montellano Yoerg
Managing Associate,
Structured Finance
Washington, D.C.
rmontellanoyoerg@orrick.com

Mitchell Naumoff
Partner,
Structured Finance
Washington, D.C.
mnaumoff@orrick.com

Katya Popoff
Senior Associate,
Structured Finance
New York
epopoff@orrick.com

Jonathan Roheim
Partner,
Real Estate
Los Angeles
jroheim@orrick.com

Emily Sabo
Senior Associate,
Structured Finance
New York
esabo@orrick.com

Adam Smith
Partner,
Energy & Infrastructure
London
adam.smith@orrick.com

John Stanley
Partner,
Public Finance
San Francisco
jstanley@orrick.com

Mergers & Acquisitions

Walter Alarkon
Senior Associate,
Energy & Infrastructure
New York
walarkon@orrick.com

David Gold
Partner,
M&A and Private Equity
New York
dgold@orrick.com

Kristin Seeger
Partner,
Energy & Infrastructure
San Francisco
kseeger@orrick.com

Patrick Tardivy
Partner,
M&A and Private Equity
Paris
ptardivy@orrick.com

Tax Considerations

John Eliason
Partner,
Energy & Infrastructure
Washington, D.C.
jeliason@orrick.com

International Trade & Compliance

Greg Hume
Economist & Compliance Specialist,
International Trade and Investment
Washington, D.C.
ghume@orrick.com

Elizabeth Zane
Partner,
International Trade and Investment
Washington, D.C.
ezane@orrick.com