pillsbury

Guide to Data Centers

MAY 2026

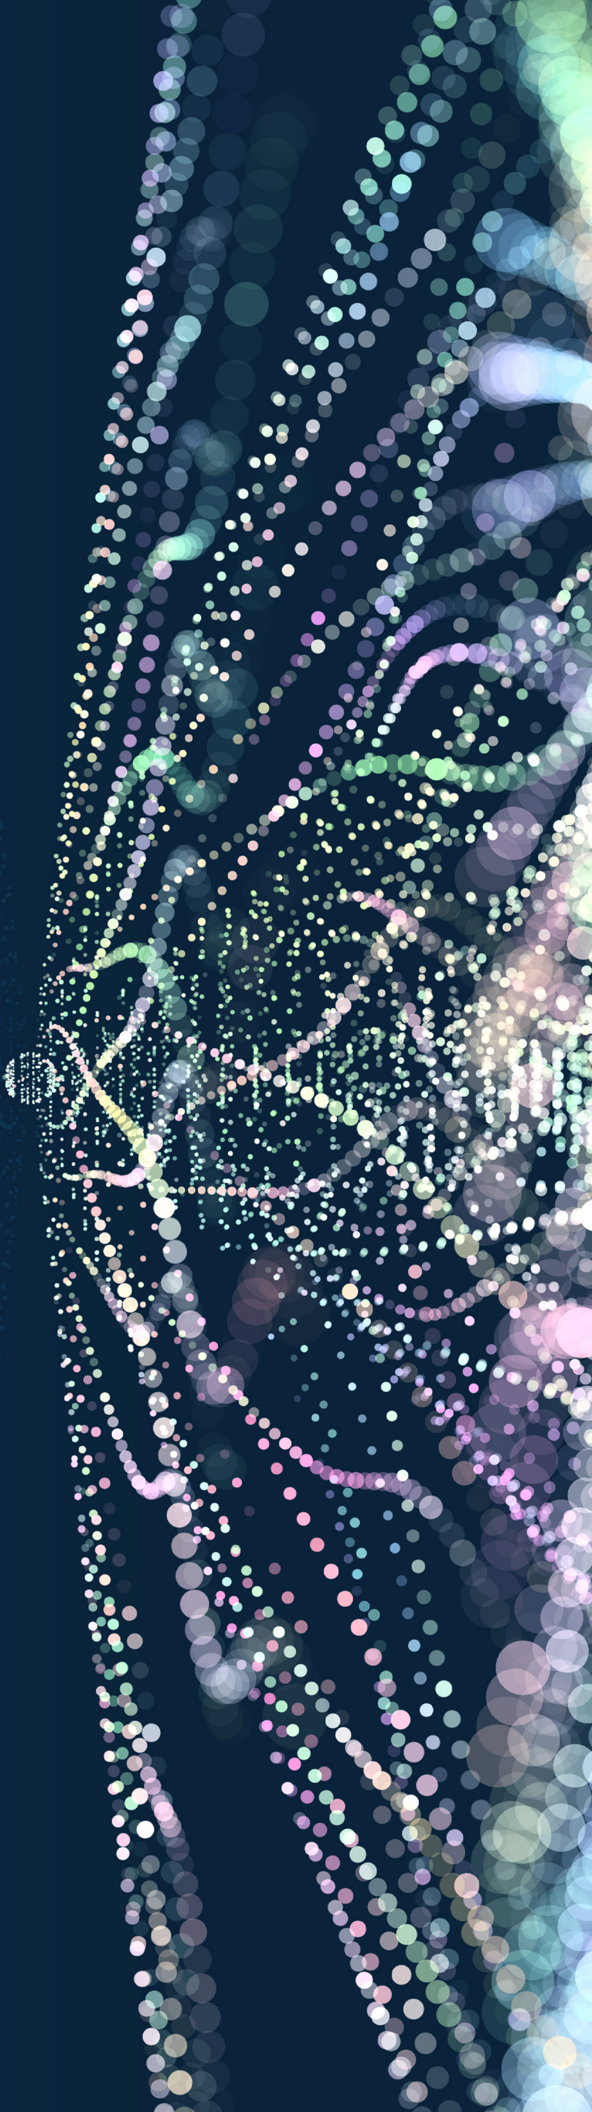

Pillsbury Winthrop Shaw Pittman LLP

Contents

Pillsbury Winthrop Shaw Pittman LLP | 31 West 52nd St. | New York, NY 10019 | 888.387.5714

Introduction

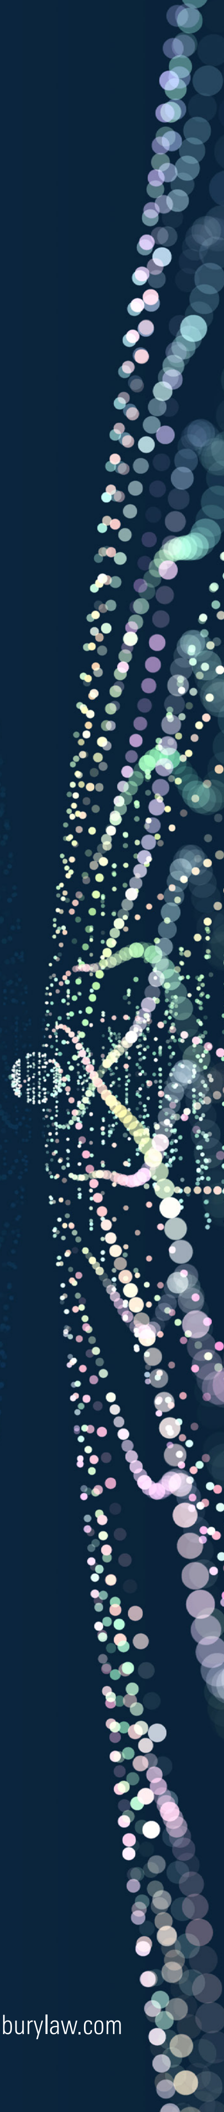

2026 Edition

Since the publication of the 2025 edition of the *Pillsbury Guide to Data Centers*, the market has continued to evolve—most notably with respect to power availability, energy strategy, tax and incentives planning, and investment activity across the sector. While many of the legal, commercial and regulatory frameworks addressed in the original *Guide* remain durable and relevant, recent developments warranted targeted updates and additions.

The 2026 edition expands and updates our energy-focused content to reflect the increasingly central role of power procurement, including an analysis of **state regulatory approaches to powering data centers**. We have incorporated new materials addressing **power purchase and interconnection agreements**, solar and other renewable energy solutions, advanced reactor designs, and nuclear-powered data centers projects, including an **updated data centers project tracker** and **permitting base checklist** for data centers and power plants.

We have also added new analysis covering **state and local tax considerations and incentive structures** relevant to data center development and operations, data center **development in the UK and EU**, and **data centers in space**.

The updated *Guide* is designed to provide a current, practical view of how data centers projects are being structured, financed, incentivized and executed today—and where the market is headed next.

— **Brittney Sandler**, Editor

Introduction

***by* Rob James, Data Centers Team Co-Leader**

As a law firm, Pillsbury's experience with data centers dates to the dot-com era, when I advised on buildouts in Silicon Valley and New York City and Jamie Bobotek advised on centers of America On-Line (then spelled out!) and others in Northern Virginia. We counseled the industry as proprietary centers were supplemented, and in some cases supplanted, by physical and cloud-based systems of larger dedicated enterprises. The data center landscape is now international in scope and diverse in speciation, with edge, modular, enterprise and hyperscale facilities each playing important yet distinctive roles.

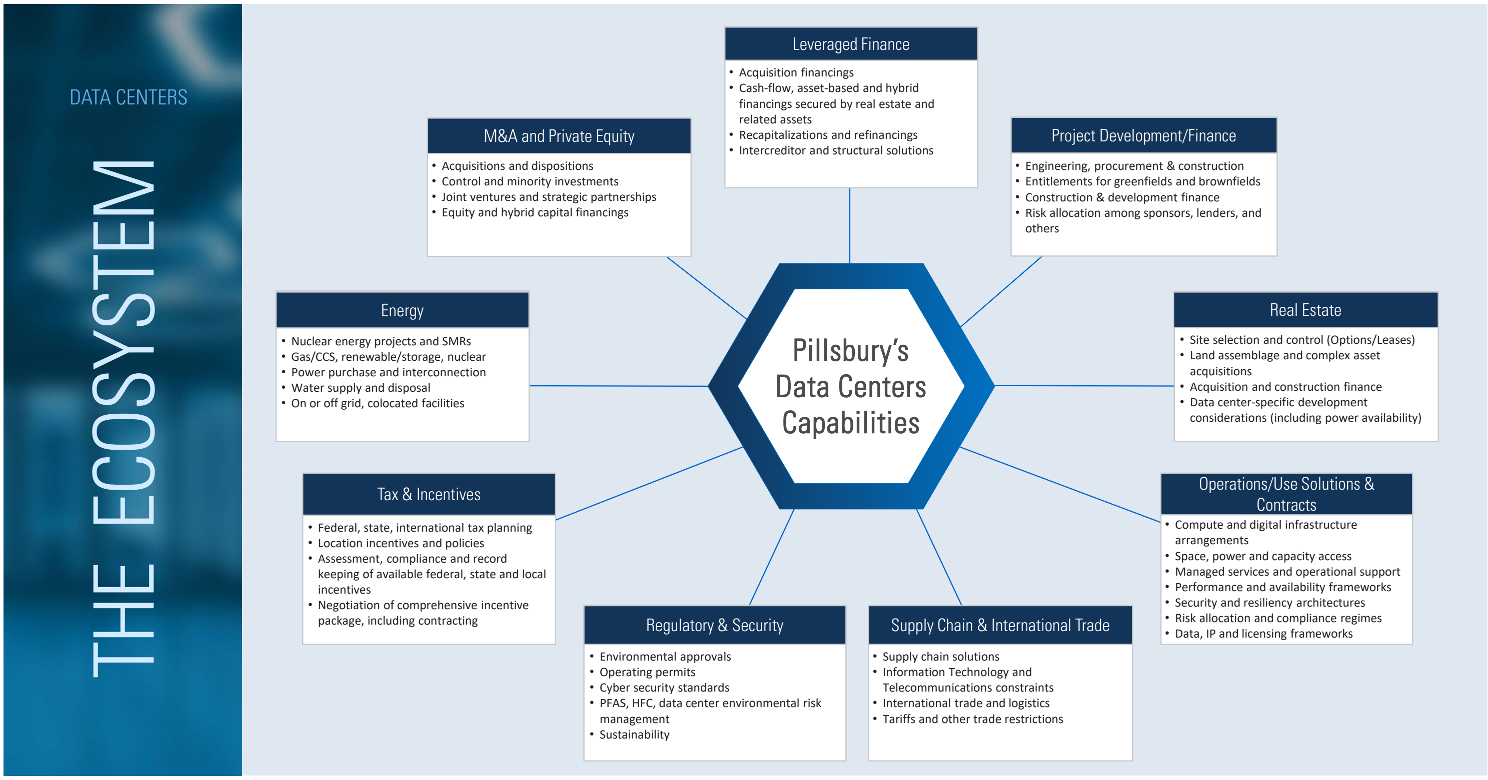

DATA CENTERS

THE ECOSYSTEM

Pillsbury's Data Centers Capabilities

Leveraged Finance
- Acquisition financings
- Cash-flow, asset-based and hybrid financings secured by real estate and related assets
- Recapitalizations and refinancings
- Intercreditor and structural solutions

M&A and Private Equity
- Acquisitions and dispositions
- Control and minority investments
- Joint ventures and strategic partnerships
- Equity and hybrid capital financings

Project Development/Finance
- Engineering, procurement & construction
- Entitlements for greenfields and brownfields
- Construction & development finance
- Risk allocation among sponsors, lenders, and others

Energy
- Nuclear energy projects and SMRs
- Gas/CCS, renewable/storage, nuclear
- Power purchase and interconnection
- Water supply and disposal
- On or off grid, colocated facilities

Real Estate
- Site selection and control (Options/Leases)
- Land assemblage and complex asset acquisitions
- Acquisition and construction finance
- Data center-specific development considerations (including power availability)

Tax & Incentives
- Federal, state, international tax planning
- Location incentives and policies
- Assessment, compliance and record keeping of available federal, state and local incentives
- Negotiation of comprehensive incentive package, including contracting

Operations/Use Solutions & Contracts
- Compute and digital infrastructure arrangements
- Space, power and capacity access
- Managed services and operational support
- Performance and availability frameworks
- Security and resiliency architectures
- Risk allocation and compliance regimes
- Data, IP and licensing frameworks

Regulatory & Security
- Environmental approvals
- Operating permits
- Cyber security standards
- PFAS, HFC, data center environmental risk management
- Sustainability

Supply Chain & International Trade
- Supply chain solutions
- Information Technology and Telecommunications constraints
- International trade and logistics
- Tariffs and other trade restrictions

There is no single locus of data center expertise in legal practice, just as there is no single locus of capabilities and users in the business environment. Pillsbury's Data Centers team accordingly deploys a one-stop, hub-and-spoke model that reflects the complex ecosystem in which these localized points of computation sit.

(To let you in on a little secret, Pillsbury employs this same hub-and-spoke approach to many other forms of infrastructure and energy projects—crypto mines, order fulfillment centers, global communication centers, you name it!)

That same philosophy and a similar structure inform this distinctive publication, the *Pillsbury Guide to Data Centers*. It contains articles that we hope are not the typical ephemeral client alert of some episodic deal, proposed bill or trend, but instead are enduring reference pieces that lawyers and clients of all levels of sophistication may wish to keep handy as they encounter questions and challenges in this field.

This eBook brings together a curated selection of articles authored by our firm's interdisciplinary team. They reflect our experience advising clients on the development, acquisition, financing, construction and operation of data centers across the U.S. and internationally. Each chapter can be tied back to one or more of the spokes on our hub-and-spoke system.

What Is a Data Center, Anway?

As my dot-com anecdote indicates, even by that name "data centers" have been around for decades. What has thrust them into the headlines today? There has been a confluence of the needs for massive data processing and communication demands of the digital economy, on the one hand, and dramatic increases in the processing speed and memory of graphics processing units (GPUs) and other hardware and software systems on the other hand. The needs of business-to-business and business-to-consumer electronic commerce and the use of massive internet platforms for external transactions and internal administration began a long time ago and show no signs of relenting. Applications such as cloud computing, the internet of things, software as a service, machine learning and generative artificial intelligence have compounded the needs for localized centers capable of receiving, processing, storing and transmitting amounts of data that were unimaginable just a few quarters ago.

At the same time as this increase in data processing, we have witnessed an increase in the physical and cybersecurity risks associated with safeguarding this information, and an explosion in the regulation of the privacy, data ownership and other aspects of the industry. These regulations emanate from specific jurisdictions, but the information technology industry spans every possible boundary, raising complex questions of compliance for every enterprise—both proprietors and users of data centers face these issues.

What exactly is a data center? If you knocked on the door of one, who would answer, and if you were admitted inside (a tall order), what would you see? Read Pillsbury's "**Anatomy of a Data Center**" for a light-hearted survey.

The Species of Data Centers

An admittedly unscientific survey says that there are about 11,000 "data centers" worldwide, about half of which are in the United States. Even within the United States, data centers began in Silicon Valley and Loudon County, Va., but we are now seeing centers and proposals anywhere energy, water and entitlements are hospitable. The global market follows similar patterns, with Digital Realty, Equinix and CoreSite among the leading colocation providers, alongside hyperscale cloud operators such as AWS, Microsoft, Google and Meta that increasingly drive demand for large-scale data center capacity.

These centers range in size from half an office building floor to some of the largest manmade structures on the planet. Read Pillsbury's "**Data Centers: A Field Guide**" to get more of a breakdown of the species in this genus.

Customer-Facing Solutions: Service and Colocation Agreements

A unique feature of Pillsbury's data centers practice is that we represent both proprietors and users. From either side, a client needs a comprehensive understanding of both the service needs of the customer and the system capabilities of the center. Accordingly, we group our services under the term **Customer-Facing Solutions** to capture the need to understand *both* sides of a harmonious deal.

The core customer relationship is typically structured either as a service contract in which the data center enterprise owns the computation facilities and provides a service, or a co-location contract in which the center hosts equipment owned by or leased to the customer. In either model, attention is focused on the service-level expectations of the customer and the corresponding qualities of the data center.

Read Pillsbury's "**From Anatomy to Action: Navigating Data Center Contracts**," authored by Pillsbury legal personnel as well as a Pillsbury business consultant, for an in-depth look at these complex commercial agreements.

Energy and Water for Data Centers

Power and cooling are the nectar and ambrosia on which data centers feed. They are also the gating factors in determining whether, when and where a data center will be built, and accordingly are top of mind not only to the hyperscalers and other proprietors, but also to the entire ecosystem of real estate developers, financers and investors in the space.

Among the most promising sources of power for data centers is nuclear generation, but few can articulate the advantages and challenges of developing advanced fission and fusion technology. In "**Powering Data Centers with Nuclear Generation**," Pillsbury's highly recognized Energy Transition nuclear and project lawyers combine to do just that.

As the data centers expand in both traditional and new locales, they face grid interconnection challenges, local utility and ratepayer concerns, and expectations of many proprietors for ultra-stable and zero-emission energy sources. One interesting alternative that can finesse many of the siting challenges is the location of both data centers and power sources on federal lands. The recent initiative to co-locate such facilities on property of the Department of Energy is one such solution. Read Pillsbury's chapter on **data centers on Department of Energy facilities** for information about the initial steps.

One tension in the competition for zero-emission energy sources is that data centers face competition from other users of that form of power—for carbon removal, hydrogen production and other facilities that are driven by green mandates or corporate commitments. Read Pillsbury's "**Big Data Meets Big Green**" to learn about that interaction among organizations with climate focuses.

Investing in Data Centers

Pillsbury represents a wide variety of real estate investors, financers, investors and advisers on the complex aspects of selecting the optimal site, including land assemblage, due diligence on land use and environmental regulations, access to the grid, energy and cooling sources, and community reactions. Tax considerations and local incentives (or disincentives) strongly factor into the analysis. Read Pillsbury's "**Investing in Data Centers**" for an in-depth survey of the players in this field. And take a look at the Appendices to this *Pillsbury Guide to Data Centers* for additional detail on the specific places where existing and proposed data centers, and the inputs essential to their development, are located or being considered.

Environmental, Privacy and Security Regulation of Data Centers

Data centers are situated at the crossroads of a bewildering number of public policy concerns. Pillsbury extensively counsels proprietors and users on data security, privacy and operational standards. An oft-overlooked area in identifying data center legal support are the environmental challenges associated with materials used in their operation. Read Pillsbury's "**PFAS, HFCs and Other Chemicals in the Data Center Industry**" to appreciate some of the considerations and the need for experienced counsel.

Designing, Constructing and Converting Data Centers and Crypto Mines

Pillsbury's experience with engineering, procurement and construction (EPC) or construction management (EPCM) of data centers stretches back decades. That body of knowledge extends now to the new build of such centers and the potential conversion of crypto mines to data centers, or vice versa. Read Pillsbury's "**Designing, Constructing and Converting Data Centers and Crypto Mines**" for an overview of the brick-and-mortar aspects of development and redeployment of these localized points of computation.

– **Rob James**, Editor

For more information contact us
at **datacenters@pillsburylaw.com**.

1

Anatomy of a Data Center

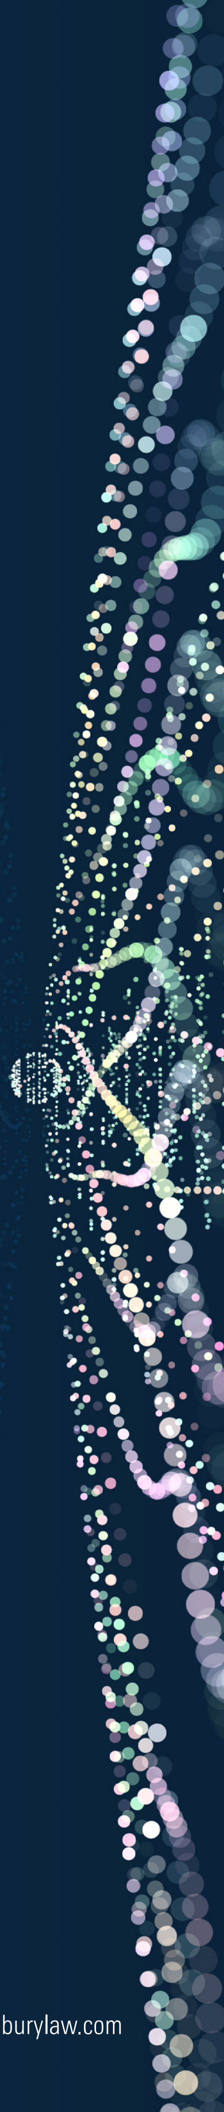

Anatomy of a Data Center

by **Rob James & Matt Olhausen**

Traditional and social media are thick with reports and predictions of the remarkable increase in size, power consumption and significance of data centers. Not only technology companies but real estate and energy developers, investment funds, lenders, and professionals of all stripes are in or determined to enter this sector. Our inboxes are full—it's data center this, data center that.

But what exactly ***is*** a data center? What infrastructure, technology and human resources come together to create and sustain one of these localized points of computation? By understanding their components, we can glean some understanding of the business, public policy and (our focus) legal issues that arise before and during their operation.

In this article, we cite key characteristics of a reference **Blackacre Data Center**, with occasional glances at other (real) structures that offer variations on themes. Blackacre is a composite of several centers we have encountered in our law practice. These facilities differ widely in size, location and functions, so your mileage will vary.

INFRASTRUCTURE RESOURCES

Blackacre Data Center is built on a 14-acre site consisting of two contiguous legal parcels that were formerly developed as Class C light industrial warehouses in the early 1970s, located on the outskirts of a major United States technology-center city. The property was acquired and developed by a real estate developer specializing in industrial projects, including data centers, and the fee owner is a special-purpose entity subsidiary of an investment fund. The acquisition and construction were **funded by equity investment (aided by tax credits and location incentives) and third-party debt financing**. The improvements were built under a **design-build contract** with an engineering and construction company joint venture. The fee estate is leased to an operating company affiliated with the developer, which then leases interior space to technology companies.

The Blackacre site was selected on the basis of many factors, including real estate cost and availability; telecommunications connectivity; availability of reliable, sustainable power; resilience to natural disasters; ambient temperatures or water supply suitable for cooling systems; taxes and incentives; and human factors to attract employee and contractor workforces. (There is a publication, **ANSI/TIA 942 Annex**, with data center site selection factors.)

To entitle Blackacre, the developer underwent a streamlined planning process with the city, since the project site is located within an existing industrial park within the city's designated "Industrial Technology and Innovation Corridor." The zoning permits a wide range of uses, including, office, business park, industrial, research and development, manufacturing, and information and technology-based uses. Fortunately for Blackacre compared with other projects, the site is bordered by undeveloped land and other industrial and commercial uses, remote from residential districts, and located near a state highway.

Data center businesses often employ one of two approaches—either an enterprise model in which the servers, storage and network switches are controlled by a tenant technology company, or a colocation model in which information technology (IT) equipment and utilities supporting that equipment are leased to individual businesses. Blackacre follows the colocation model.

The outside perimeter is an imposing 10-foot-tall security fence, with gated site access staffed 24/7 and the first of many closed-circuit television (CCTV) cameras. The grounds include parking for 73 cars and a dozen bicycles, outside electricity stepped down to an onsite 75 MVA (mega volt-ampere), 25,000-square-foot substation, a 25,000-square-foot utility switching station, and telecommunications outlets (both satellite dishes and a dedicated T1 fiber optic line). The substation has two transformer lines allowing one to be taken out of operation without interruption of service.

There are two data center structures located on the campus: a three-story main building (350,000 gross square feet) and a one-story auxiliary building housing storage media (15,000 gross square feet). There is also a backup generator yard; a 3,000-square-feet security building; and infrastructure improvements such as fire detection and suppression devices, access driveways, stormwater facilities, and water storage tanks.

On the entry floor of each data center building, there is a lobby to verify access and monitor entry and exit by authorized individuals. Security measures include biometrics and badge access into and around the facility (including the equipment areas), key access to specific racks and servers, logs for employee and visitor/vendor access, physical escorts for all non-employees, video surveillance, and 24/7 on-site security personnel. **Cybersecurity** is maintained through firewalls, virus detection software, encryption, disaster and breach security, backup and recovery systems, and regular tests and audits. Blackacre security personnel are mindful of industry security standards such as **ISO/IEC 27001** and **data privacy regulations and best practices**.

Authorized entrance from the lobby leads to an operations center with monitors of all the data center activities. From that center authorized personnel can next enter the computer rooms, where the servers, storage, network switches and other equipment are located, operated and maintained. When you enter a computer room, besides the hum and whir of storage devices, you will hear the whoosh of cool air being fed and hot air being withdrawn in separate currents called “cold aisles” and “hot aisles.”

Blackacre has visions of greater and greater sustainability. Over the longer term, **zero-emission power** is to be supplied for the data center from advanced energy systems; its owners are in negotiations with developers of small modular reactor (SMR) nuclear fission, natural gas with carbon capture, advanced geothermal, sodium-based storage from renewables, and even nuclear fusion, to be located at an adjacent undeveloped project site. This power may be delivered directly physically to the site, or **virtually through a power purchase agreement**.

INFORMATION TECHNOLOGY RESOURCES

The heart of Blackacre Data Center is its information technology, which falls into three broad categories: servers, storage devices, and network devices.

Servers are the computers that perform the customers’ desired functions with respect to the data that come to and leave from the data center. Server operation entails the retrieval, storage, and processing of data supplied by the technology company or generated by its user base in engaging with internal and customer-facing application software. These servers have a range of internal memory, processing power, and other specifications, requiring customized support by Blackacre staff. Servers can be self-contained computers installed in racks, or slimmed-down “blades” more dependent on Blackacre support and utilities.

Ancillary to the servers are large storage devices, both hard drive and solid state. They store data in block increments with many terabytes of data capacity. (Lawyers should start getting used to the next thousand-power prefixes: peta, exa, zetta and yotta.) Lastly, a data center relies on network infrastructure. Switches transfer data between nodes on a network; routers transfer data from and to networks; and firewalls and load balancers complement their functions.

Blackacre also has monitoring computers, system software, and application software. The various IT components and technology are partly proprietary to the operating company and partly licensed to them and their customer base. With virtualization and greater employment of cloud resources, the lines separating server, storage, network and other devices are blurring, and the Blackacre configurations will change with the times to keep up with the larger enterprise “hyperscale” data centers.

POWER AND COOLING RESOURCES

Data centers are well known voracious consumers of electricity, not only for powering the computers but also for cooling them down and maintaining the rest of the integrated facility. A metric called Power Usage Effectiveness (PUE) is used to compare the power needs of the IT equipment with the overall power needs of the center. Within that energy budget, there is a target power usage of so many kilowatts per rack or per blade within the facility. The processing power of individual chips and the energy density of servers are increasing at a torrid pace, raising Blackacre's energy needs rack by rack and upgrade by upgrade.

The absolute requirement for a data center is guaranteed power at "five nines" (99.999%) or even higher levels of assurance. The offsite utility power line comes into a transfer switch with the generator circuit, so the generators can step in within several seconds of a grid outage. That is not good enough, though, so an Uninterruptible Power System (UPS) is also connected. The UPS may be a long-life (over 12 years useful life) lead-acid or lithium-ion battery system or a continuously spinning flywheel. (See **APC White Paper No. 92** for their respective merits.) The UPS can keep critical load in steady operation until the generator comes on or the grid source resumes.

The cooling requirements of Blackacre are daunting. Even with built-in fans, the IT equipment would overheat quickly were it not for the complex circulatory systems in the main building. The American Society of Heating, Refrigerating and Air Conditioning Engineers (ASHRAE) recommends data center IT room temperatures and humidity be kept in a controlled range. While some larger centers have chilled water systems with Computer Room Air Handlers (CRAHs), others use a compressor-based system with Computer Room Air Conditioning (CRAC), in which the familiar **refrigerant cycle** of evaporation, compression, condensation and expansion of a fluid with low boiling point drives the airflow.

For Blackacre, by contrast, Data Hall Air Handling Units (DAHUs) are installed in dedicated galleries. Outside air is drawn from each building's perimeter. The DAHUs use a "flooded room" design, meaning that no ductwork or raised flooring systems direct the cooling air to the IT racks' air intakes. Instead, all DAHUs in each mechanical gallery discharge cooling air into the adjacent common supply air header. The potable water demand for the entire site is relatively modest, with half dedicated to evaporative cooling and the balance to landscaping and domestic needs.

Continued expansion of the data center sector to accommodate AI and other demands will depend on ever more efficient designs and operations of server components, innovative power systems, and cooling systems that minimize energy losses. Pillsbury advised on contracts for Stanford University's **award-winning central energy plant** using heat exchangers to capture excess heat from data centers to warm dormitories. At locations more remote than Blackacre, centers may be able to take advantage of **geothermal temperature gradients** for both power and cooling.

HUMAN RESOURCES

Employee and contractor workforces are a critical part of Blackacre's operation. The campus currently employs 52 local residents on a full-time basis, with monitoring, security and maintenance shifts around the clock 24/7. Representatives of the industrial real estate company, the customer base, and equipment suppliers visit frequently and upgrade the technology assets and security systems. Maintenance contractors are also engaged on a periodic basis, all with the objective of uninterrupted performance at highest industry levels.

Personnel (sometimes even working with their attorneys on **Pillsbury's data centers team**!) also maintain compliance with laws, contracts, permits and industry practices. Relevant standard-setting organizations include ASHRAE, the Uptime Institute, the Telecommunications Industry Association (TIA, which has a comprehensive family of data center specifications **TIA-942**), and the Building Industry Consulting Service International (BICSI). Available sustainability certification sources include LEED, Energy Star, and Building Research Establishment Environmental Assessment Methodology (BREEAM).

CONCLUSION

Everyone talks about data centers, but few can fully describe one. They indeed sometimes even *look* like "black boxes"; the architects do not typically splurge on windows and aesthetic touches. The need for security limits the amount of visibility into how they are built and operated. There are many variations on Blackacre, and this article is no substitute for learning about any given facility's features. We nonetheless find it useful to have a reference to compare to our new projects in this rapidly expanding field. ❖

2

A Field Guide to Data Centers

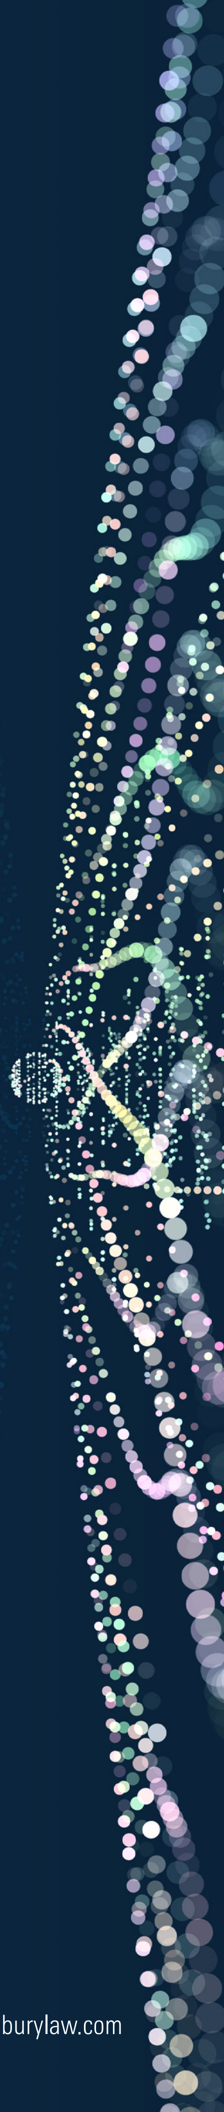

A Field Guide to Data Centers

***by* Rob James, Matt Olhausen & Sam Bucher**

By **one count**, worldwide there were some 11,800 data centers in early 2024. Within that census are facilities so small that they fit in office building closets, while others are among the largest manmade structures on the planet. How are we to make sense of this diverse population?

Data centers house servers, storage devices and network devices to store, process and disseminate large quantities of data of organizations and their customers and supply chains. The size and complexity of these facilities vary with their functions in the business ecosystem. Much like the well-known depictions of the **evolution of horses** from tiny brush creatures to mighty stallions, the overall category cries out to be broken down along multiple dimensions.

This post provides a naturalist's field guide to data center types and features.

Enterprise Data Centers

Enterprise data centers are operated by a single organization, either through direct ownership or a leasing arrangement, to support its internal information technology (IT) operations. These are the traditional and longstanding center types, often running on a private cloud. They have customized infrastructure to meet one company's needs, giving it full control over security, compliance standards, and system optimization.

Organizations that can afford to operate an enterprise data center benefit from the ability to fine-tune the environment, ensuring governance over sensitive data, and adhering to industry-specific regulations. They pose challenges such as the need for dedicated staffing and robust backup solutions to mitigate downtime, as well as generally high costs not spread across other users.

Enterprise data centers can be either on-premises or off-premises of the organization they serve. On-premises centers offer greater assurance of reliability and security. In contrast, off-premises centers can be a cost-effective alternative for organizations with less stringent security and speed demands, reducing upfront infrastructure and real estate costs, and alleviating the workload for in-house IT staff.

These enterprise centers can be very modest in scale. One private research university, for example, runs a 12,000-square-foot facility—6,000 square feet of data center space and 6,000 square feet of infrastructure space—on 450 KW of power and cooling. The center costs $12.4 million and hosts the institution's private cloud that supports campus servers.

There are enterprises, and then there are *enterprises*—how about the NSA? The National Security Agency operates a high security center powered by 65 megawatts of electricity in Camp Williams near Bluffdale, Utah. The $1.5 billion campus encompasses one million square feet with a 100,000-square-foot mission critical data center. The other 900,000 square feet in the campus is used for technical support and administrative space. This center facilitates the NSA's efforts, as the executive agent for the Office of the Director of National Intelligence, to monitor, strengthen and protect the nation with regard to digital data.

Colocation Data Centers

Since operating an enterprise data center is costly and requires dedicated expertise, many users turn to the **colocation data center** model. Colocation offers storage and support for computation essentially as a service. These third-party facilities lease space to businesses for their servers, often at a lower initial and potentially reduced operational cost compared to on-premises data centers. The centers host multiple businesses as tenants, each renting the amount of space according to its specific needs—either a full cabinet or often just a fraction of a cabinet.

The centers also supply essential infrastructure, including power, cooling, maintenance services and network connectivity. For organizations lacking around-the-clock expert staff or seeking to minimize upfront investment, colocation can be a practical choice. Additionally, these facilities can offer scalability to meet market demand, enabling businesses to expand their operations as space and resources allow in that specific colocation center.

Colocation can be specifically attractive to e-commerce and financial services companies through faster load times and handling increased traffic and demand due to the site owner's specialized knowledge of the IT industry. This model can allow businesses to adapt to ever-changing demands in the market and scalability if the site has additional space to lease. Health care providers and governmental bodies can depend on these multi-tenant data centers if they lack the expertise in developing and operating their own enterprise sites.

A composite example of a colocation data center is what we called the Blackacre Data Center in northern California. The site houses two data center structures, a three-story main building consisting of 350,000 square feet and a 15,000-square-foot one-story auxiliary building. For an inside look at exactly what goes on in Blackacre Data Center, see our recent post, **Anatomy of a Data Center**.

Hyperscale Data Centers

Hyperscale data centers are vast, highly scalable facilities—often exceeding a thousand cabinets, 10,000 servers, and a million square feet. By one 2024 census there were about a **thousand centers** worldwide in this category, half of them in the United States. Most are owned or operated for the benefit of cloud industry giants, and similar enterprises in China and elsewhere.

These facilities commonly function as enterprise data centers, operated solely by the provider for their own needs, but they can sometimes function as colocation data centers. Regardless of the ownership and operation structure, they are designed to support massive workloads and bring to bear enormous infrastructure capabilities—and enormous infrastructure demands including power and cooling.

Hyperscale facilities are engineered for resilience, featuring robust backup systems and automated failover mechanisms that activate immediately in the event of equipment failure or power loss. While these facilities can be purchased or leased, they are typically custom-built and specifically designed to meet the unique operational needs of the owner and located in rural environments.

It is difficult to convey a sense of the immense scale of these facilities. The world's **largest hyperscaler**, a $3 billion project, is located in China's Inner Mongolia region. This facility consumes 150 MW of power and spans almost 11 million square feet, or the size of 165 regulation U.S. football fields.

A Reno hyperscale is also the single largest colocation data center facility in the world, occupying 1.3 million square feet drawing up to 130 MW of power. It is the first facility to be built within a new data center campus targeting up to 7.2 million square feet of data center space and 650 MW of renewable power.

A hyperscale facility in Mesa, Ariz., occupies 1.3 million square feet, consumes 50 megawatts of renewable power, and cost a reported $2 billion. And **another center** under development in Louisiana occupies land that would constitute a sizable chunk of Manhattan.

Edge Data Centers

Some cutting-edge applications, such as autonomous vehicles, cloud gaming and instant grocery delivery, require data to be hosted close to users to achieve high-speed network functionality. Edge data centers offer a relatively new solution to these challenges. They are often smaller, decentralized facilities located close to the "edge" of a larger telecommunications and IT network.

The unique location and footprint of these centers allow organizations to reduce latency and optimize bandwidth for new applications. Further, because edge data centers lower network bandwidth usage, they can significantly reduce the cost of data transmission and routing—lessening the required amount of expensive circuits and interconnection hubs. Edge data centers are often deployed in various urban locations, such as at a telecommunications central office or at the base of cell towers throughout a major city.

One services provider recently launched a number of edge data centers designed to extend their infrastructure closer to large urban populations. This venture allows visual effects, video and other applications that require single-digit millisecond latency or local data processing to achieve their goals by bringing infrastructure closer to the end user.

Evolutionary Trends

Data centers are categorized not only by size but also by stability of uptime. A **Tier Classification System** runs from Tier 1 with 99.671% uptime and a maximum of 28.8 hours of annual outages to Tier 4, for the most mission-critical applications, with 99.995% and only 26.3 minutes of permitted outage. Getting that last bit of assured uptime is very expensive and demands considerable redundancy of all systems.

An alternative for enterprise and edge centers is the **modular data center**, which can be added or relocated to existing facilities. These prefabricated facilities can finesse issues with real estate acquisition and permitting and allow users to fill interim needs while larger permanent facilities are being developed.

Data centers are running up against economic and land use constraints in their traditional habitats. The increasing size of each center accentuates this challenge; from 2022 to 2024, the average data center parcel grew by 144% to 224 acres. As of October 2024, the weighted average cost of data center land stands at $244,000 per acre. Moreover, the ability of the centers or the local utility grid to support the prodigious power requirements of the data center population is leading to constraints and delays on entitlements and approvals of centers in overheated markets. Data centers are engaged in active management of community, government and public relationships to address the externalities associated with the facilities.

The headwinds in some locations are leading the data center investment and proprietor community to look far beyond the original settings. Thanks to improvements in processing speed and telecommunications capabilities, the largest centers are being situated anywhere in the world that they can find sufficient power, accommodating land acquisition and permitting conditions, and workforces and supply chains suitable for large scale data processing.

Conclusion

Multiple species can thrive in distinct parts of the biosphere. Data centers exist in a diverse range of sizes, locations, operation models, and ownership structures, each tailored to meet a particular set of demands. Selecting the right location, size and model is one of the greatest strategic decisions any organization can face. ❖

3

A Permitting Base Checklist for Data Centers and Power Plants

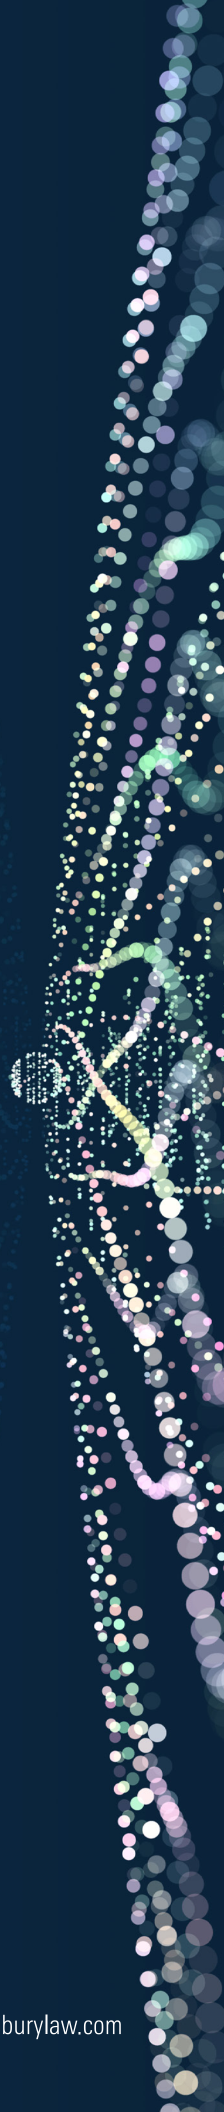

A Permitting Base Checklist for Data Centers and Power Plants

***by* Stacey C. Wright, Michael S. McDonough & Stephen J. Humes**

There is a lot of talk these days about "license to operate" for data centers, meaning management of the relationships with stakeholders and broader communities concerning both the benefits and adverse consequences of locating a facility in a particular locale. Here, we are speaking of "license to operate" more literally—namely, the legal and regulatory permitting and approval requirements for a privately owned data center whether by itself or colocated with a power generating plant.

Our Base Checklist includes generally and potentially applicable permitting requirements for development and operation, using California as an example. (Taking legal authority Frank Sinatra out of context, "if you can make it there, you can make it anywhere.") The actual requirements for a given facility would depend, in part, on local law, including planning and zoning laws and plans, and the environment of the site. Just as examples, additional permitting and mitigation requirements might apply if sensitive receptors are located nearby (e.g., noise mitigation for residential dwellings), if sensitive and protected biological resources (e.g., jurisdictional waters and/or protected species) would be impacted, or if the present or former land uses require additional measures (e.g., hazardous materials remediation, mitigation for conversion of prime farmland, or protection of cultural resources). The scope of permit requirements would ultimately be determined by the applicable regulatory agencies and by the lead and responsible agencies under the applicable state environmental land use regime—in our reference case here, the California Environmental Quality Act (CEQA).

To clarify, this Base Checklist is just that—a base. We hope (well, let's say our clients hope) never to encounter a California private project that runs into all of these permitting requirements. Naturally, entirely different concerns may apply to a project in another state, or to a project anywhere that is launched by or with a public agency or public body. We nonetheless hope that building or subtracting from this reference is of use to our future selves and to others.

At the federal level, President Trump issued an executive order (EO) on July 23, 2025 entitled "Accelerating Federal Permitting of Data Center Infrastructure" that directs federal agencies to support and expedite project development and permitting of data centers and supporting power infrastructure. "It will be a priority of my Administration to facilitate the rapid and efficient buildout of this infrastructure by easing Federal regulatory burdens," President Trump declared. "In addition, my Administration will utilize federally owned land and resources for the expeditious and orderly development of data centers. This usage will be done in a manner consistent with the land's intended purpose—to be used in service of the prosperity and security of the American people."

As a result of the EO, we have reason to believe that the federal government will be expediting environmental reviews of major data centers and streamlining or accelerating issuances of applicable permits. The federal initiatives prompted by the Executive Order should bring some enhanced predictability around permitting facilities that support especially large data centers (defined as a facility that requires greater than 100 megawatts (MW) of new load dedicated to AI inference, training, simulation or synthetic data generation).

More recently, separate **actions** by FERC and by **federal and state officials in the PJM electricity transmission market** have spurred the advancement of proposals to co-locate data centers and power generation facilities as a means to reducing impacts of the extraordinary energy demands for computation and cooling of the centers. Whether spurred by governments or by economic efficiencies or both, permitting must take into account the impacts of data centers and power plants as integrated projects, an increase in both integration and sheer scale which may resolve some public policy issues and introduce others.

The effect of these initiatives on the permit and approval processes is unclear. Political policies could smooth things at any of the jurisdictional levels, or it might lead to greater scrutiny and potential opposition, especially on the state, local and Tribal fronts.

For more information on permitting and approvals of these or other facilities, please reach out to the authors at the addresses provided in this Pillsbury Guide to Data Centers.

PERMITTING BASE CHECKLIST

Pre-Construction and Construction

Federal Permits/Approvals/Regulations

Air Quality

- Clean Air Act (CAA) implemented by local air districts (see below), including new source review and Title V operating permits.

Biological/Water Resources

- U.S. Fish and Wildlife Service (USFWS) "take" authorization, if required under the Endangered Species Act (per ESA Section 7) for protected plant and animal species, or their critical habitat; USFWS biological opinion, if required
- USFWS authorization under the Bald and Golden Eagle Protection Act or Migratory Bird Treaty Act, if applicable
- U.S. Army Corps of Engineers permit under the Clean Water Act for discharge into federal jurisdictional waters, including wetlands, if applicable (see also Regional Water Quality Control Board (RWQCB) certification requirements)
- Clean Water Act permits may be needed from US EPA to the extent the project diverts water from streams or aquifers for cooling, discharging wastewater and impacting wetlands.

Hazards and Hazardous Materials

- Federal Aviation Administration (FAA) notification or approval required for certain construction near an airport/runway or within navigable airspace
- Waste Management Permits for routine transport, use or disposal of hazardous materials (e.g., diesel fuel)

National Environmental Policy Act (NEPA) Review

- Projects on federal lands or involving federal funding require agencies to assess environmental impacts (but per EO above, agencies are directed to expedite these reviews)

Others

- Federal Communications Commission (FCC) (telecommunication)
- FAA and local airport regulations (potential interaction with air transportation, as noted above)
- National Oceanic and Atmospheric Administration (NOAA) (satellite communication)

State Permits/Approvals/Regulations

Air Quality

- Air Resources Board (ARB) Airborne Toxic Control Measure limits might apply (e.g., for emergency standby diesel-fueled engines)
- ARB Distributed Generation Certification Program (electrical generation exempt from local air district permit requirements), if applicable
- Additional requirements if ground disturbance occurs in an area of naturally occurring asbestos

Biological Resources

- California Department of Fish and Wildlife "take" authorization under the California Endangered Species Act, if applicable

Energy

- California Energy Commission (CEC) certification for power plants exceeding 50 MW capacity, or
- CEC Small Power Plant Exemption, if applicable (local permitting would apply)

Hazards and Hazardous Materials

- Hazardous Materials Business Plan, if handling or storing threshold amounts of hazardous materials during construction
- Permits for routine transport, use, or disposal of hazardous materials (e.g., diesel fuel), if necessary during construction
- Site Management Plan and/or Health Safety Plan for contaminated soil, if applicable during construction

Hydrology and Water Quality

- National Pollutant Discharge Elimination System (NPDES) General Permit for Storm Water Discharges Associated with Construction and Land Disturbance Activities (Construction General Permit), including a construction Storm Water Pollution Prevention Plan (SWPPP) administered by the State Water Resources Control Board

Local Permits/Approvals/Regulations

Air Quality

- Local Air Pollution Control/Air Quality Management District Air Permits to Construct (or Pollution Control District Determination of Compliance/Small Power Plant Exemption for projects licensed by CEC – e.g., centers incorporating onsite generators > 50 MW)
    - District Prevention of Significant Deterioration/New Source Review and permitting could apply to new/modified emitting units (e.g., best available control technology requirements and/or emissions offsets might be required)
- District (and possibly local) requirements for emissions created during construction (e.g., exhaust and dust control measures)

Biological Resources

- Regional Water Quality Control Board (RWQCB) permit for discharge of pollutants/waste discharge requirements during construction, if applicable

Building/Trade Permits

(typically handled by local planning/building/public works department)

- Building/construction permits, grading permits, occupancy permit/certificate
- Electrical/plumbing/mechanical permits
- Seismic safety/soil analysis, if applicable
- Other permits required by local ordinances or planning documents (e.g., fire department permit, tree removal permit, etc.)

Energy

- Local permitting, if sufficiently small or otherwise exempt from CEC jurisdiction
- Utility grid interconnection agreement and approval of interconnection costs with local utility subject to utility tariff and tariff of the Independent System Operator (ISO)
- FERC approval (i.e., market-based rate authority) needed if project involves wholesale transmission of power or wholesale sales of power

Land Use, Planning, and Zoning

(by local planning agency unless otherwise noted)

- Any approvals or amendments to planning documents necessary from the local planning department/agency to ensure consistency with local plans (e.g., General Plan, Master or Specific Plan, if applicable)
- Certified CEQA document (EIR or equivalent or mitigated negative declaration)
    - CEQA process may require adoption of impact mitigation measures during construction
- Zoning Clearance, or Amendment if proposed uses are inconsistent with allowed uses
- Specific Use Permit or Conditional Use Permit, if required
- Site Plan/Design Review (Local planning department/agency or architectural review committee)
- Site Development Permit or other local permit, if required
- Telecommunications permit, if required

- Coastal Development Permit, if necessary (California Coastal Commission (CCC))

Public Utilities and Services

- Water supply assessment, if applicable
- Satisfaction of all local agency and utility requirements for water source, usage and disposal
- Local approvals for new utility connections, such as potable water, fire water, recycled water, sanitary sewer, and storm drain

For some sites, additional agencies might have permitting or review authority (or broader authority than listed above), including U.S. Environmental Protection Agency (EPA), National Marine Fisheries Service (NMFS), U.S. Army Corps of Engineers (USACE), California Department of Water Resources (DWR), California Department of Toxic Substances Control (DTSC) and/or the county environmental health agency, California Coastal Commission (CCC), California Natural Resources Agency, California Public Utilities Commission (CPUC), California State Lands Commission (SLC), California Department of Conservation (DOC), California Department of Forestry and Fire Protection (CAL FIRE), California Department of Parks and Recreation, Federal Transit Administration (FTA), U.S. Department of Transportation (DOT), California Department of Transportation (Caltrans), California Highway Patrol (CHP), California Native American Heritage Commission (NAHC), California Department of Industrial Relations, Division of Occupational Safety and Health (Cal/OSHA), California Department of Conservation Division of Mines and Geology, and, for the Bay Area, the San Francisco Bay Conservation and Development Commission (BCDC).

Potentially applicable local and regional plans and policies vary by jurisdiction, and policies might include requirements for air quality, groundwater sustainability, habitat, community conservation plans and other issues. Local planning, building, design and code enforcement depend on location, as do municipal permits (grading, building, design review, occupancy and use permits, and other requirements, including potential hazardous materials storage, historic preservation, tree removal approvals/arborist report) and other reports (geotechnical, environmental site assessment, engineering flood potential, noise and vibration, fire water, water supply assessment, etc.).

Operational

All permits and approvals for operational activities should be renewed and maintained in good standing, including compliance with all conditions of approval.

Federal Permits/Approvals/Regulations

- To the extent required above, comply with terms and conditions during operation

State Permits/Approvals/Regulations

Air Quality

- Air Resources Board (ARB) Airborne Toxic Control Measure limits might apply (e.g., for emergency standby diesel-fueled engines)

Hazards and Hazardous Materials

- Hazardous Materials Business Plan, if handling or storing threshold amounts of hazardous materials during operation
- Permits for routine transport, use, or disposal of hazardous materials (e.g., diesel fuel)
- Compliance with California Health and Safety Code regulations for storage tanks
- Site Management Plan and/or Health Safety Plan for contaminated soil, if applicable

Hydrology and Water Quality

- NPDES General Industrial Stormwater Permit and SWPPP (State or Regional Water Quality Control Board)

Local Permits/Approvals/Regulations

Air Quality

- Regional Air Quality Management District/Pollution Control District Permit to Operate (typically renewed annually) for regulated sources

Land Use

- CEQA process may require adoption of impact mitigation measures during operation

Other

- ISO27000 and related standards (information security, quality and environmental management)
- Financial management, data transfer, and data broker requirements

Considerations for Colocated Power Plants

Most of the above checklist is equally applicable to a power generating plant that is colocated with a privately owned data center. In addition, there are many more sources of regulatory authority that may have specific application to generation and interconnection facilities. These facilities vary widely in energy source and operational impacts, so this part of the Base Checklist is merely illustrative.

- California Energy Commission (approval of larger generation facilities)
- California Independent System Operator (as to interconnection and grid dispatch)
- California Public Utilities Commission (as to satisfying or demonstrating exemption from certificate of public convenience and necessity or tariff requirements)
- Federal Energy Regulatory Commission (as to interconnection with the interstate grid serving the Western United States and to the extent of any wholesale sales of power to the grid or other customers)
- Other affected federal regulatory bodies, such as Nuclear Regulatory Commission (NRC) for nuclear generation ❖

4

Navigating Data Center Contracts

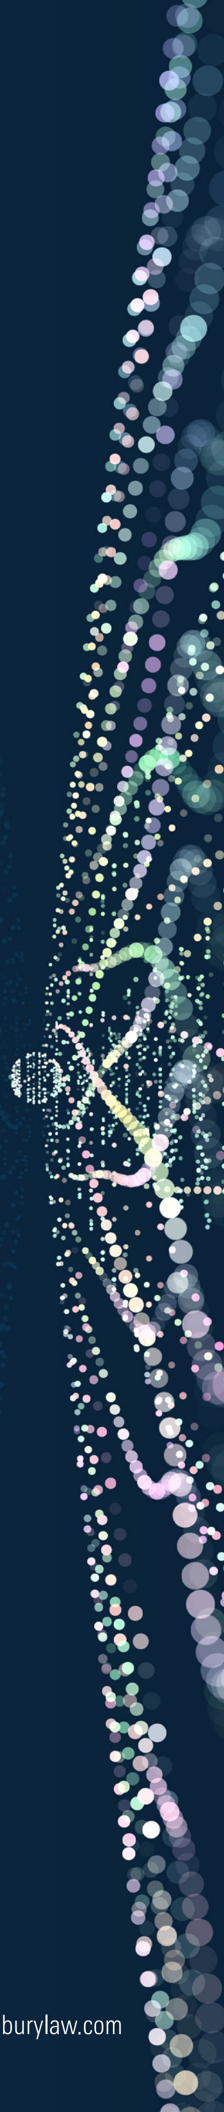

Navigating Data Center Contracts

by Brittney Sandler & Phil Evans

Our colleagues recently provided a comprehensive overview of the **anatomy of a data center**, which explored the structural, energy and real estate implications of these essential facilities.

As more businesses rely on external data centers to house their critical infrastructure, the importance of a well-negotiated data center services agreement becomes paramount. These contracts require more than just finding space—they require careful consideration of the technical, operational and financial details that will support a customer's business needs.

Although terms may vary depending on whether the contract is with a dedicated data center, shared colocation ("colo") data center, and/or managed services provider (MSP) that subcontracts with a data center to provide space and services, the following are key considerations when negotiating any data center deal.

Space and Specifications: Defining the Environment

The space a customer leases within a data center can vary greatly in size and specifications. When contracting, it is important to detail:

- **Size and Power Requirements:** How much physical space and power the customer's equipment will need. Power density (watts per square foot) is critical, and must align with the customer's anticipated growth. Accordingly, parties should specify initial size and power requirements and negotiate a customer's rights of first refusal for adjacent space and additional power.
- **Cooling and Heating:** Climate requirements to ensure that the data center maintains appropriate temperatures for customer equipment. Excessive heat can damage servers, while proper cooling enhances performance. Accordingly, contracts often specify acceptable temperature ranges and require continuous climate monitoring and alerts for any deviation beyond the acceptable range. The parties can also monitor these data center environment attributes as part of a Service Level Agreement (SLA) to incentivize the provider to meet those tolerances.
- **Connectivity and Redundancy:** Terms that specify failover systems and procedures to ensure business continuity if there is a failure (e.g., backup cooling systems, redundant power and network connections, the ability to connect to different carriers or cloud providers). Often, this takes the form of a detailed backup and recovery plan that meets the customer's requirements and is tested regularly. Thoughtful SLAs that guarantee response and restoration times also help minimize the risk of downtime if there's an incident.

Physical and Data Security: Protecting Your Assets

Security in a data center is twofold: protecting physical assets and ensuring data integrity. The contract should address both, including:

- **Access Controls:** The parties should negotiate access protocols, including who can physically enter the customer's space and how security will be enforced and monitored. For instance, is 24/7 biometric access control available, and how long are video recordings retained? Is there an authorized list of entrants? Who is responsible for escorting approved visitors to a customer's cage (and if it is someone other than the customer, is there a separate charge for this service)? These are some of the many details that should be defined.
- **Data Security:** If a third-party provider has access to a customer's data or equipment, the parties need to ensure that data protection regulations (like the EU's General Data Protection Regulation (GDPR) or the California Consumer Privacy Act (CCPA)) are met, and each party's

obligations and rights are clearly defined. The contract should also include clauses covering encryption and other security requirements, along with each party's obligations with respect to changes in regulations and compliance audits.

Service Delivery: Managing Expectations

Particularly where a customer is leasing space in a colo facility or working with an MSP, the parties need to clearly define what equipment, office/storage space, and services will be delivered, such as:

- **Equipment:** The contract should specify any equipment (e.g., racks, servers, power units) that will be provided or made available for rent by the data center provider or MSP. Where this is the case, the agreement will also need to clearly outline (i) the provider's maintenance and replacement obligations, (ii) who is responsible for the equipment, including liability for damages, and (iii) each party's ownership status during and after the term of the agreement. For example, does the customer have the right or obligation to purchase the equipment at the end of term?
- **Smart Hands Services:** If the customer does not have onsite personnel to manage its equipment, it may need "smart hands" services, whereby the provider's technicians handle troubleshooting and maintenance within the customer's space. The parties will need to specify the hours of operation, access protocols and costs associated with such services.
- **Additional Space:** If an MSP provides smart hands services for a customer, the parties will need to determine whether the MSP will have office space in the data center to work from when the personnel are not on the data center floor. In a colo data center, a customer may also need to rent additional storage space for spare parts and equipment. In addition to accounting for these additional space needs, the contract must specify the corresponding costs and use restrictions for such spaces.
- **Service Level Agreements (SLAs):** To incentivize adequate service delivery, it is crucial to include in SLAs appropriate provider obligations for things like uptime, response times and repair timelines, along with meaningful credit amounts due to the customer if the provider fails to meet the service level commitments.

Pricing Structures: Breaking Down the Costs

Pricing in data center deals often involves multiple layers and can be complex. Understanding the market standards can help avoid surprises.

- **Space and Power:** These are typically charged based on square footage and the amount of energy consumed (often in kilowatts). Some data centers offer tiered pricing depending on power density needs.
- **Energy Costs:** The parties must specify how energy consumption is calculated and whether costs are fixed or variable based on market rates.
- **Services:** Managed services, whether from the colo provider or an MSP, come at additional cost. These include smart hands, monitoring, backups and managed IT services.
- **Flexibility and Scalability:** The contract should specify whether the customer can expand or reduce space or power requirements, along with the associated costs. Including these pre-negotiated terms and rates can help avoid disputes as customer needs evolve.

Other Key Issues to Consider in Contract Negotiations

Finally, the parties will need to address these common issues during contract negotiations:

- **Termination/Extension Rights and Exit Strategy:** The parties should define their options, rights and obligations (and associated costs) if the customer needs to (i) exit the data center early, (ii) extend the term, or (iii) receive assistance transferring to another location/provider ("disengagement assistance"). This is typically done through appropriate termination, extension and disengagement assistance terms. What's "appropriate" depends on the specifics of each deal; however, because exiting a data center can be a long and complex process, the parties should ensure adequate time to exit, ideally with rights to extend for predetermined periods (e.g., three-month blocks) at predetermined costs. This provides certainty for all parties if a customer experiences unexpected (and often, inevitable) data center exit delays. Having the right to a clean exit strategy, with rights to remove equipment and ensure continuous service elsewhere, will help a customer negotiate all other aspects of the services relationship.

- **Energy Efficiency and Sustainability:** With increasing focus on sustainability, the parties may want (or may be required by their regulators or investors) to negotiate terms related to energy efficiency, renewable energy use and the provider's carbon footprint. In some cases, monthly reporting of carbon footprint sustainability may be required, including whether there is a zero-carbon footprint or carbon offset purchases.
- **Liability and Insurance:** The parties will need to define liability in the event of equipment damage, service failure, security incidents and other events. The parties should also define and understand the limits of the provider's insurance coverage and whether additional coverage is required.
- **Regulatory Compliance:** Certain regulated customers may be obligated by their regulators to only use data center facilities that are certified according to applicable standards (e.g., NIST 800-53, ISO/IEC 27001, etc.), in which case the contract must include such flow-down requirements.

Conclusion: Tailoring Your Deal to Your Needs

A successful data center contract requires more than just securing space. It requires aligning the physical, operational and financial aspects of the arrangement to meet a customer's specific needs. Addressing these and other key considerations in data center contracts will help protect each party's interests and ensure a smooth operation. ❖

5

Powering Data Centers with Nuclear Generation

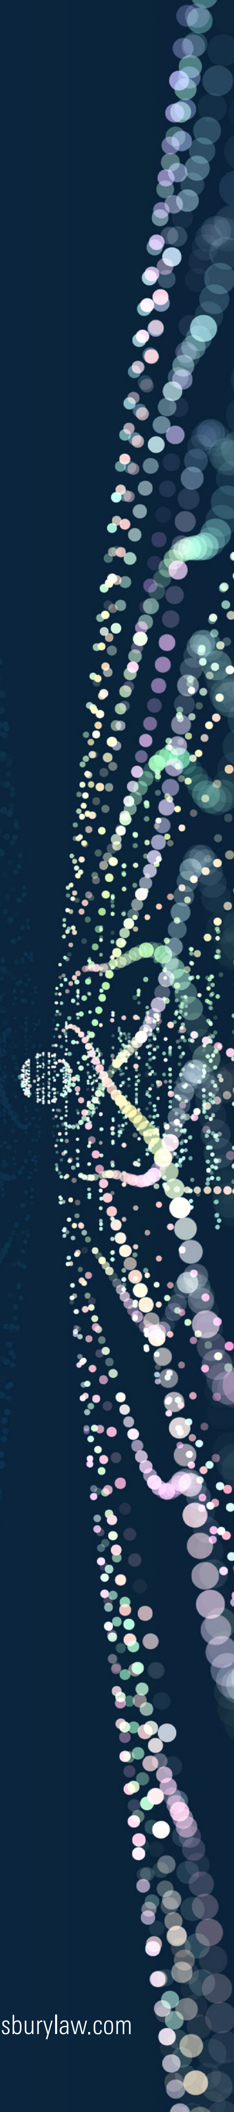

Powering Data Centers with Nuclear Generation

by **Robert A. James, Elina Teplinsky, Jeffrey S. Merrifield & M.C. Hammond**

The rapid growth of electricity demand from data centers has emerged as a major challenge for the U.S. power sector. Much of this demand is being driven by the deployment of large learning models (LLMs) and generative artificial intelligence (AI). These workloads require large-volume, high-uptime computational infrastructure, and correspondingly large, reliable power supplies.

Combined with broader electrification across transportation, buildings and industry, this surge is pushing generation planning as well as grid capacity toward and beyond their limits across many national and international jurisdictions. For utilities and grid operators—many of which have faced decades of ho-hum flat or predictable load growth—this new environment demands rapid, large-scale planning. Build too little, and the risk of blackouts increases. Build too much or in the wrong places or sequence, and ratepayers are stuck funding stranded assets.

Data centers are approaching and exceeding gigawatt-magnitude loads and demand exceptional reliability—up to 99.999% online service—due to the computational and financial costs of power interruptions, and many data center developers are committed to identifying carbon-free generation for their facilities. Traditional generation sources have their limitations. Natural gas-fired projects have been delayed by the current long lead times to procure turbines, and carbon capture is presently an expensive pathway to zero-emission power. Variable renewable energy (VRE) sources like solar and wind, paired with lithium-ion battery storage, have attractive near-term economics but face longer-term impediments for data center applications when taking into account the useful life of panels, turbines and especially batteries, and the large amount of land required for renewable deployment at this enormous scale. **Advanced geothermal** is promising for some quantity of generation or heat content focused in certain geographies and geologies.

The challenge with any mix of existing generation sources is that the demand for power occasioned by data centers is dwarfing the capabilities of gas, renewables/storage and geothermal for baseline operation. “It’s crazy because […] the [entire] state of Oregon is about 6 gigawatts (GW),” **says one Amazon manager**, “and you have these large hyperscalers [asking,] ‘Can I get 6 GW too?’”

Enter Reliable and Carbon-Free Nuclear Energy

Nuclear power in the United States has long served as a cornerstone of the country’s low-carbon energy mix. As of 2025, there are 94 commercial nuclear reactors operating across 28 states, collectively providing nearly 20% of the nation’s electricity and almost half of its carbon-free generation. These reactors offer unparalleled reliability, with average capacity factors exceeding 94%, far outperforming most other generation sources. Increasingly, the nuclear fleet is being viewed not only as clean energy infrastructure but as a strategic asset for powering AI workloads and data centers.

In this context, several shuttered nuclear plants are being actively pursued for restart, reversing what had been a trend of premature closures due to market pressures. Most notably, the Palisades Nuclear Plant in Michigan—shut down in 2022 and acquired by Holtec International—is now on track to become the first U.S. commercial nuclear plant to restart after closure, with support from the U.S. Department of Energy’s Loan Programs Office (LPO). The plant could return to service as early as 2025.

Another notable example is the 2024 power purchase agreement between Microsoft and Constellation Energy, structured to support the restart of the Three Mile Island Unit 1 (TMI-1) reactor in Pennsylvania. (This is not the TMI-2 unit that was the subject of the 1979 incident.) The deal is among the first of its kind to link a corporate clean energy buyer directly to the revitalization of a retired nuclear asset and reflects growing interest in pairing existing nuclear infrastructure with long-term digital energy demand.

These restart efforts are not sufficient to meet the scale and speed of demand driven by the explosive growth of AI. Even if all viable U.S. reactors currently under consideration for restart were successfully brought back online, their combined capacity would represent only a fraction of projected demand growth—let alone offer the geographical diversity or scalability needed to support the decentralized digital infrastructure emerging across the country. Similarly, many existing nuclear units will be seeking power uprates—potentially adding several thousand megawatts to the grid—but this too will not meet AI-based demand. To close this gap, companies are pursuing a new wave of nuclear construction focused not only on traditional large-scale plants but also on Small Modular Reactors (SMRs). To understand the attraction of SMRs for this application, it is useful first to step back and understand what these reactors are and how they are distinctive in two dimensions—different from prior nuclear builds and suited for the data-center context.

SMRs have the capability to have significant portions of the plants built in a factory environment, where economies of scale can be captured and weather-related impacts are avoided. While today's large nuclear power plants can involve 30-35% fabricated content, many SMRs are **expected** to have 50-60% of such content. The modules are also smaller and designed to be shipped either entirely within a standard rail or truck container or in a modest number of containers for ready assembly on site. The more modules built in a given factory with the same workforce, the better the cost and schedule expectations.

SMRs are generally considered to be reactors between 50 and 300 megawatts of electricity output or MWe, though there are examples, such as the Rolls Royce 440 MWe design, that are larger. Designs smaller than 50 MWe are considered to be microreactors. The range of designs encompasses (a) compact versions of the "Gen III+" widely deployed technology based on light water cooling (either pressurized water reactors (PWR) or boiling water reactors (BWR)), as well as (b) new "Gen IV" designs using non-water cooling and moderating materials (including molten salt, sodium and lead, among others) at higher temperatures for greater efficiency.

Here is the **Pillsbury Guide to Advanced Reactor Designs**, so you can see a subset of the dozens of modular designs being evolved in the regulatory and commercial domains.

SMRs are especially compelling for the data-center use case. They offer:

- Baseload reliability critical for high-throughput AI workloads;
- A small physical footprint compared to land-intensive VRE;
- Flexibility for siting near data centers to minimize latency and interconnection issues. This is because the U.S. Nuclear Regulatory Commission, which normally requires an Emergency Planning Zone (EPZ) of at least 10 miles, supports SMRs having a smaller EPZ that extends only to the plant's site boundary (as little as 1,000 yards from the plant);
- Support for "behind-the-meter" installations that allow developers to bypass increasingly lengthy grid interconnection timelines. (**FERC proceedings are currently under way** that will help determine the contributions colocated facilities will make to share in the grid costs.)

New Nuclear Economic and Regulatory Outlook

Can the promise of new reactor deployments be realized to meet the energy demand occasioned by the data center and other drivers for electrification?

According to a recent study by the Idaho National Laboratory, which facilitates nuclear research, powering a 300 MWe data center entirely with VRE and batteries could cost more than a similarly sized SMR over time. That study recognizes that to address issues of intermittency, wind or solar generation would have to be overbuilt to meet the demand need. Additionally, due to the degradation curves of lithium batteries, their limited useful lives will require more capital investment during the lifespan of a typical data center facility. Admittedly, the study also assumes that SMRs can be built on budget and in about 4.5 years, as is currently envisioned.

Other studies hedge their bets on the possible cost, schedule and efficiency performance of SMRs, given the early stage in their development. Nonetheless, there are certainly conceptual attractions of having a **large amount of zero-emission power** generated with **high uptime** on a **small footprint**. An SMR should be able to achieve 57,000 MWh/acre/year at only 38 GW/acre, better than either gas or VRE on the one hand or large nuclear on the other.

Each SMR design outlined on the chart above is pursuing its own distinctive path to regulatory approvals and commercialization in the United States and beyond. Kairos Power obtained a construction permit from the NRC in December 2023 and completed the first safety-related concrete pour for its Hermes reactor earlier this month. And in April 2025, Ontario Power Generation (OPG) in Canada obtained a construction permit from the Canadian Nuclear Safety Commission (CNSC) for the construction of the first GE-Hitachi BWRX-300 reactor on the site of its Darlington nuclear power plant. TerraPower started site preparation activities last year for its commercial scale reactor in Wyoming. NuScale has now achieved NRC standard design approval for its larger 77 MWe design. Other designs are well along in the regulatory approval and pre-application processes. The various hyperscalers are aligning with several of the modular reactor developers, placing bets in most cases on more than one horse.

The **Pillsbury Nuclear-Powered Data Center Project Tracker** is **available here**. (Fusion generation for data centers deserves, and will receive, its own Pillsbury article.) Pillsbury is active on multiple nuclear projects—both large reactors and SMRs. Please contact the authors for more information and assistance.

De-Risking First-Mover Investment

The on-again, off-again history of new nuclear construction in the U.S. has led to little stability in construction workforces and an inability to validate cost estimates. This has made sponsors and lenders wary of the economic viability of nuclear energy—whether large reactors or SMRs. Any lender is likely to require that project funding plans include some form of large and readily available financing reserves (e.g., cash, letters of credit, or funding availability) to cover unplanned costs. This contingency adds meaningfully to the overall capital commitment. Even with such project cost buffers, there remains some probability that costs will exceed committed financing.

To combat this risk, the Energy Futures Finance Forum—a program within the EFI Foundation—recently published a policy framework for a publicly funded cost stabilization facility (CSF) to address the risk of potential cost overruns to sponsors for early-stage projects, with the intent to mitigate a key hurdle to new nuclear energy projects.

Together, **EFI and Pillsbury developed a model term sheet** for such a CSF. Under this structure, a guaranteed loan would be drawn to support potential cost overruns of at least three SMR projects using the same underlying technology. The CSF could be backed by either a private lender or a public one, such as the Department of Energy's LPO.

The risk of cost overruns is greatest for the earliest projects, and later projects will benefit from identifying the issues encountered in prior projects and how to avoid or mitigate them. Thus, the sponsors should agree to share the repayment of the CSF, again on an equitable basis. In recognition that the first projects may face the greatest challenges, the model term sheet does not specifically allocate the CSF, so the first project could potentially absorb all of the capacity, unless the sponsors choose to specify an allocation.

The model term sheet is by its nature only a starting point, and eventual agreements for a CSF of this type may differ in various ways. However, this model can provide the required conceptual approach needed in order to share the risks of cost overruns across multiple projects and over an extended payback period.

Uncertainty Between Congress and the White House on Nuclear Financing

The Trump administration and Energy Secretary Chris Wright have continued to champion new nuclear development. Indeed, on May 23, 2025, President Trump **signed four executive orders** to speed up commercialization of nuclear power in the U.S., with a goal to quadruple nuclear power capacity by 2050.

The enthusiasm shown by the executive branch needs to be aligned with legislative proposals. The House budget bill proposed shortening the eligibility window for key federal incentives. Under the proposed change, construction would need to begin before the end of 2028 to qualify for Production Tax Credits (PTCs) or Investment Tax Credits (ITCs). This accelerated timeline could disqualify new nuclear projects without a current order book.

The House also proposed rescinding all unobligated credit subsidy funding for the DOE LPO. If enacted, this would require borrowers—particularly those developing first-of-a-kind projects—to bear the full cost of the loan risk, making LPO-backed financing less affordable and accessible. This is at odds with Secretary Wright's indication that LPO monies should be directed toward nuclear projects.

If these proposals were to become law, nuclear project development will rely even more heavily on anchor customers like the hyperscalers. These companies may need to provide upfront equity, development funding, or long-term power purchase agreements (PPAs) to catalyze early deployment.

The Senate is actively considering modifications to the budget bill. Senators Dave McCormick (R-PA) and Chris Coons (D-DE) recently introduced the International Nuclear Energy Financing Act to encourage more nuclear energy financing for projects that would create jobs in the United States. Senator John Barrasso (R-Wyo.) is the Senate Majority Whip and has historically supported nuclear development.

For tech companies aiming to scale AI data centers, nuclear offers firm, clean, scalable power. But making this opportunity a reality will take creative approaches to project finance, including new forms of risk-sharing.

Pillsbury's Energy Transition Group is actively engaging with clients on innovative funding structures for new nuclear. From engaging with the DOE and national laboratories to forming development consortia, we are supporting sponsors, developers, utilities and tech firms navigating this next chapter in energy infrastructure. ❖

Pillsbury Guide to Advanced Reactor Designs

Design	Capacity	Details & Status
Pressurized Water Reactor (PWR)		
Nuscale VOYGR	77 MWe	NRC certified 50 MWe design in 2022 and issued Standard Design Approval for uprated 77 MWe US460 module design in May 2025. Established supply chains with Doosan in Korea, and ability to use commercial fuel.
Hadron Energy Halo MMR	10 MWe	Ongoing pre-application engagements with the NRC. Regulatory Engagement Plan and amendments submitted in May and July 2025. Publicly observed pre-application meeting on licesing pathways with NRC in July 2025. Revised Quality Assurance Program Description (QAPD) Topical Report accepted for NRC review in December 2025.
Holtec SMR-300	300 MWe	$116 million risk reduction award through DOE ARDP program. Two units to be deployed at Palisades site in Michigan. Pre-application activities with the NRC.
Westinghouse AP300	300 MWe	Smaller version of AP1000. Pre-application Regulatory Engagement Plan submitted to NRC in May 2023 and updated in February 2024. NRC held public meetings on licensing strategy in December 2024.
Rolls Royce RR SMR	470 MWe	Competitively selected by Great British Energy – Nuclear in June 2025 as preferred SMR technology. Ongoing UK Generic Design Assessment since 2022. Wylfa confirmed in November 2025 as first planned deployment site.
Deep Fission Borehole Reactor 1 (DFBR-1)	15 MWe	Regulatory Engagement Plan submitted to NRC in March 2024 and updated in January 2025. Ongoing pre-application activities with the NRC. Selected for DOE Reactor Pilot Program in June 2025.
Last Energy PWR-20	20 MWe	In February 2025, announced plans to build 30 microreactors in Haskell County, Texas as well as nonbinding commercial agreements to deploy 80 units in UK and EU markets. NRC pre-application activities ongoing. UK regulators completed Preliminary Design Review in June 2025. Selected for DOE Reactor Pilot Program in June 2025.
Boiling Water Reactor (BWR)		
GE-Vernova Hitachi Nuclear Energy BWRX-300	300 MWe	Ontario Power Generation (OPG) submitted construction license application for Darlington New Nuclear Project in Lake Ontario in April 2025. Canadian Nuclear Safety Commission (CNSC) approved construction in May 2025. TVA submitted USA's first Construction Permit application for the Clinch River site near Oak Ridge, TN in May 2025, which NRC accepted for review in July 2025.
High-Temperature Gas Reactor (HTGR)		
BWXT BANR	15-20 MWe	Contracted with Wyoming Energy Authority in September 2023 to evaluate placement at mining/industrial sites. Signed a collaboration agreement with Tata Chemicals Soda Ash Partners (TCSAP) in October 2023 to explore the potential deployment of the BANR SMR at the TCSAP's Green River site in Wyoming.
NANO Nuclear Energy KRONOS MMR	15 MWe	Working with University of Illinois at Urbana-Champaign (UIUC) to go through NRC licensing approvals since May 2025. NRC issued final Safety Evaluation approving reactor's Fuel Qualification Methedology Topical Report in April 2025. Won Direct-to-Phase II SBIR contract from AFWERX in September 2025. Geotethnical drilling at UIUC host began in October 2025. Planned Construciton Permit application submission around Q1 2026.
Radiant Industries Kaleidos	1.2 MWe	Selected by DOE in April 2025 to receive HALEU for first reactor test. Completed prototype testing at INL in November 2024. Selected for DOE Reactor Pilot Program in June 2025.
Valar Atomics Ward 250	100 kWt	Began testing at Utah San Rafael Energy Lab (ISREL) in September 2025. Experimental NOVA core achieved cold criticality at DOE Los Alamos Naitonal Laboratory (LANL) in November 2025. Selected for DOE Reactor Pilot Program in June 2025.
X-Energy XE-100	80 MWe	Selected for DOE ARDP in October 2020 – Two units at Dow Texas site. Four units at Amazon/ EnergyNorthwest with directed $500 million investment from Amazon in October 2023. DOE allocated initial HALEU quantities in 2023 and 2024.
Molten Salt Reactor (MSR)		
Natura Resources MSR-1	Non-power research 100 MWe commercial	NRC completed environmental review and issued construction permit for initial prototype research reactor deployment at Abilene Christian University in September 2024. Selected for DOE Reactor Pilot Program in June 2025.
Terrestrial Energy Integral Molten Salt Reactor (IMSR)	195 MWe	Undergoing NRC pre-application & CNSC review. Completed Phase 2 of Canadian Nuclear Safety Commission Vendor Design Review in April 2023. Selected for DOE Reactor Pilot Program in June 2025.
Kairos Power KP-FHR (Fluoride-Cooled High Temperature Reactor)	150 MWe (2x75 MWe plants)	Selected for up to $303 million risk reduction award by DOE. NRC construction permit for Hermes demonstration reactor issued in December 2023. Signed Master Plant Development Agreement with Google in October 2024 targeting deployment of ~500 MW by 2035. DOE allocated conditional HALEU commitments in April 2025 in first round of HALEU Availability Program.
Liquid Metal-Cooled Reactor (LMR)		
TerraPower Natrium Sodium Fast Reactor	345 MWe 500 MWe (5.5 hrs)	Selected for DOE ARDP demonstration award in October 2020. Non-Nuclear construction permit for coal-to-nuclear site in Kemmerer, Wyoming granted by Wyoming Industrial Siting Council in January 2025. Construction Permit application submitted to NRC accepted for docketing in May 2024. DOE allocated conditional HALEU commitments in April 2025 in first round of HALEU Availaiblity Program.
ARC Clean Technology ARC-100	100 MWe	Selected in July 2020 by NB Power, the primary electric utility in the Canadian province of New Brunswick, for deployment at their Point Lepreau Nuclear Generating Station.
Oklo Aurora Sodium Fast Reactor	1.5 MWe 15 MWe	Submitted Combined License application to NRC for 1.5 MWe design in March 2020. NRC denied aplication without prejudice in January 2022. Selected for DOE Reactor Pilot Program in June 2025. Announcd in May 2023 plans to deploy of two 15 MWe units on Southern Ohio Diversivication Initiative (SODI) lands by ~2028. Selected for DOE Reactor Pilot Program in June 2025.
High-Temperature Sodium Heat Pipe		
Westinghouse eVinci Microreactor	5 MWe	Selected for $7.4 million DOE ARDP risk reduction cost share awared in October 2020. DOE allocated conditional HALEU commitments in April 2025 in first round of HALEU Availability Program.

Pillsbury Nuclear-Powered Data Center Tracker

May 2026 Edition

Commercial Structure	Project Details	Offtaker	Anticipated MW in Deal	Expected Online Year	Reactor Designer/ Type
Long-Term PPAs	Google and NextEra signed a 25-year PPA to purchase carbon-free nuclear energy from NextEra Energy's Duane Arnold plant in Iowa. Duane Arnold, which was shut down in 2020, is expected to be online and delivering electricity onto the grid by the first quarter of 2029, pending regulatory approvals.	Google	615 MW	1Q2029 (Plant Restart-Duane Arnold)	GE BWR-4
	NextEra Energy and Google have also signed an agreement to explore the development of new nuclear generation to be deployed in the U.S.				
	Plant restart + 20-year PPA at approximately $100/MWh between Microsoft and Constellation Energy Generation's Crane Clean Energy Center (formerly Three Mile Island Unit 1).	Microsoft	~830 MW	Mid-2028 (Plant Restart - Three-Mile Island)	Babcock & Wilcox/ Pressurized Water Reactor (PWR)
	Additionally financed by $1 billion loan guarantee from DOE Loans Program Office.				
Long-Term PPAs	On January 9, 2026, Meta announced a deal with Vistra for support 20-year PPAs covering more than 2.6 GW of nuclear capacity from its Perry and Davis-Besse plants in Ohio and the Beaver Valley plant in Pennsylvania, including 433 MW of uprates—the largest nuclear uprates ever backed by a corporate customer—while providing the financial certainty needed to pursue subsequent license renewals and extend plant operations for decades.	Meta	2.6 GW + 433 MW of uprates	Operational/ Existing Plant (Perry & Davis-Besse)	GE BWR-6 (Perry) Babcock & Wilcox/ Pressurized Water Reactor (PWR) (Davis-Besse)
	Idaho Falls Power signed MOU on September 16, 2024, for a PPA which would provide the "right to eventually purchase energy" from Aalo Atomics. Under the PPA, Aalo "would lease land for the life of the project, up to 80 years," at the site, where the public power company also plans to complete construction of a 17.5-MW natural gas-fired peaking plant by the end of 2025.	Idaho Falls Power	75 MW (7 Units at 10 MWe)	After 2030	Aalo Atomics/ Aalo-1
	17-year PPA for $18 billion between Talen Energy and Amazon from Talen's Susquehanna nuclear plant.	Amazon	1,920 MW from the 2.5 GW Susquehanna plant.	Operational/ Existing plant (Susquehanna)	General Electric/ BWR
	At full contract quantity, Talen will provide Amazon with 1,920 MW through 2042, with options to further extend duration. Power delivery schedule will ramp over time, expecting to achieve the full volume no later than 2032.				
	Meta secured a 20-year agreement with Constellation Energy to extend the life of the Clinton Clean Energy Center in Illinois, beginning in 2027. The deal includes a 30 MW plant uprate and potential development of advanced reactors at the site.	Meta	Plant life extension + 30 MW uprate of facility.	2027	General Electric/ BWR
Direct Investment + Future PPAs	Amazon provided $500 million of direct investment into X-Energy through its Climate Pledge fund.	Amazon	n/a	n/a	X-Energy/ Xe-100
	Amazon to pay for the development of SMRs at Dominion North Anna Power Station.	Amazon	320 MW	Early 2030s	X-Energy/ Xe-100
	Amazon to fund initial feasibility phase for EnergyNorthwest to develop and deploy X-Energy SMR technology in Washington State; Amazon to enter into PPA for output of SMR.	Amazon	320 MW with option to scale up to 960 MW	Early 2030s	X-Energy/ Xe-100
Prepayment Against Future PPA	Datacenter colocation firm Equinix has made a $25 million prepayment to Oklo.	Equinix	Up to 500 MW	TBD	Oklo/Aurora Powerhouse
	Equinix stated that the prepayment is "expected to be used toward the purchase of power through future PPAs, anticipated when Oklo begins operating their powerhouse fast fission plants" though it can also be used as an option for Oklo stock. A Letter of Intent (LOI) provides Equinix the right of first refusal for 36 months for the output of certain reactors for power capacity of no less than 100 MWe and up to a cumulative maximum of 500 MWe. Future PPA rates to be determined.				
	Meta provided a funding commitment to Oklo that "provides a mechanism for Meta to prepay for power and provide funding to advance project certainty for Oklo's Aurora powerhouse deployment" in Piketon, Ohio. Oklo plans to use the funds to "secure nuclear fuel and advance Phase 1 of the project—supporting the development of clean, reliable power in Pike County that can scale up to 1.2 GW." The Meta-backed prepayments will secure fuel, initiate site work, and enable phased deployment beginning as early as 2030 directly into the PJM market.	Meta	1.2 GW	2030	Oklo/Aurora Powerhouse
	Meta agreement with TerraPower is for "funding" to support the development of two new Natrium units—a combined 690 MWe of firm power—whose delivery is targeted as early as 2032. But as notably, the Meta–TerraPower agreement will also provide Meta "with rights for energy from up to six other Natrium units capable of producing 2.1 GW and targeted for delivery by 2035."	Meta	690 MW with an option for an additional 2.1 GW	2032–2035	Terrapower/ Natrium
Microreactor Preorder Agreement	Datacenter colocation firm Equinix has preordered 20 of Radiant's Kaleidos microreactors units.	Equinix	24 MW/ 1.2 MW/ unit	TBD	Radiant/Kaleidos
Master Plant Development Agreement + PPA	Tennessee Valley Authority (TVA) PPA between Kairos Power and TVA for delivery of 50 MW to the TVA grid that powers Google data centers in Tennessee and Alabama.	Google	50 MW	2030	Kairos Power/ Hermes 2 Plant
	Part of a Master Power Plant Development deal between Google and Kairos Power for up to 500 MW of reactors, designed to create an orderbook and certainty for investment and future deployment.				
Partnership with a Nuclear Development Company	ENTRA1 will develop up to 6 GW on TVA existing nuclear power sites. ENTRA1 will oversee the deployment, financing, investment, development, execution, and management of ENTRA1 Energy Plants with NuScale's SMRs inside. This likely means that TVA and ENTRA1 Energy will now jointly identify sites, coordinate development activities, and share planning efforts for the nuclear projects as ENTRA1 handles the financing and construction obligations.	Various on TVA Grid	Up to 6 GW, through module packs of 77 MWe per unit.	TBD	NuScale/ VOYGR
	Google has partnered with Elementl Power, an advanced nuclear development company, to identify, acquire and prepare locations to accommodate nuclear reactors.	Google	2.4 GW (at least 3 sites will deploy 600 MW each)	TBD	TBD
	Elementl Power will work with Oak Ridge National Laboratories (ORNL) OR-SAGE siting tool to identify 3 ideal SMR plant sites and Google will provide early-stage capital.				

6

Power Purchase and Interconnection Agreements for Data Centers

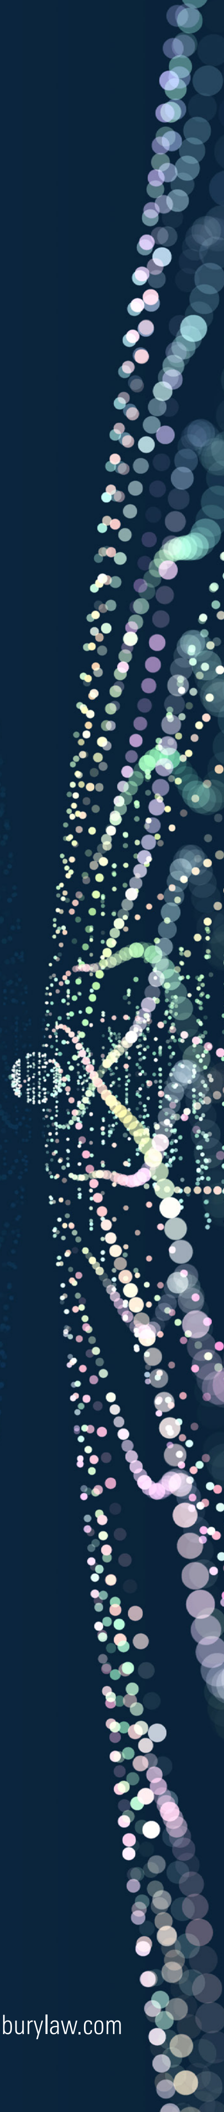

Power Purchase and Interconnection Agreements for Data Centers

***by* Alicia M. McKnight & Rob James**

When structuring these agreements, it is important to consider how siting decisions, load characteristics and regulatory constraints shape risk allocation and commercial strategy.

Data centers—particularly those supporting generative artificial intelligence (AI)—are rapidly emerging as one of the most significant sources of electricity demand globally. In response, power procurement and grid access have become mission critical considerations for developers and tenants alike. As these facilities seek scalable, uninterrupted and cost-effective power, the negotiation of power purchase agreements (PPAs) and interconnection agreements plays a central role in determining both operational reliability and long-term economic viability. This article outlines key considerations for structuring these agreements, with a focus on how siting decisions, load characteristics and regulatory constraints shape risk allocation and commercial strategy.

Background

Data centers are facilities that house large numbers of computer systems for the purpose of storing and allowing access to large quantities of data. With the rapid growth of generative AI, the development of data centers and their concomitant power demand is expected to dramatically increase. In 2014, data center demand was approximately 1.5% of electricity consumed in the United States; that has since more than doubled, and today it accounts for approximately 4.4% of global demand. In the United States alone, current forecasts predict that this demand will double or even triple over the next decade.

These data centers require a significant and reliable power supply in order to function efficiently to ensure uninterrupted service to customers. That power can either come from the grid via utility service or through negotiated PPAs with electric generating resources. Companies are looking at siting these facilities in markets where the power is relatively cheap and interconnection queues are relatively short. Other siting considerations include mitigating environmental risks in particular locations (e.g., avoiding areas that are prone to

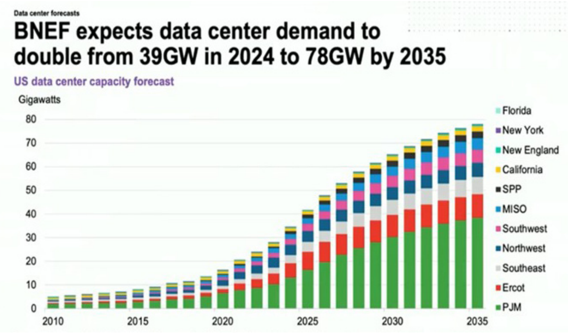

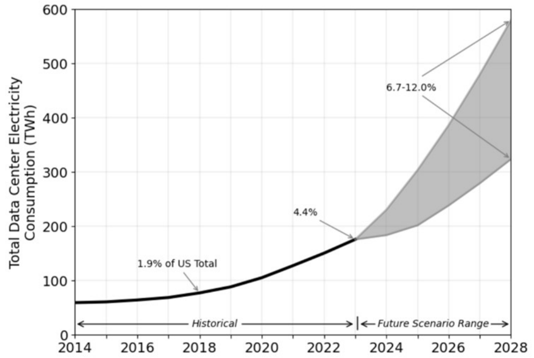

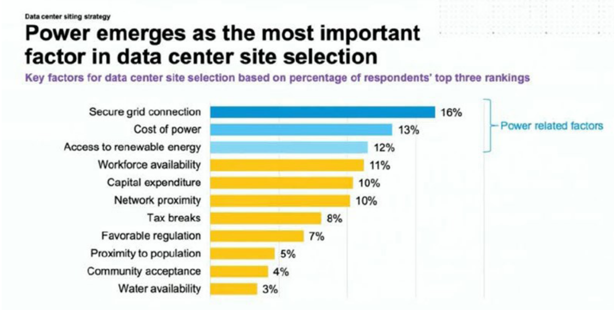

Source: BloombergNEF, DC Byte

natural disasters such as earthquakes, hurricanes or wildfires), choosing a temperate climate (given that cooling accounts for a significant portion of data center energy use), and managing land costs. Siting will likely impact the considerations applicable to the PPA and Interconnection Agreement.

In addition to siting, the purpose of the center may also impact the parties' strategic drivers with respect to the PPA and Interconnection Agreement. Generative AI data centers must rely more on having continuous power with no room for interruption, while crypto miners may not be subject to such rigid demands. Because of this, generative AI data centers are more likely to require strong interconnection and be located in areas that can provide this, whereas crypto miners may have more flexibility on where they can locate. Crypto miners can be more flexible when it comes to negotiating PPAs, while generative AI data centers will have more specific needs and must ensure that they can receive long-term PPAs that can provide them with a reliable and sustainable power supply.

Power Purchase Agreements

PPAs are long-term contracts establishing the terms and conditions under which electricity generated by a power plant will be sold. In the data center space, the PPA can be either (i) between the power generator and the data center developer, who then provides the energy as a service to the tenant as part of the data center lease, or (ii) directly between the generator and the data center tenant. The former is more common for a multitenant or colocation data center, and the latter is more common where there is a single or substantial anchor tenant in a data center, such as a hyperscaler (i.e., a company that builds and operates massive-scale computing infrastructure). In fact, many of the large-scale hyperscalers prefer to procure the energy for their data centers directly. For example, in May 2024, **Microsoft entered into an agreement with Brookfield** to deliver over 10.5 gigawatts of new renewable energy capacity between 2026 and 2030 to Microsoft in the United States and Europe. Similarly, **Meta recently signed a deal with developer Invenergy** to support four new projects which will secure 791 megawatts of renewable energy.

There are several different flavors of PPAs. First, there is the distinction between physical PPAs and virtual PPAs (or VPPAs). Physical PPAs allow buyers to take title to the actual electricity produced by the generating facility through physical delivery via the grid and, for renewable projects, acquire the renewable-energy certificates produced by the generating facility. In contrast, VPPAs do not result in the buyer taking physical delivery of electricity, but rather (i) the buyer procures their energy from their local utility company, (ii) the buyer pays a fixed price to a power generator for power that is produced by the generating facility and delivered into the wholesale market, and if the market price for the delivery of power at the specified grid delivery location is higher than the fixed PPA price, then the generator pays the difference, and vice versa, and (iii) in the case of renewable-energy projects, the buyer acquires the renewable-energy certificates produced by the generating facility (and to further complicate things, some VPPAs allow the buyer to elect a physical delivery option). A more detailed discussion of VPPAs can be **found here**.

Because there is no actual physical delivery of energy for VPPAs, there is more geographic flexibility for buyers, which can be beneficial to companies with multiple locations. A physical PPA may not be a viable option for a distributed load or load in a regulated market. Additionally, VPPAs provide a partial hedge against wholesale price volatility. However, because of this, VPPAs also pose a market price risk for buyers since they involve the exchange of a fixed price for the unpredictable wholesale market price, which may result in buyers owing significant amounts of money to the generator even if their own electricity costs remain stable.

On the other hand, physical PPAs allow buyers to receive power directly from the generating facility, which can reduce the risk involved as a specific price or price formula is agreed upon upfront. However, location is more important with physical PPAs as the buyer must be located in the same grid region as the generating facility. Physical PPAs also require coordination with grid operators for delivery and can entail significant interconnection costs. Finally, physical PPAs may not be available in regulated electricity markets or may need to be in the form of a sleeved PPA where the utility acts as an intermediary between the generator and the data center or hyperscaler.

In each of these different types of PPAs, there are some consistent considerations and risks that must be allocated.

- **Term of the PPA.** PPAs can span 10 to 20 years to provide long-term price stability and support the financing for the construction of the power generating facility. This duration may also drive more competitive PPA pricing for the buyer. However, some data center operators or hyperscalers may prefer a shorter five-to-ten-year term to allow them to adapt to evolving technologies or market conditions. There is an inherent tension as hyperscalers may need to deploy rapidly in order to take advantage of technological innovations, while power generators face development, construction, ownership and operation processes that frequently span decades.

- **Counterparty Credit Risk and Credit Support.** The power generator will likely be financing the construction of the generating asset on the long-term revenue streams under the PPA, and the project finance providers will be acutely attuned to the credit worthiness of the buyer. For that reason, many power generators may prefer a PPA directly with the data center tenant, if that tenant is one of the major hyperscalers. In contrast, if the data center tenants are smaller companies or startups and the data center developer has a stronger credit profile, then the power generator may prefer a PPA with the data center developer. In contrast, the buyer will be concerned with ensuring that the power generator (which is typically a special purpose entity) provides sufficient credit support in the form of a parent company guaranty or letter of credit to secure its long-term obligations under the PPA.

- **Development Milestones.** If both the data center and the power generating asset are greenfield developments, it will be critical to ensure that the key development milestones and guaranteed commercial operation date for both the data center and the power generating asset are aligned, and there are liquidated damages for any delay in either the data center of the power generating asset that are adequate to cover anticipated losses if they become misaligned. Those liquidated damages should dovetail with the remedies available to either party under its engineering, procurement and construction contracts.

- **Scheduled and Unscheduled Outages; Curtailment.** The PPA will need to address scheduled and unscheduled outages of the power generating facility, including curtailment directed by a grid operator or by the buyer. The data center may want the ability to procure (or cause the power generator to procure) replacement or backup power during any scheduled or unscheduled outages or grid-directed curtailment. Conversely, the power generator may want the ability to operate the power-generating facility and to sell that power either into a wholesale market or to other buyers if the data center is not capable of taking that power.

Power Supply and Interconnection Agreements for Data Centers

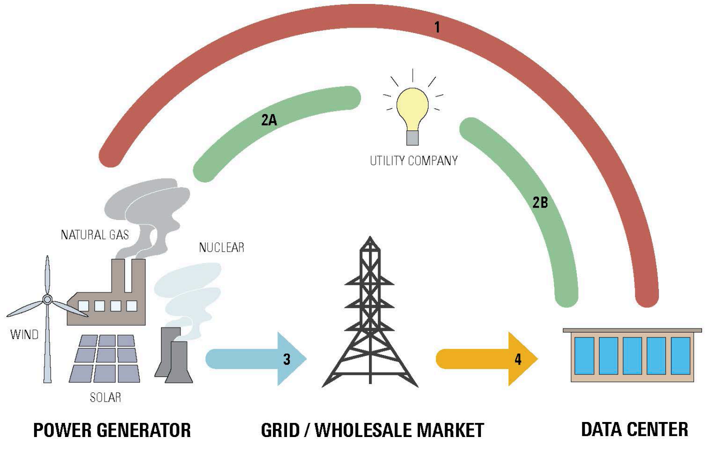

① Power Purchase Agreement (Physical or Virtual) with Data Center operator or tenant
(2A) Power Purchase Agreement with utility and (2B) Electricity service at utility tariff rates
③ Interconnection Agreement (Generation)
④ Interconnection Agreement (Load)

- **Shape and Volume Risks.** These risks relate to a potential mismatch between the energy generation profile of the power generating asset and the data center consumption pattern, which are most acute in VPPAs for intermittent clean energy generation projects. Shape refers to the differences between the hourly generation profile of a renewable project and the hourly consumption profile of the data center—for example, a wind project typically has the strongest generation at night, but a data center may have its load peaking during the day. Volume risk refers to the risk that the overall production of the generating project over a specified period (frequently an annual basis) will be either greater or less than overall electricity usage by the data center during that period. In recent years, two different types of hedge instruments have been developed to help mitigate the financial consequences of these risks. First, a volume firming agreement can address shape risk, resulting in the combined volume of power on which the buyer settles under the VPPA and under the volume firming agreement being roughly the same amount of power used by the data center. Second, a settlement guarantee agreement to lock in a long-term-fixed cost with respect to the VPPA (either as a fixed $ per MWh or fixed $ per settlement period). For a data center, mitigating the financial risks are only half of the practical solution; the data center will want to have backup generation or the ability to acquire power from the grid to ensure continuity of the operation of the data center, and then these financial instruments can offset the costs of that backup generation.

- **Assignment and Change of Control.** The conditions for transferring the PPA to a third party or upstream change of control of either the power generator or the customer will be specified, as well as any obligations to replace credit support and any qualifications of the applicable transferee.

Interconnection Agreement

Interconnection refers to the system structures that allow electricity generators to feed power into the grid and for electric loads to draw power from the grid in a synchronized and reliable fashion. When considering issues of interconnection, it is imperative to consider both the interconnection of the power generating asset and the interconnection of the data center. These involve different approval processes and most independent system operators have separate interconnection queues, with associated timelines. Typically, both types of interconnection requests will be initiated by the submission of an interconnection application or request to the relevant grid operator. Upon submission of the initial interconnection application or request, together with the payment of any associated fees, the applicant will be placed in the queue. The grid operator will then conduct an early-stage feasibility or screening study to determine if the proposed generating asset or data center load can interconnect at the proposed location. If the feasibility or screening study is positive, generally the next step is to conduct one (or more) fulsome study (or studies) to determine the costs associated with the relevant interconnection—including both the interconnection facilities themselves, as well as any system impacts and related network or deliverability upgrades.

Most grid operators require either a deposit or posting of some type of financial security in a relatively modest amount to commence the full interconnection studies. The final step is to enter into an interconnection agreement with the independent-system operator, participating transmission owner or local distribution company. The timeline for the interconnection process varies based on the applicable grid, but it often ranges from 18 to 48 months.

Frequently, the independent system operator specifies the form of interconnection agreement. However, there are several items that can be negotiable:

- **Interconnection Cost Allocation.** The interconnection agreement will outline the responsibility for each party to construct the necessary interconnection and system upgrade infrastructure and allocate the associated costs among the parties. Frequently, direct interconnection costs are paid by the generating asset or load, whereas system costs or network upgrades may be shared. Frequently, the sharing mechanism is defined by the rules of the independent system operator or tariff of the local utility company.
- **Modification to Facilities.** The rights of the power generating asset or data center to add additional capacity or load will often be specified, as well as the utilities obligation to reserve transmission or distribution capacity. For data center interconnection, minimum demand thresholds, ramp-up periods and penalties for under-utilization will need to be negotiated.
- **Milestones and Scheduling.** The timelines for achieving permitting, equipment procurement, construction testing and commercial operation should be clearly set forth, as well as any associated flexibility and grace periods to address regulatory or construction delays. Force majeure provisions and other events excusing delay or non-performance, as well as associated duration limits and required notices, will also be included. Consideration should be given to delay-damages and termination rights.
- **Ownership of Interconnection Facilities.** Who owns, operates and maintains the interconnection facilities will be set forth in the interconnection agreement, including the timing for the transfer of ownership of the applicable facilities. Ownership typically depends on location and function, and maintenance responsibilities and costs are frequently negotiated.
- **Credit Support.** The credit support requirements, including the amount, form (letter of credit, cash, parent company guaranty) and duration will be set forth in the interconnection agreement. Consideration should be given to reductions in credit support over time, release mechanisms and triggers for drawdown.

- **Curtailment Rights.** Any dispatch, must-run requirements or curtailment rights of the independent system operator will be set forth in the interconnection agreement. Data centers may seek guarantees of firm or non-interruptible service. For both types of interconnection, consideration should be given to whether compensation is owed for non-emergency (i.e., economic) curtailment, duration and settlement methods.
- **Metering and Data Access.** Location, specifications and access rights for metering equipment and operational data will be set forth. Both generators and customers will often want real-time and historical access to meter data to ensure accuracy. Consideration should be given to meter calibration and testing, data ownership and cybersecurity provisions.
- **Assignment and Change of Control.** Similar to a PPA, the conditions for transferring the interconnection agreement to a third party or upstream change of control will be specified, as well as any obligations to replace credit support and any qualifications of the applicable transferee.

CoLocation with Power Generation Facilities

Colocation of a data center with a power generating facility can offer benefits like reduced transmission costs and losses and improved reliability. Many companies are currently exploring behind-the-meter colocation, meaning that the power generating asset can deliver power directly to the data center without being interconnected with the grid; however, even when there is behind-the-meter, colocated power generation and data center load, it is likely that both the power generator and the data center will want independent interconnections to the grid to mitigate the risk of stranded assets. This has raised significant regulatory concerns, including potential impacts on grid reliability and cost-shifting to other consumers. For example, the Federal Energy Regulatory Commission (FERC) has initiated a **show-cause proceeding in PJM** after it rejected an agreement between Amazon Web Services and Talen Energy to expand a data center colocated with a nuclear plant in Pennsylvania. The key issues in that proceeding involve (i) whether FERC has jurisdiction over loads that are fully isolated from the grid; (ii)whether colocated loans should be required to pay for transmission services and other ancillary benefits they receive from the grid; (iii) reliability and resource adequacy concerns related to ensuring grid stability; and (iv) clarity on the rates, terms and conditions applicable under the PJM tariff.

While colocation may still be attractive, both power generators and data center operators and tenants should be fully aware of the complications and engage regulatory counsel early in the process to reduce risk that their behind-the-meter arrangements inadvertently expose them to regulatory scrutiny, additional fees, charges or other penalties from grid operators.

Conclusion

As data centers—particularly those supporting generative AI—emerge as dominant energy consumers, the importance of structuring PPAs and interconnection arrangements cannot be overstated. These contracts directly impact cost, reliability and regulatory risk. Power developers, data center developers and hyperscalers alike must align project timelines, credit structures and flexibility mechanisms to balance infrastructure buildout with market volatility and technological change. Interconnection processes must be navigated early and strategically, particularly where both load and generation assets are greenfield developments. Finally, while colocation offers potential efficiencies, it also presents complex regulatory questions that demand early legal engagement. Tailored, forward-looking agreements are not just a contractual necessity—they are critical for ensuring the long-term viability of the data center itself. ❖

7

State Energy Regulatory Approaches to Powering Data Centers

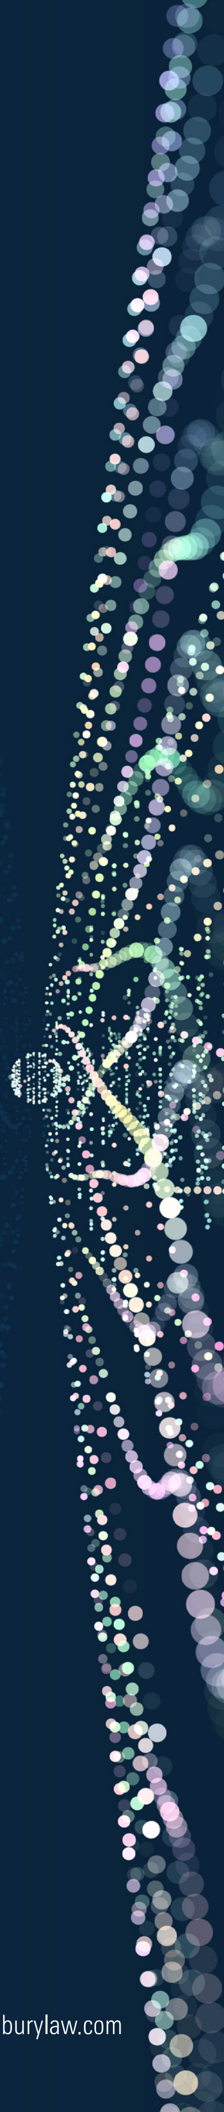

State Energy Regulatory Approaches to Powering Data Centers

***by* Stephen J. Humes, Alicia M. McKnight & Andrew H. Jacobs**

The rapid expansion of data centers—driven by cloud computing, artificial intelligence and hyperscale digital infrastructure—has transformed what were once localized land-use and utility ratemaking concerns into issues of statewide and federal economic and energy policy.

While our **recent commentary** focused on an emerging federal regulatory framework to power data centers, including in pending proceedings before the Federal Energy Regulatory Commission (FERC), states are simultaneously moving to legislate, regulate, incentivize, and in some cases constrain data center development. These state-level actions are ever more consequential to the economics of data center projects, impacting everything from site selection and interconnection timelines to long-term operational risk.

Across the country, legislatures, public utility commissions, governors' offices, electric utility executives and ratepayer advocates are responding to a common set of pressures from the expansion of data centers: unprecedented load growth, concerns about grid reliability, and water and land use impacts. At the same time, policymakers are weighing those pressures against the significant economic growth and tax base expansion that data centers provide. The result is a rapidly evolving energy policy landscape that is redefining how large loads are integrated into grid.

This article analyzes key emerging themes across states rather than cataloging any individual state's statutes or dockets. State policy frameworks are evolving to address the scale and energy impacts of data center growth, even while ongoing considerations **occur** at the federal level to redefine the boundaries of federal versus state authority over powering data centers. As a result, until FERC intervenes and successfully defends its pending new rules in federal courts, these state-level frameworks will materially shape data center projects.

The State Role in Electricity Regulation

Understanding current state authority over and policy responses to data center load growth requires an appreciation of the traditional role states play in electricity regulation. Under the cooperative federalism balance of power set forth in the **Federal Power Act** (FPA), FERC has **authority** over wholesale sales of electricity in interstate markets, interstate transmission, and reliability standards. States retain **jurisdiction** under Section 201(b) of the FPA over retail electricity sales, local distribution facilities and the obligation of utilities to serve customers within their service territories—areas that are central to how data centers obtain power and connect to the grid.

State public utility commissions historically **regulate** retail electric service to end-use customers on a cost-of-service basis, approving tariffs and capital investments of utilities to ensure that rates are just and reasonable while allowing utilities to recover prudently incurred costs. States also **oversee** utility planning and infrastructure development, including the review of integrated resource plans and the approval of transmission and distribution investments that are recovered through retail rates. In addition, states, and often local governments, **exercise** authority over the siting and permitting of electric generating, transmission and distribution infrastructure. Since large load data centers are ultimately end-use customers, states retain primary regulatory authority and responsibility over the requirements for powering such facilities.

Against this backdrop, most state legislation and regulation affecting data centers operate within long-established areas of state authority: retail rates, distribution interconnection, utility planning and infrastructure siting. This authority gives states meaningful influence over powering data centers. As data center load grows in scale and concentration, states are reassessing assumptions embedded in traditional retail rate regulation and planning frameworks.

With respect to on-site or adjacent power generation, the regulatory framework depends on whether or not the generating asset is connected to the interstate transmission system. Increasingly, with encouragement from FERC and given the practical reality that artificial intelligence-scale data centers are experiencing four- to seven-year delays for interconnection to the grid, many operators are pursuing on-site generation (especially natural gas) under "Bring Your Own Generation" (BYOG) models that operate either fully islanded or in hybrid configurations with the grid. Projects that avoid transmission-level interconnection can reduce time to power to roughly three to six years, making BYOG a prevalent hyperscale strategy to power data centers. Such unplugged systems would only be subject to state energy regulatory jurisdiction, but they may present stranded asset risk and greater financing challenges.

Emerging State Policy Themes

State legislatures and public utility commissions are aligning in their responses to rising electricity demand from data centers, though approaches vary across the country. Recent enactments and proposals increasingly reflect a broader shift toward cost causation principles, grid reliability safeguards, and policy-driven constraints tied to clean energy and resource use. Trends among the roughly dozen states that have demonstrated regulatory approaches to data centers show the following themes:

- The states whose regulatory frameworks enable data centers to be paired with on-site or adjacent natural gas generation or behind-the-meter systems include North Dakota, Oklahoma, Texas and Wyoming.
- The states whose regulatory frameworks facilitate renewable powered hybrid models include California, Colorado, New Jersey (with BYOG) and Virginia (increasingly restrictive but favoring BYOG).
- The states with faster pathways for regulatory approvals for BYOG currently include Oklahoma, Texas and Wyoming (for industrial siting).
- The most consumer-protective states are California, New Jersey and Virginia.
- Other states with increasingly strong pushes to protect or insulate ratepayers from costs of powering data centers include Florida, Montana and Pennsylvania.

Prohibition on Cost Shifting

A central theme emerging across states is a renewed emphasis on ensuring that data centers bear the costs they cause the electricity system to incur. This policy objective shows up repeatedly in state legislative activity **addressing** large load interconnection standards and ratepayer protection. States are increasingly focused on preventing "cost shifts," or the allocation of costs to upgrade the grid on residential or small commercial customers when those investments are driven primarily by new large data center load.

Minnesota provides a clear example of a state using legislation to establish guardrails around costs associated with large data center development. In 2025, the state legislature **passed a bill** that included ratepayer protection measures and required enhanced review of utility arrangements for large load. Infrastructure investments required to serve large data center loads—including those to support upgrading the transmission and distribution system—must be directly allocated to those customers rather than socialized across residential or small commercial customers. Similarly, Florida's **proposed framework** would require large load customers to fund system upgrades triggered by their interconnection and to enter into service arrangements that mitigate the risk of stranded utility investment.

Many politicians and ratepayer advocates argue that new large load will not be broadly accretive to the electric grid. On the other hand, increased demand can actually **help lower consumer prices** because the transmission system already requires upgrades, meaning new load highlights an existing infrastructure need rather that creates one. However, many states are moving toward a model in which data center development is structured to avoid cross-subsidization of projects and to protect other customer classes.

Clean Energy and Reliability

A number of states are also focusing on the source and system impacts of the electricity serving data centers. In jurisdictions with aggressive decarbonization goals, policymakers are considering measures that would require large data centers to demonstrate alignment with clean energy targets.

California, for example, has **proposed legislation** that would require new large load facilities to procure or develop clean energy resources sufficient to serve their incremental demand. Alternatively, these facilities may demonstrate that their load will not impede the state's greenhouse gas emissions reduction goals. In contrast, in Texas, the policy approach is less related to clean energy requirements and more a reliability-driven framework for large loads. Texas **passed** a law in 2025 that would expand oversight over large loads and require that large loads be curtailed during electric system emergencies to ensure the stability of the electric grid. These initiatives reflect the view that large load growth may be accompanied by both supply-side or demand-side solutions depending on a state's policy priorities.

Other Environmental Constraints

In parallel with electricity-related reforms, states are increasingly scrutinizing water and environmental impacts of large data centers. In water-constrained regions across the West, particularly in parts of California and Texas, policymakers have **raised** concerns about the cumulative impact of data centers on local water supplies used to cool data centers. Virginia has introduced **legislation** that would require data center companies to submit detailed water consumption estimates as part of their planning and zoning applications. Utah is **pushing** for similar requirements. Resource constraints—especially water—are moving from secondary considerations to material siting and permitting issues for data centers. In some jurisdictions, they are now being directly embedded into state-level policy frameworks governing large data center projects.

Regulatory Agencies as Active Policymakers

Importantly, these trends are not confined to legislative action. State commissions and agencies are exercising existing authority to implement similar principles through tariff revisions, interconnection policies and integrated resource planning processes. For example, Colorado regulators **adopted** guiding principles that highlighted the need to address cost allocation and grid impacts associated with large load customers. Commissions in multiple states have **directed** utilities to separately model data center load growth and assess its reliability impacts. In Virginia, which has the **largest concentration** of data centers in the world with more than 650, regulators have **introduced** specialized large load tariffs or contractual commitments to mitigate risk to existing ratepayers. These tariffs are becoming so common that tariff design for large load is now an active, **tracked area** of policy development rather than a purely utility-by-utility commercial exercise.

The convergence of legislative and regulatory initiatives reflects a broader shift: States are not merely reacting to data center growth but proactively reshaping the rules under which large loads are served.

Strategic Implications for Developers and Investors

Taken together, these trends at the state-level have concrete implications for data center development:

- **Upfront diligence is critical.** Understanding state regulations toward large loads, utility planning processes and political dynamics within a state is as important as ever.
- **Cost allocation decisions are occurring earlier in the project lifecycle.** States are increasingly clarifying cost allocation principles and operational obligations before projects are approved or constructed.
- **Adaptive power strategies can create competitive advantages.** Projects that can adapt to evolving regulatory frameworks may face fewer obstacles. For example, facilities capable of pairing load with on-site or adjacent BYOG, incorporating energy storage, or participating in demand response programs may align more easily with state policies around reliability and cost causation.
- **Tariff structures matter.** Demand charges, curtailment provisions and cost allocation associated with system upgrades may have a greater impact on project economics than the advertised per-kilowatt-hour price.
- **Integrated planning reduces delays.** Coordinating utility upgrades, procuring generation and pursuing grid studies in parallel can help prevent sequential bottlenecks that extend development timelines.

Convergence with Federal Policy and the Risk of Preemption

State-level developments are also unfolding against the backdrop of significant federal regulatory activity. FERC is actively **considering** whether large loads that interconnect at the transmission level should fall under expanded federal jurisdiction, potentially bypassing traditional state oversight. In December 2025, FERC **signaled** a more assertive federal role in this debate by ordering PJM Interconnection, the regional grid operator for 13 states in the Mid-Atlantic region, to create three new transmission services that would facilitate large loads and colocated generation where large loads are physically connected to generators on the same site. That direction was **reinforced** in January 2026 through a Statement of Principles backed by the White House and multiple state governors. PJM was urged to hold an emergency backstop auction and take other measures to support reliability in the region given forecasted demand from data centers.

This federal-state intersection is critical. If FERC ultimately adopts such a framework nationwide, certain aspects of state regulation—particularly those governing interconnection of large loads over 20 megawatts as well as cost allocation—may be preempted. While states would still retain authority over siting, land use and retail rate design, the balance of energy regulation of large loads could materially shift to the federal government, essentially preempting states from more comprehensive energy regulation. Any potential realignment creates legal and regulatory uncertainty, particularly for projects structured around existing interconnection frameworks or cost allocation assumptions.

Conclusion

State regulation of data centers is entering a transitional phase. The era of largely passive accommodation is giving way to a more deliberate policy approach. While states continue to compete for investment, they are placing greater emphasis on protecting ratepayers from shifted costs, maintaining grid reliability and limiting resource impacts.

At the same time, the ultimate contours of state authority may depend on how far federal regulators go in redefining jurisdiction over large loads. Until that picture becomes clear, data center developers, operators and investors should expect continued policy development at the state level and should plan accordingly. ❖

8

Incentives and SALT Considerations for Data Centers

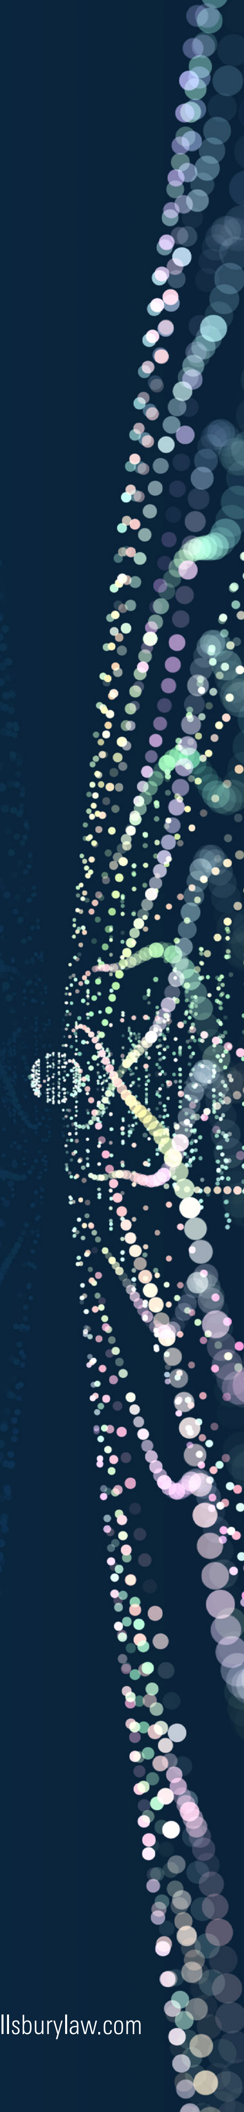

Incentives and SALT Considerations for Data Centers

***by* Aruna Chittiappa, Aimee P. Ghosh, Evan Hamme, Zack Atkins, Jack Camillo, Breanna Zagorski & Phil Wolf**

Global investment in data centers is projected to reach nearly $7 trillion by 2030, with 40% in the United States. As demand for cloud computing, artificial intelligence and digital infrastructure continues to accelerate, competition among some states and localities to attract data center projects has intensified. In response, jurisdictions have increasingly relied on incentive packages—most notably state and local tax incentives—to influence where and how data centers are built and operated.

For developers and operators, understanding incentives for data centers requires more than evaluating headline tax rates. Successful projects depend on navigating a complex landscape of state and local tax (SALT) regimes, statutory exemptions, negotiated abatements and operational requirements that can materially affect both upfront costs and long-term profitability. At the same time, the rapid growth of data center development has prompted increased public scrutiny of incentive programs, particularly with respect to energy consumption, infrastructure strain, and the allocation of costs between operators and ratepayers.

This article provides a practical overview of the incentives and tax considerations most relevant to data center development, with a focus on how tax policy, economic development programs, and operational planning intersect throughout the lifecycle of a project

Incentives and Site Selection Framework for Data Centers

Data centers are uniquely attractive to many state and local governments due to their capital intensity, long asset lives, and reliance on fixed infrastructure. Although data centers typically generate relatively few permanent jobs, their substantial upfront investment and long-term presence often make them candidates for economic development incentives. In particular, tax-based incentives typically form the core of incentive packages for data center projects, reflecting both the capital-heavy nature of these facilities and the limited applicability of traditional job-based incentive models.

Selecting a data center site therefore requires a comprehensive review of the tax environment in a given jurisdiction, alongside a broader assessment of economic development policies and infrastructure capacity. While some incentives are broadly available by statute, many of the most valuable tax benefits—particularly property tax abatements and sales and use tax exemptions—may be structured to attract data center investment at scale. These incentives frequently intersect with state and local tax rules, making early coordination between tax planning and site selection essential.

Incentives for data centers generally fall into three categories:

- Tax-based incentives, which are typically the most economically significant and include property tax abatements, sales and use tax exemptions, and income or franchise tax credits;
- Non-tax economic development incentives, such as infrastructure investment, workforce training programs, and regulatory or permitting efficiencies; and

- Federal incentives, including energy-related tax credits and opportunity zone benefits that may be layered onto state and local programs.

Understanding how these incentives interact—and how they are conditioned on investment thresholds, operational commitments and compliance obligations—is critical during both site selection and ongoing operations.

State and Local Tax Incentives and Considerations

Property Tax Incentives and Considerations

Property tax is a central incentive and tax consideration for choosing a data center site and often represents the single most significant long-term tax exposure for a facility. Because property taxes are imposed annually, negotiated abatements and valuation treatment can materially affect project economics and frequently form the cornerstone of local incentive packages.

Property tax analysis requires careful review of property classification rules, valuation methods and available incentives. The distinction between real and personal property is a fundamental property tax principle, and the resulting classification can significantly affect a company's tax liability. Data center owners (or lessees, if contractually responsible for property taxes) generally prefer to classify as many assets as possible as personal property, which typically results in faster depreciation and lower effective tax rates. Nevertheless, each situation requires careful analysis of the state's property classification rules.

For data centers, states typically treat servers as personal property; however, jurisdictions differ as to whether cabling, power systems and cooling systems are considered personal property or real property. Each asset type requires fact-specific analysis, and classifications vary by state. The determination often depends on the item's attachment, adaptation and intended permanence. To support favorable positions, owners should document installation methods, mobility, and useful lives through engineering or cost-segregation studies.

Assessors often use cost, income and market approaches to value assets. Ensuring that assessors choose the appropriate method is especially important for data centers because misapplications can result in overpayment of taxes. Cost-based approaches, which assessors regularly use for personal property, often overstate value by ignoring rapid depreciation and technological obsolescence—factors especially significant for servers and related assets that may lose value quickly. Owners can potentially reduce assessments by demonstrating functional or economic obsolescence or the right to accelerated depreciation. Maintaining detailed, current asset inventories with acquisition and retirement data can support accurate valuations and successful appeals.

Finally, property tax incentives often drive data center site selection. Many jurisdictions, including states such as Texas and Virginia, offer negotiated property tax abatements or payment in lieu of tax (PILOT) agreements. A PILOT is a contractual arrangement under which a taxpayer agrees to make specified payments in place of otherwise applicable property taxes, typically for a defined period and often as part of a broader economic development initiative. The structure and value of a PILOT can vary significantly depending on property classification, valuation methodology and the degree of assessment certainty provided. As a result, a comprehensive understanding of property tax rules is critical to evaluating the true benefit of a PILOT or abatement program. These incentives are implemented at the state and local level and are subject to jurisdiction-specific approval and compliance requirements.

Sales and Use Taxes

Sales and use taxes are among the most impactful SALT issues for data centers due to the high cost of equipment and infrastructure. Because these are generally destination-based taxes, the jurisdiction where equipment is first delivered or used typically governs taxation. As a result, a data center's most prevalent sales and use tax obligations will be those in the state and locality where the facility is located.

Most states impose sales tax on tangible personal property (TPP) and certain services, while complementary use tax obligations may require self-assessment on out-of-state purchases used locally that were not subject to sales tax at the time of purchase. Combined state and local rates range from about 5% to over 10% depending on the jurisdiction, making location a major cost driver for data center development. Typical taxable items include servers, racks, cabling and environmental controls, but these items can be exempted through either targeted or general exemptions, as discussed below.

Targeted Exemptions

To attract investment, many states offer data center-specific exemptions tied to capital investment, job creation, or the receipt of certification from the state. Project-specific benefits often flow through to the operator's customers, enabling them to claim related sales tax exemptions tied to the facility's certification. While states such as Alabama, Minnesota and Nevada limit the availability of such incentives to a period of 10 to 20 years, certain states, such as Ohio, do not impose time restrictions.

Eligibility for these exemptions is typically conditioned on statutory investment or job-creation requirements. Because certification is granted in advance of project completion, states may have clawback provisions or other corrective measures if a taxpayer ultimately falls short of its commitments. Required capital investment levels vary by state and depend on factors such as whether the facility is located in an economically distressed area, the number of jobs created, and whether the investment involves new construction or expansion of an existing facility.

In addition, several states also impose job creation or wage requirements tied to the number of new jobs and compensation levels relative to local or statewide averages and may require that these jobs be maintained for a specified period. And states may condition incentive eligibility on other operational factors such as minimum facility size, required IT power capacity, or compliance with environmental or sustainability benchmarks.

General Exemptions

Beyond project-specific incentive programs, many states offer general sales and use tax exemptions for construction materials, and equipment used in qualifying data center operations. The most common exemption applies to purchases of TPP used to construct or equip a facility. States like Virginia and Illinois limit the incentive to qualifying TPP (such as hardware, enabling software and infrastructure systems) that support the processing, storage, retrieval or communication of data in a certified facility. Arizona extends its exemption to equipment and component parts used to outfit or upgrade a data center.

Some states go further by exempting not only core infrastructure but also key operational inputs such as electricity or other energy sources. Texas, Washington and Minnesota exempt power consumed in qualifying facilities, whereas states like Wisconsin generally treat electricity as taxable TPP unless used in a data center—creating significant upfront cost variance across jurisdictions. Given the wide variation in state definitions, scope and conditions, a careful review of each cost element and the applicable exemptions is essential.

Corporate Income and Franchise Tax Considerations

Although income and franchise taxes typically represent a smaller component of data center incentive packages than property or sales tax benefits, they remain an important consideration for long-term operating costs and structuring decisions. States vary widely in whether they impose a corporate income tax and in how they apply combined reporting regimes. States like South Dakota, Wyoming, Nevada, Texas, Washington and Ohio do not impose a traditional corporate income tax (though some have gross receipts or margin-based taxes). Among states with income taxes, rates vary significantly—from 2.5% in North Carolina to nearly 11% in New Jersey.

Income and franchise taxes can also encompass affiliated entities. Data center operators evaluating potential facility locations must also consider how a state's "combined reporting" regime could affect overall tax exposure. In mandatory combined reporting states, entities that comprise a unitary business are required to file a single combined return that aggregates their income and apportionment factors. This approach can significantly influence where and how much income is taxed. By contrast, several states do not require combined reporting and instead use separate-company reporting, which generally limits apportionment to the income and activities of the in-state entity.

Apportionment

Data center operators must also evaluate each state's apportionment, sourcing and reporting rules to understand how income will be taxed. Historically, states have apportioned multistate business income using a three-factor formula based on property, payroll and sales. Locating facilities in states that still include payroll or property in their apportionment formulas can increase the share of income apportioned to—and thus taxed by—that state.

States now, however, often apportion income using a single sales factor formula. To calculate the sales factor, states typically attribute sales to the state using either a "cost-of-performance" or a "market-based" sourcing methodology. States applying a cost-of-performance methodology, such as Virginia, source receipts to where the income-producing activity is performed, while states using market-based sourcing, such as New York and Texas, generally assign receipts from digital products and services to the customer's location. Because data centers generally incur their costs of performance at the data center site, selecting a site in market-based jurisdictions may help moderate income tax in the state where the data center is located.

Tax Incentives

Many states provide tax credits for investing within the state, which may apply broadly to a variety of investments or specifically to data center projects. Although no federal income tax credit applies specifically to data centers, certain qualifying projects may benefit from energy, clean electricity, or energy-efficient building incentives under Internal Revenue Code (I.R.C.) sections 48, 48E and 179D. Many states conform to these federal incentives and may also offer job creation or investment credits for qualifying projects.

And federal opportunity zone benefits under I.R.C. section 1400Z-2 are an important consideration. Although a creature of federal tax law, states work closely with the federal government to determine where opportunity zones will be located and how the federal benefit interacts with any state-level benefits.

Operational Considerations

Once operational, data centers must maintain compliance with evolving SALT obligations, which can vary depending on structure (hyperscale, colocation, or managed services) and geographic footprint.

Sales and Use Taxes

A data center's own sales and use tax obligations are only one part of the sales and use tax equation. Operators must collect and remit these taxes from their customers, requiring a close understanding of sales and use tax rules in jurisdictions where customers are located, not merely the state in which the data center is located.

The Expanding Reach of Sales and Use Taxes

The type of data center directly affects nexus and therefore makes a meaningful difference in sales and use tax obligations. For example, in colocation arrangements, the colocator typically establishes nexus simply by owning and operating a facility in the state. States such as Texas, Virginia, New York and California treat the lease or rental of TPP to a lessee or renter in the state as sufficient to establish nexus. In some jurisdictions, including Washington, nexus may also be established by maintaining on-site personnel—whether employees or independent contractors—to provide customers with installation or maintenance services at a third-party's facility. By contrast, providing purely virtual or managed services (e.g., cloud hosting, data storage, network monitoring, or other remote IT management functions) to customers located in another state, without owning or leasing hardware or other equipment located in the state and without in-state employees or independent contractors, generally does not create physical presence nexus.

But even if the customer does not have physical presence at the data center site, the operator still needs to review its customers' locations for nexus purposes. The U.S. Supreme Court's 2018 decision in South Dakota v. Wayfair fundamentally changed how states determine when a business has sufficient connection to establish nexus with a state. The Court eliminated the long-standing physical presence standard for sales and use taxes and greatly expanded states' ability to tax remote and digital businesses. In Wayfair's aftermath, states rapidly adopted sales tax economic nexus thresholds, typically tied to receipts derived from in-state sales (often as low as $100,000) or transaction volume (100 or 200 sales). Although some states may measure these thresholds based only on taxable sales, data center operators can easily exceed them through digital or subscription-based transactions with customers in the taxing jurisdiction.

Tax Classification and Sourcing

Once nexus is established, the critical question becomes what is being taxed, an area where state approaches diverge significantly. In addition to taxing the sale or lease of TPP, a number of states tax digital, cloud-based transactions. States differ sharply in classifying digital and cloud-based transactions, using labels such as "digital services," "data processing," "information services," "computer services," or "digital products." For example, New York, Texas and Washington impose sales tax on various digital, data-processing and cloud-based services. Virginia, by contrast, exempts these types of services if provided electronically, offering a cost advantage for sourcing such sales there. This patchwork creates significant variation in tax exposure and compliance obligations across jurisdictions.

Taxpayers must also determine where sales tax applies (i.e., sourcing). Sales of tangible property, such as servers, racks, or cooling equipment, are generally sourced to the location where the property is delivered and first used—typically the data center. By contrast, digital products and cloud-based services, are often sourced to the customer's location, reflecting where the customer receives the benefit of the service rather than where the hardware resides.

Finally, even minor structuring differences can materially affect tax exposure such as bundling fees for taxable and nontaxable sales, which can create unexpected liabilities. For example, when installation, monitoring, or managed service fees (typically nontaxable) are combined ("bundled") with taxable software or hardware fees, the entire charge may be treated as taxable unless the fees are separately stated.

Corporate Income and Franchise Taxes

As with sales and use taxes, operators may also establish economic nexus with other states for corporate income taxes based solely on sales to customers in a state. The proliferation of market-based sourcing requires data center operators to carefully evaluate how each state sources service receipts for apportionment purposes.

For example, colocation centers may deliver services to colocators only at the site location, whereas hyperscalers and managed services may be delivered where the customer uses or benefits from the service. Complicating the analysis, some states apply "look-through" sourcing rules, requiring taxpayers to consider the location of their customer's customer, or to consider the ultimate "end user," when determining the market. Thus, a data center's operations require a broad review of not only the income or franchise tax rules at the data center site but also other states where significant customers may use or benefit from the data center services.

Non-Tax Economic Development Incentives

In addition to tax-based incentives—which typically represent the most economically significant component of data center incentive packages—states and localities also employ a range of non-tax economic development strategies that can materially affect the feasibility, timing and long-term viability of data center projects. Although these measures generally play a secondary role to tax benefits, they remain critical considerations in site selection and project planning because they directly influence infrastructure availability, operating costs and regulatory certainty.

Infrastructure Investment

Data centers place significant demands on electric power, water, fiber connectivity and transportation infrastructure. As such to attract data center investment, states and localities may invest in or expand core infrastructure as part of broader economic development strategies. Such efforts may include upgrades to transmission and generation capacity, expansion of substations, increased water and wastewater capacity, road and site access improvements, and expanded broadband or fiber deployment. The availability, timing and reliability of this infrastructure can directly affect a data center's ability to come online as scheduled and to expand capacity over time.

Workforce and Training Programs

States may offer workforce development programs that support data center operations, including programming to upskill a regional workforce and partnerships with educational institutions. While these incentives may be modest in scale, they may contribute to operational efficiency and workforce stability. Developers should evaluate how workforce programs align with anticipated staffing needs and whether ongoing eligibility or reporting requirements apply.

Zoning, Land Use and Permitting Efficiencies

Regulatory flexibility can materially affect both project timelines and development risk. Some jurisdictions offer expedited zoning approvals, streamlined permitting processes, or data center–specific zoning classifications intended to reduce uncertainty and accelerate development. These measures can be particularly important for large or phased projects where delays in entitlements can have significant cost implications.

Practical Strategies for Managing SALT Exposure

Property Tax Considerations

- Record classification rationale and installation details.
- Ensure assessors apply appropriate methods and account for obsolescence.
- Evaluate property tax abatements and PILOT agreements.

Sales and Use Tax Considerations

- Monitor both physical and economic presence across states.
- Evaluate taxability of all inputs, services and sales on a state-by-state basis.
- Maintain documentation including exemption certificates and eligibility records.
- Separately state taxable and nontaxable charges on customer invoices.

Corporate Income and Franchise Tax Considerations

- Reassess regularly as customer and equipment locations evolve.
- Consider both statutory rates and apportionment factors.
- Evaluate federal and state incentives in tandem to maximize benefit realization.

Non-Tax Economic Development Incentive Considerations

- Assess energy cost allocation, grid impact fees and infrastructure funding policies early in site selection.
- Track compliance obligations tied to workforce, zoning or permitting benefits.
- Evaluate clawback and performance risks before committing to long-term operational assumptions. ❖

9

Investing in Data Centers

Investing in Data Centers

***by* Adam Weaver, Dave Lillevand & Allan Van Vliet**

It seems like such a simple question. Who owns data centers?

Ownership structures in the digital economy are more varied than might appear on the surface. While the largest computing and cloud service providers, such as Apple, Amazon (AWS), Microsoft (Azure) and Google (Cloud)—also known as "hyperscalers"—do own and operate a significant portion of their global infrastructure, they are increasingly partnering with third-party developers and investors, including real estate investment trusts (REITs), to expand capacity and deploy capital quickly and efficiently. This article provides a guide to how the most prominent strategic and financial players are engaging in this sector.

The Variety of Investment Approaches

Hyperscalers

Hyperscalers at the largest scale often own core campuses in strategic global regions and carry most if not all of the capital cost on their own books. Publicly reported examples include **Microsoft's data center campus** near San Antonio, Texas, Google's facilities in **The Dalles, Oregon**, and **Northern Virginia**, and **Meta's Kuna data centers** in Kuna, Idaho. Most hyperscalers own and operate a dozen or more of their own data center facilities in **locations around North America**. These companies acquire land (often without leveraging debt), self-fund construction costs, hire employees, and manage and operate these data centers to support their own cloud computing business lines; the economics of hyperscale cloud computing services make these major capital investments worthwhile for these large companies.

This model is less common among smaller (non-hyperscaler) operators. The capital requirements involved in acquiring land and in building and energizing the facility are significant, and non-hyperscaler operators may not have the financial capacity or risk tolerance to pursue direct ownership.

REITs, Developers and Infrastructure Funds

Many large-scale data centers and campuses are owned by infrastructure funds, REITs or private equity-backed developers who specialize in developing, owning and operating data center assets. These data center landlords lease space to hyperscalers, enterprises and other cloud service providers. Though hyperscalers do own campuses and facilities of their own (as described above), many still lease data center capacity from data center-specific REITs (like Equinix or Digital Realty), or partner with developers to build on a sale-leaseback model, which provides the developer/landlord with guaranteed rental income to help mitigate the large financial risk of data center development. REITs and developers on the other hand specialize in acquiring and constructing these facilities, and operators are often happy simply to lease the completed facility.

These arrangements often include long-term, triple-net leases that provide steady, predictable cash flow for the facility owners and reduced capital costs for operators. A hyperscaler or other operator leasing a data center will typically pay a base annual rent, and will reimburse the landlord for all (or a pro rata share, if there are multiple tenants or buildings within a campus) of the landlord's operating costs, such as maintenance, security, cleaning, landscaping and other management activities. Leasing rather than owning gives an operator flexibility to more rapidly adapt to market and technological changes, tying up capital for only a defined lease term rather than an extended period of ownership, and in some cases providing **rights to terminate a lease early**. Leasing also helps operators scale up more quickly—when additional capacity is needed, leased space can be brought online much more quickly than newly constructed space.

Other Investors

It is possible for others to invest in data centers, but options depend on financial capacity, risk appetite and access to institutional networks. Some avenues through which interested investors (including individuals) might gain exposure to the data center market are:

- **Public Markets (REITs).** This is the most accessible route. Publicly traded data center REITs like Equinix (EQIX), Digital Realty (DLR) and newer entrants like CoreSite (now part of American Tower) offer direct exposure to the sector. These REITs provide liquidity, dividends and built-in diversification across geographies and tenants.
- **ETFs and Mutual Funds.** While there aren't many mutual funds focused exclusively on data centers, several tech and infrastructure ETFs include heavy allocations to REITs and hyperscaler-linked real estate. Examples include the Pacer Benchmark Data & Infrastructure Real Estate ETF (SRVR).
- **Private Equity and Infrastructure Funds.** Access to private equity or other infrastructure funds that invest in data centers is more limited. These types of funds often require minimum investments in the seven-figure range and are typically open only to institutional investors and high-net-worth individuals. However, some private wealth platforms are beginning to offer limited access to infrastructure-focused funds.
- **Direct Investment or Syndication.** A more niche option is to invest directly in regional operators or through syndication platforms offering co-investment opportunities in smaller to medium-sized data center facilities. These are higher risk but potentially higher return, often suited for investors with both domain knowledge and capital.
- **Investment in Power Suppliers.** One interesting way to participate in the data center phenomenon is to invest in utilities and independent power producers that have a **particular specialty** in projects to supply energy for power and cooling for data centers.

The Data Center Capital Stack: Equity, Debt and Hybrids

Data centers are capital-intensive to build and operate. A hyperscale facility can cost hundreds of millions of dollars, and even smaller regional data centers often require tens of millions in up-front capital. Data center developers may seek financing for these projects from a range of sources, and may mix and match these sources, depending on the type and scale of the facility as well as regional leasing activity and demand.

Private Equity (PE) and Infrastructure Fund Equity

PE and infrastructure investors have poured capital into both developers and operators, funding greenfield builds and platform acquisitions. For example, Microsoft and BlackRock are publicly reported to have teamed up with technology-focused private equity firms to create a Global AI Infrastructure fund which aims to raise $30 billion for the construction and energizing of data centers; xAI and Nvidia have since **joined the fund**. Investment funds of this kind can deploy capital relatively quickly when opportunities arise, and can leverage debt resources on a project-by-project basis. Once a project is completed and operational, the fund (or other investment vehicle) may seek to sell the project at a profit.

REIT Equity and Debt

Data center REITs access both public equity and debt markets to finance new developments or acquisitions. Their access to capital is one reason hyperscalers increasingly partner with REITs. REITs acquire or invest in the land and facility directly, either independently or by entering a joint venture with an operator. REITs may work with developers to construct new projects or may acquire already-completed facilities and lease them out to operators. As with PE and infrastructure investors, REITs can raise funds from a variety of sources and deploy those funds quickly when opportunities arise. Unlike hyperscaler-owned projects, REITs typically have a specific investment period or horizon, after which they will look for exit opportunities rather than continuing to own a project or facility indefinitely.

Corporate and Project-Level Debt

Securing long-term debt to finance large infrastructure projects such as data centers is common, particularly in stabilized assets with secure tenant leases. For example, a developer may negotiate with a hyperscaler to enter an agreement to lease a yet-to-be-completed data center and then provide a lender with the finalized lease to make the lender comfortable that the project will provide sufficient cash flow as a means to secure more favorable loan terms. Because data centers tend to have long occupancy periods and low turnover as compared to other asset classes, cash flows can be easier to predict, and vacancy risks can be minimized. This in turn attracts potential lenders to underwrite these sizeable projects. Alternatively, hyperscalers themselves can use debt to **defray their own capital outlays** for new projects; the lender can rely on cash flow projections of the hyperscaler itself rather than potential lease income to gain comfort lending money in these cases.

Debt terms can range from traditional bank loans to infrastructure debt and even green bonds when sustainability targets are involved.

Hybrid Approaches

The investment options are also ripe for consideration of hybrid approaches. While some states are straining under the power and other requirements of data centers, other localities are looking to lure them in with tax abatements, grants and other incentives. Government users will increasingly demand data center capacity of their own in the 21st century. These alternative approaches to investment may entail public-private partnerships and joint ventures and other forms of collaboration and equity, convertible debt and conventional debt instruments.

The sheer size, diversity and sophistication of the modern data center can test the most experienced of investors and most robust of capital stacks. For those contemplating investment in data centers and the related digital economy, knowing where to start, what to do on the way, and what comes next after you get there requires its own substantial stack of sector-specific legal, business and regulatory guidance. ❖

10

PFAS, HFCs and Related Chemicals in the Data Center Industry

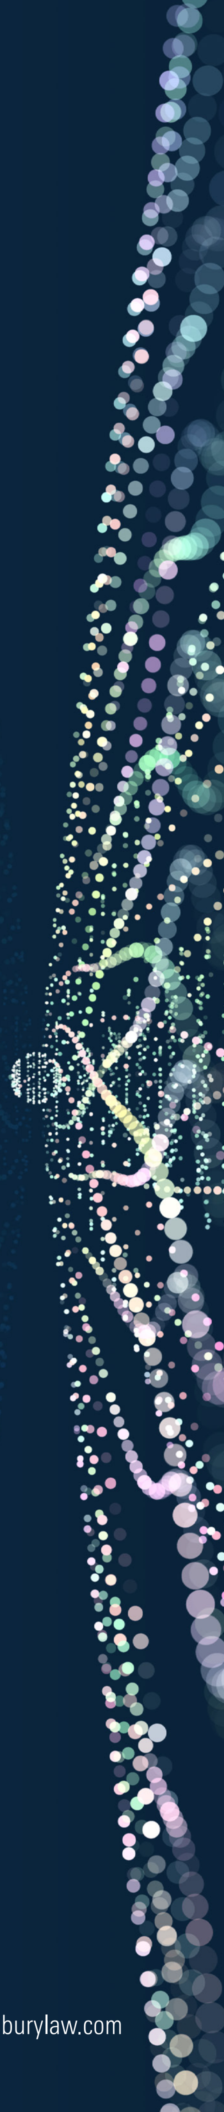

PFAS, HFCs and Related Chemicals in the Data Center Industry

by **Sid Fowler & Reza Zarghamee**

Data centers use various chemicals that have recently been the focus of regulatory efforts at the federal and state level. The historic or future use of these chemicals may create liabilities, obligations, or new costs for both existing and planned data centers.

PFAS Use in Data Centers

Per- and polyfluoroalkyl substances (PFAS) are a class of approximately 15,000 synthetic fluorinated organic compounds (by EPA's **most recent reckoning**) that have been widely used for decades in a wide range of products and industries. The carbon-fluorine bond is among the strongest in organic chemistry and gives PFAS their physical properties: fire, water and grease resistance. These properties—chemical stability, resistance to degradation, and oil and water resistance—make PFAS both useful and concerning from an environmental and health perspective, as they tend to persist in the environment and bioaccumulate in living organisms.

As a result, PFAS have become a major focus of regulatory action and litigation in recent years, with governments and plaintiffs increasingly targeting their use, disposal and presence in consumer products and water supplies. Regulatory definitions of PFAS vary depending on the regulatory body. The most expansive definitions of PFAS are those adopted by **certain states** that classify as PFAS any chemical featuring a carbon atom that is saturated with fluorine.

PFAS that meet this definition serve several necessary functions with respect to data centers. For example, certain server components and cables, including switchgears, either integrate or are coated with PFAS, particularly fluoropolymers. The presence of fluoropolymers in such equipment is hardly unique to the data center industry. Because of their relative stability and molecular size, fluoropolymers are not as bioaccumulative or toxic as long-chain alkyl PFAS (i.e., straight-chain PFAS with molecular backbones that are at least seven carbons long).

On the other hand, more focused applications of PFAS occur with respect to one of the main concerns at data centers—namely, the propensity for servers and other equipment to malfunction as a result of overheating. Because of their low boiling point and fire-resistant qualities, PFAS (as defined per the state definitions) are used as refrigerants and cooling agents within chillers and air conditioning centers, which are an important aspect of most data center designs. In addition to impairing server operability, overheating can give rise to a risk of fire, which is exacerbated by the concentration of electrical equipment and combustible materials like cabling and plastics. As such, data centers are equipped with fire suppression systems, which may contain PFAS.

PFAS in Fire Suppression Systems

Because data centers are packed with server racks, they require fire suppression systems that will not destroy the equipment. The best suited are dry pipe fire sprinkler systems and a clean agent fire suppression system.

- Historically, dry pipe fire suppression systems have used Class B aqueous film-forming foams (AFFF) and automatic reactivation AFFF. These AFFF formulations may contain PFAS, and the activation of fire suppression systems may result in the sudden discharge of large quantities of PFAS-containing AFFF. Furthermore, the long-chain alkyl PFAS present in AFFF—perfluorooctanoic acid (PFOA) and perfluorooctane sulfonic acid (PFOS)—are dispersible in the environment and (a) subject to stringent de facto cleanup standards in the form of maximum contaminant levels of 4 parts per trillion (ppt) in groundwater, orders of magnitude lower than those for other industrial chemicals, and (b) formally designated as hazardous substances under the Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA), meaning that sites contaminated with these chemicals are subject to strict, joint and several liability. Accordingly, some of the largest PFAS settlements to date have occurred in **multidistrict litigation involving the discharge of AFFF to the environment**. To date, 15 states (AK, CA, CO, CT, HI, IL, ME, MD, MN, NH, NJ, NY, RI, VT and WA) have banned the sale of firefighting foam containing PFAS. There is an expanding market of synthetic fluorine-free foams (SFFFs). However, SFFFs do not have the same efficacy as AFFF technologies.

Clean fire suppression systems use gaseous agents to quickly extinguish fires, leaving no residue behind, which makes them ideal for protecting sensitive equipment and areas where water damage would be undesirable. The most commonly used clean agents like FK-5-1-12 (Novec 1230), HFC-227ea (FM-200), HFC-125, HFC-236fa, and 2-bromo-3,3,3-trifluoropropene are PFAS under various definitions, though they have different properties from PFOA and PFOS and are not subject to CERCLA regulation. Even so, data centers may be regarded as a target for litigation in areas where contamination with fluorinated organic compounds has been identified.

HFC-Refrigerants

Hydrofluorocarbons (HFCs) are synthetic chemicals commonly used in commercial refrigeration and cooling systems. Data centers produce large amounts of heat and often need complex cooling systems to ensure their systems do not overheat.

Starting in the 1990s, HFCs came to replace traditional refrigerants, classified as ozone-depleting substances, in the aftermath of the 1987 Montreal Protocol. However, HFCs themselves present environmental risks as greenhouse gases. Indeed, studies exist to suggest that the global warming potential of HFCs may be many times greater than that of carbon dioxide, and the data center industry is coping with an increasingly stringent regulatory framework focused on HFCs. Hydrofluoroolefins (HFOs), which are increasingly being seen as alternatives to HFCs, have a lower global warming potential but may carry their own risks such as flammability.

Congress passed the **American Innovation and Manufacturing ("AIM") Act of 2020**, mandating an 85% phasedown of HFC production and consumption by 2036. Under the Biden administration, the U.S. Environmental Protection Agency (EPA) moved aggressively to implement the AIM Act through a suite of interconnected regulations:

- A cap-and-trade system, which became effective on November 4, 2021, and which is intended to **phase down** over time the total volume of HFCs in the U.S. market.
- The Technology Transition Rule, 88 Fed. Reg. 73098 (October 24, 2023), restricting the use of most HFC-based refrigerants in new equipment.
- The Emissions Reduction and Reclamation (ER&R) Rule, 89 Fed. Reg. 82682 (October 11, 2024), which mandates leak prevention, repair, and in some cases the replacement of HFC-based cooling systems.

This regulatory approach means that both new and existing data centers may be impacted. Existing data centers that rely on HFCs may be impacted by the ER&R rule, as it imposes **extensive and likely costly compliance obligations** on such systems. For example, operators of existing systems may be required to install automatic leak detection equipment, meet detailed recordkeeping requirements, promptly repair leaks that exceed specified thresholds, and even retrofit or retire systems that cannot be repaired within certain timeframes. New data centers may be impacted by the limitations on new cooling technology available, as the Technology Transitions rule removes HFC-based systems from the market. For example, new data centers systems will be legally prohibited from using systems that rely on R-410a, a popular refrigerant in data center cooling systems.

Despite the increasing availability of alternative refrigerants coming onto the market, it may be difficult to retrofit existing refrigeration systems to use these products. Additionally, the new refrigerants—including some that are HFOs and hydrochlorofluoroolefins—may present other risks not customarily associated with HFCs. For example, A2L refrigerants may raise new concerns due to their flammability, CO2-based systems may require higher energy consumption particular in warmer climates, and adiabatic cooling (temperature reduction via pressure drop) may be water-intensive and therefore unsuitable for drier regions. ❖

11

Designing, Constructing and Converting Data Centers and Crypto Mines

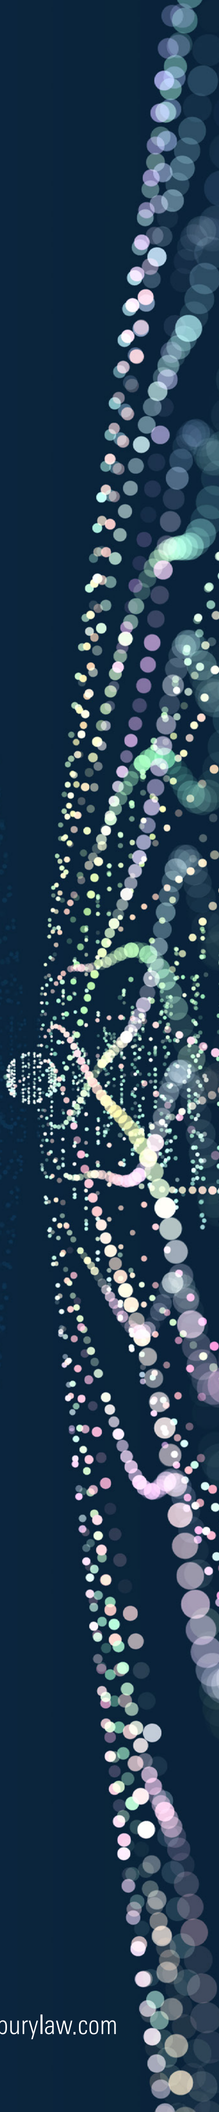

Designing, Constructing and Converting Data Centers and Crypto Mines

***by* Jamie Bobotek, Rob James, Arielle Murphy & Polly Gomez**

The ever-increasing demand for digital infrastructure, coupled with continuing cryptocurrency demand volatility, has generated significant interest in transforming building improvements housing crypto mining farms into modern high-capacity data centers. Crypto mines and data centers share certain foundational elements—facilities with a large footprint and the need for high power capacity and cooling—but they have different functions, and their operational models, technical requirements and regulatory considerations are correspondingly diverse.

Below are some key distinctions between the two, as well as issues developers should anticipate, when designing and constructing either type of facility, or considering a crypto farm-to-data center conversion.

Built for Different Purposes

Crypto mines and data centers share a fundamental requirement—both demand enormous computational power, resulting in the need for massive amounts of power and cooling. Their distinct operational goals, however, result in facilities that look and operate differently. Despite their common infrastructure backbone, their distinct purposes result in differing technical and architectural configurations.

It is useful to review the basics. A crypto mine is a facility that solves complex mathematics problems to secure digital currency transactions, particularly for cryptocurrencies akin to Bitcoin. Those transactions are then applied across a large number of decentralized blockchain networks. In return for this work, the "miners" are rewarded with new cryptocurrency.

The Greek poet Archilochus wrote, "the fox knows many things, but the hedgehog knows one big thing." **Crypto mining facilities** are hedgehogs, optimized for one task and one task only: performing vast numbers of computations, often through **proof-of-work** algorithms, to validate transactions on a blockchain. These operations are typically powered by application-specific integrated circuits (**ASICs**) or high-performance graphics processing units (**GPUs**) running nonstop to maximize **hash rate** and, by extension, potential earnings.

By contrast, **data centers** are foxes. They are facilities, often resembling a complex of warehouses, that hold computers, equipment and other multipurpose digital ecosystems designed to store, manage and disseminate large volumes of digital data. They serve a large number of purposes for many audiences, powering everything from websites, apps and cloud services to online banking platforms and streaming media. While their design and complexity may vary depending on operator functions and support needs, most require strong telecommunication capability to interact with the outside world with no lag, and uber-redundancy to maintain service reliability and keep systems running 24/7.

Key Operational Differences

The most fundamental operational distinction between crypto mines and data centers rests in how they are expected to perform under pressure. Crypto operations are inherently tolerant of downtime—if a machine fails, the network adjusts without significant disruption. This tolerance allows crypto mines to operate without the costly layers of **redundancy** common in enterprise-grade data centers. Crypto mines can take advantage of power price differences by operating at optimal times of day or in response to other market signals.

Data centers, on the other hand, are typically built for near-continuous **uptime**. Any interruption can result in **financial losses**, data loss and corruption, negative customer impact, reputational damage, and regulatory and compliance risks. These facilities rely on multiple backup systems, dual power feeds, failover networking, and strict monitoring to ensure performance and availability.

Network speed is another point of distinction. Crypto mines do not require high-speed bandwidth because they only deal with **small bits of transaction data**. In contrast, data centers are constantly sending and receiving large amounts of information between users and systems, so they require fast, reliable internet through fiber-optic and satellite telecommunications.

Another key difference is how intensively each facility works its hardware. Crypto mines usually run their machines—like ASICs and GPUs—at **full power all day, every day**, without much effort to balance the load. This heavy use can **wear out the hardware faster**. Data centers, on the other hand, take a more **measured approach** to extend the life of their cutting-edge hardware.

Shared Infrastructure and Design

Despite these differences, crypto mines and data centers share a common need for underlying infrastructure. That commonality is a large part of the reason that potential conversions are gaining traction.

Both facility types rely on massive computational power—whether through ASICs, **GPUs** or **CPUs**—to deliver nonstop performance. As a result, both require large-scale, efficient power delivery systems and advanced thermal management to keep hardware from overheating.

In addition, crypto mines and data centers often share similar physical layouts, starting with a large warehouse-type space. These facilities typically include raised floors or dedicated conduits for cabling, reinforced structures to accommodate heavy equipment, and secure access points to protect valuable hardware. Generally, the same building codes and industry standards will apply to the infrastructure, including **ANSI/TIA 942** and **APC No. 92**. These basic shared design elements make conversions feasible at a structural level.

Technical Challenges in Conversions

Because of their shared structural infrastructure, converting a crypto mine into a data center—or vice versa—has become an **appealing potential option to developers** as a cost-efficient reuse of an existing facility. However, developers must be cognizant of the technical challenges involved in such a conversion that may affect their bottom line. Below are several to consider.

Power Redundancy

In general, while crypto mines are typically optimized for maximum computation density with less of an emphasis on redundancy and resiliency, resiliency is perhaps the primary consideration when building and operating data centers. Converting a crypto mining facility to meet data center standards requires significant upgrades to power distribution units and the addition of **uninterruptible power supplies** (UPS) and backup generators to ensure reliability and compliance with data center uptime standards.

Cooling Systems

Crypto mining rigs generally operate under sustained maximum load to produce concentrated heat, which often exceeds the load legacy data center HVAC systems were designed to withstand. Conversion to a data center will likely require upgrades from basic or immersion cooling systems to precision computer room air conditioning/computer room air handling **(CRAC/CRAH) units**. Conversely, standard data center server racks are not suitable to accommodate the form factor or thermal output of ASIC miners or high-performance GPUs. Operators must also account for the different power delivery profiles of mining hardware, which often draw higher amperage at lower voltages than traditional information technology equipment.

Network Infrastructure

In either direction, reconfiguring network infrastructure to fit the requirements of each specific use is another key challenge. Crypto mines rely on minimal network throughput and often lack the robust switching and routing equipment needed for enterprise data workloads. A conversion to a data center will likely require procurement and installation of high-speed fiber, routers and redundant connections. Conversely, existing data center fiber and bandwidth capacity may remain underutilized in a mining operation, leading to a potential stranded investment.

Security Requirements

Security requirements can also shift when converting a crypto mine to a data center. Crypto mines prioritize physical security to prevent theft of valuable mining equipment over cybersecurity controls. Data centers typically require both physical and cybersecurity protocols, such as biometric access controls, 24/7 surveillance, audit logs for compliance, multi-factor authentication, and compliance with frameworks such as **ISO 27001** or **SOC 2**. Data centers run on a colocation or leasing model requiring physical and systems segmentation and segregation of customer assets and the utilities directed to serving those specific assets. A crypto facility may not have been originally designed and constructed with that need for segmentation in mind.

Noise Control

Noise control is another aspect to consider when contemplating a potential conversion, particularly depending on the facility's location. Crypto mines, which often operate in industrial or remote areas, generate substantial noise from ASICs, high-speed fans and dense clusters of mining rigs, with little concern for noise abatement. By contrast, data centers are often planned to operate near commercial or residential zones, especially the edge data centers that are becoming prevalent. Facilities in these locations are typically subject to **stricter noise regulations** and community impact standards. Conversion may necessitate installation of sound dampening infrastructure, such as soundproofing or upgraded HVAC systems, to comply with local ordinances and mitigate operational disruptions.

The bottom line? Converting an existing crypto mine into a data center, or vice versa, may be a viable and effective way to repurpose an existing but no longer fully utilized facility. But it is a complicated process. Developers contemplating these conversions should engage engineering, fire protection, power, acoustic and cooling consultants early in the process to assess feasibility, along with counsel to evaluate contractual, zoning and regulatory implications that may (and likely will) arise from the material change in facility use. ❖

12

AI and Data Centers: A View from the UK and EU

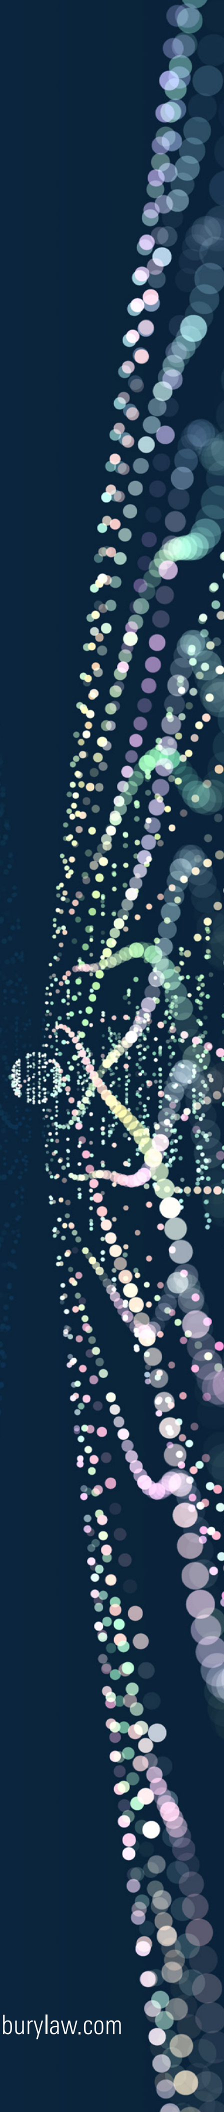

AI and Data Centers: A View from the UK and EU

***by* Lee Rubin, Johanna Lipponen & Scott Morton**

As demand for data processing and storage surges—increasingly driven by AI technologies—the need for robust, scalable data centers is becoming increasingly urgent. Ambitious plans from the UK government and strong private sector investment offer encouraging signs for stakeholders in the nation's data center and AI sectors.

Since the 2024 UK general election, artificial intelligence (AI) has been a top priority on the government's agenda. The King's Speech in July 2024 signaled a cautious approach, focusing on AI regulation and safety, but by October, the government had made sweeping announcements of foreign investment in data centers at the UK's International Investment Summit. The government has since introduced incremental policy developments addressing data center planning and expansion. This momentum continued into 2025, with the release of the government's **AI Opportunities Action Plan** (the Plan), published on January 13, 2025, detailing new initiatives and strategic direction.

The Plan recommends a range of policies and actions for the government to take as part of its overarching aspiration to kickstart broad-based economic growth. The Plan commits to three goals:

- Investing in the foundations of AI
- Pushing hard on cross-economy AI adoption
- Positioning the UK to be an AI maker, not an AI taker

The first of these goals is the most dominant section in the Plan—capturing 30 of the 50 total recommendations. This section focuses on, among other things, building sufficient, secure and sustainable AI infrastructure through recommendations aimed at facilitating the accelerated build out of AI data centers, mitigating the sustainability and security risks of AI infrastructure, and long-term planning for the UK's AI infrastructure needs.

Bold Bureaucracy: AI Growth Zones (AIGZs), AI Energy Council and "Critical" Designations

Building the UK's AI Infrastructure

The government's **response** to the Plan highlights a strong commitment to investment in the foundations of AI through infrastructure support and policy reform. All 50 recommendations set out in the Plan have been adopted by the government with promised delivery timelines.

To drive the government's goal of building sufficient, secure and sustainable data infrastructure, the Department for Science, Innovation and Technology will publish a long-term plan for the UK's AI infrastructure needs and is committed to setting out a 10-year roadmap. While specific details remain under wraps, key focus areas are expected to include (i) addressing security concerns, (ii) sustainability and energy, (iii) supply chain resilience and (iv) the pursuit of sovereign AI compute—public-sector-owned or allocated infrastructure to ensure national capability and strategic independence.

AIGZs: Unlocking Investment

The government has also introduced AIGZs, a public-private initiative to consolidate AI infrastructure development in strategic locations within the UK. This initiative is designed to streamline data center planning, while offering fast-tracked regulatory approvals and priority access to clean energy—incentives aimed at making the UK more attractive to investors. Through this, the government aims to unlock private capital, support technological innovation and drive local economic rejuvenation through job creation and regional growth, ensuring the UK remains globally competitive in the rapidly evolving AI landscape.

In January 2025, the government announced that the first pilot AIGZ would be established in Culham, Oxfordshire, just 1.5 hours from London and home to the UK's Atomic Energy Authority, selected in part for its high-capacity connection to the UK's national grid, available land and proximity to highly skilled human capital. The process for identifying a private-sector partner for the pilot AIGZ will commence in Spring 2025.

On February 10, 2025, the government opened bidding by inviting local authorities and data center developers to submit expressions of interest for the development of additional AIGZs. The formal selection process commenced during Spring 2025 and remains open for submissions. The initial proposals will help to inform the final selection of sites and broader policy decisions.

AI Energy Council: Addressing Energy Constraints

Energy access and grid capacity remain critical bottlenecks for data center expansion in the UK, with operators now actively exploring on-site gas generation and private grid solutions to bypass connection delays, which can exceed 10 years. Currently accounting for approximately 1-2% of the UK's total electricity usage, energy demand for data centers is projected to rapidly expand alongside demand for AI. To tackle this challenge, the government announced in its response to the Plan the formation of an AI Energy Council, co-chaired by the Technology Secretary and the Energy Secretary and formed of industry leaders. The AI Energy Council, which has met twice since its formation, will guide energy policy investment in infrastructure, devising innovative energy solutions with a focus on clean energy.

Energy is one of the greatest areas of skepticism in the government's overall strategy, with the UK having some of the most expensive energy prices in the world, and aspirational net-zero targets, which have been labeled as unrealistic by industry leaders. A report published on February 10, 2025, by the Social Market Foundation has aimed to put pressure on the government to introduce zonal pricing to enable cheaper energy access where there is greater supply—an arrangement which could make Scotland's energy prices the lowest in Europe and spread regional growth and opportunities to remote areas. With the new AI Energy Council available to direct influence towards, industry voices could have more sway on policy going forward.

Data Centers as Critical National Infrastructure (CNI)

In September 2024, the Technology Secretary designated data centers as "Critical National Infrastructure"—the first new CNI classification in almost a decade since the space and defense sectors were added to the CNI register alongside other existing critical sectors, such as transport, energy and health. The designation means that data centers in the UK will receive greater support from the government to monitor, prepare for and respond to threats, such as natural disasters and targeted cyber-attacks. A key focus of the CNI designation is cybersecurity. Among other support, data centers will now enjoy prioritized access to the National Cyber Security Centre—the government agency which bridges industry and government on all things cyber security.

Although the CNI designation does not directly alter the legislative framework, it serves as a strong indicator of the evolving direction of the UK's cybersecurity regulatory landscape. Currently, this area is governed by the UK Network and Information Systems Regulations 2018 (NIS), which implement the EU Directive 2016/1148 (NIS 1). In the EU, Directive 2022/2555 (NIS 2) officially replaced NIS 1 in October 2024, which expanded the types of in-scope entities and enhanced the powers of enforcement and supervision available to authorities in each EU Member State. The UK government is expected to reflect the expanded scope of NIS 2 in the highly anticipated Cyber Security and Resilience Bill which will be introduced in Parliament this year. As with NIS 2, the Cyber Security and Resilience Bill is likely to apply to data center service providers who could become subject to detailed security and reporting requirements. For a detailed discussion on NIS 2, see our earlier **Client Alert**.

The government hopes that improving the resilience of data infrastructure and the industries it supports will not only protect critical services but also attract foreign investment and job creation.

Planning Reform: Addressing Development Hurdles

Due to the scale of modern data centers and their acute infrastructure requirements, obtaining planning permission has presented delays and hurdles to new developments. In July 2024, the government updated its National Planning Policy Framework, which sets out the government's planning policies for England and how the government expects local authorities to apply them. The updated framework introduces explicit support for data center expansion which will pave the way for more accommodating decisions and policy on green-belt development and the identification of suitable locations for data centers.

The government has also confirmed that it will include data centers in its Nationally Significant Infrastructure Projects (NSIP) consenting regime, which streamlines the consenting process for infrastructure projects over certain thresholds.

While the development of planning reforms and policy shifts as they relate to data centers have thus far been quite piecemeal, stakeholders will be paying close attention to the upcoming Planning and Infrastructure Bill (currently at the committee stage in the House of Lords) to see how the government ties these actions together and enables accelerated data center expansion.

Global Investments in AI and Data Centers Presents Competition

The UK government has taken a front seat approach in channeling private sector investment to strengthen its position in the global data center landscape. Between July and October 2024, over £25 billion was committed to new data center projects across the country. The government appears confident that its policy work will encourage even greater investment, announcing over £14 billion in new investment since the launch of the Plan just last month.

The substantial investments in the UK, however, are modest compared to the United States' "Project Stargate," a $500 billion initiative launched in January 2025 by OpenAI, SoftBank, Oracle and MGX to build AI infrastructure. For a discussion focused on progress in the United States, see our recent **Client Alert**. Similarly, ahead of the AI Action Summit in Paris, on February 9, 2025, France announced a €109 billion ($112.6 billion) investment in its AI sector, which includes a €30 to €50 billion commitment from the UAE to finance a 1-gigawatt data center, almost four times the power capacity of the UK's largest operational facility. For our discussion on the AI Action Summit, see our **blog post**.

Regulation in Europe: EU AI Act and the EU Data Act

A key aspect of the EU's approach to AI is a focus on regulation. The EU has advanced regulation significantly in key areas, including the introduction of comprehensive legislation governing AI. Regulation (EU) 2024/1689 (the EU AI Act) establishes a harmonized framework that categorizes AI systems into prohibited, high-risk, limited-risk and minimal-risk, with corresponding regulatory obligations. High-risk AI systems encompass applications critical to the safety of essential infrastructures, such as data centers. For example, AI systems used to monitor water pressure or manage fire alarm controls in cloud-computing centers are classified as high-risk and must comply with stringent regulatory requirements. In contrast, AI systems developed exclusively for cybersecurity purposes may not fall within the high-risk classification. Organizations deploying high-risk AI systems must conduct rigorous risk assessments, adhere to strict safety standards and implement robust data governance measures to ensure the security and reliability of AI in critical sectors. For more insights into the EU AI Act and its implications, see our recent **Client Alert**.

As well as the EU AI Act, Regulation (EU) 2023/2854 (the EU Data Act), the provisions of which mostly come into force from September 12, 2025, introduces new rules aimed at improving fair competition in cloud services by making it easier and more cost effective for customers to switch providers. The EU Data Act will also mean that unfair terms in standard form business-to-business data licenses will not be binding, preventing providers from imposing unfair conditions that could hinder portability or create excessive exit costs. These changes are particularly relevant for data center operators and cloud service providers, which will need to review existing contracts and adapt their infrastructure to comply with the EU Data Act's requirements.

Conclusions: Momentum Meets Uncertainty

Despite the UK's strong push to establish itself as a global AI and data center hub, significant hurdles remain. Energy constraints, planning difficulties and maintaining investment growth in the competitive global arena all pose challenges to scaling AI infrastructure at the pace required. While initiatives such as AIGZs and regulatory reforms provide positive signals, critical details — including the long-term compute strategy — are yet to be announced. Until then, uncertainty persists around how these policies will translate into actionable change for developers, investors and operators. As the UK navigates this crucial period, maintaining confidence in the sector will require continued government engagement, private sector collaboration and decisive action to remove bottlenecks standing in the way of AI-driven growth.

Strategic Considerations for Businesses

As the UK and EU AI and data center landscapes evolve, businesses and investors should consider the following to successfully leverage these opportunities:

AIGZs

- **For data center operators.** Begin evaluating potential AIGZ locations and prepare expressions of interest and applications for government-backed incentives.
- **For investors.** Identify opportunities to collaborate with developers in AIGZs to benefit from fast-tracked approvals and energy incentives.

Energy Supply Constraints and AI Energy Council

- **For operators and hyperscalers.** Conduct energy feasibility studies and engage early with policymakers to secure priority access to clean energy resources.
- **For investors.** Monitor developments in zonal pricing, as this could significantly impact long-term cost structures—consider engaging with policymakers or industry groups to advocate for zonal pricing incentives, particularly for regions where energy abundance could lower operational costs (e.g., Scotland).

Planning Reform

- **For developers.** The UK's inclusion of data centers in the NSIP regime could streamline the approvals process. Operators and investors should begin assessing whether their projects qualify under NSIP and engage with local authorities early to take advantage of fast-track processes.
- **For investors.** The forthcoming Planning and Infrastructure Bill may provide additional incentives or hurdles—consider engaging with policymakers or industry groups to shape outcomes favorable to AI infrastructure development.

Regulatory Landscape and Compliance Preparedness

- **For UK data center operators.** Begin aligning cybersecurity strategies with the anticipated Cyber Security and Resilience Bill, particularly for providers expected to fall within its expanded scope.
- **For EU-based providers.** Assess the impact of NIS 2 and the EU Data Act on data governance and contract terms.

Global Positioning and Competitive Advantage

- **For multinational AI infrastructure investors.** Compare the UK's incentives against U.S. (Project Stargate) and French initiatives (€109 billion investment) to determine where long-term capital deployment is most favorable.

As the UK's policies continue to evolve, organizations must take proactive steps to secure a competitive edge. Whether assessing investment opportunities in AIGZs, mitigating regulatory risks from evolving cybersecurity requirements or evaluating global AI infrastructure strategies, companies should begin preparing now. Pillsbury's Data Centers team is well-equipped to provide tailored guidance, helping businesses navigate these shifts, structure compliant operations and optimize investment strategies in the UK, EU and U.S. markets.

Our team delivers to operators, developers, investors, power suppliers and users the strategic advice required to address all aspects of the data center landscape, leveraging experience with real estate, construction, project development and finance, sourcing, cybersecurity, consumer protection and intellectual property issues directly translated to acquiring, developing, financing, powering and operating data centers, as well as protecting the data and privacy interests associated with them. ❖

13

Satellite Data Centers

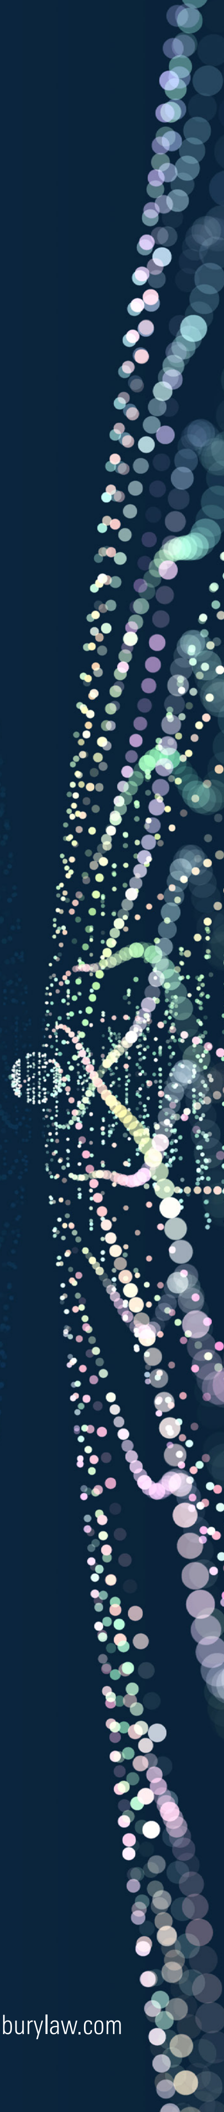

Satellite Data Centers

***by* Jodi A. Goldberg, Robert A. James & Clarence H. Tolliver**

Data centers … in space.

How did the "final frontier" become a genuine consideration for siting and constructing next-generation data centers? Perhaps it is the inevitable result of demand greatly outstripping supply in two of the great pillars of the data center ecosystem: real estate and power supply. Some 11,800 data centers were reported to already be operating worldwide in 2024. Yet the demand for centers is expected to **rise by as much as 22% annually** until 2030, placing significant constraints on the ability of operators to locate sufficient real estate to build and operate them all, obtain the necessary permits for construction and operation on a timely basis, and, significantly, ensure the availability and reliability of electricity to power and cool the components. Or perhaps growing concerns over security, resiliency and environmental impacts are driving operators and users to look for alternative solutions.

What is indisputable is that players across the data center ecosystem are searching for new solutions that can simultaneously help meet the exponentially growing demand for data volume and circumvent the increasing challenges facing expanding terrestrial infrastructure. While some players are looking seaward—to the harbors, shores and even the ocean deep as future sites of permanent or mobile solutions—others are looking up. The 21st Century's commercial space renaissance has led to orbital and extraterrestrial solutions emerging as legitimate—and arguably frontrunning—contenders to resolve the data center dilemma.

Powering Up

Access to consistent power is the lifeblood of any data center. It is the top entry in any site selection checklist—the gating factor in determining whether, when and where a data center will be built. The power source's vitality is twofold—it enables the main components to operate, and it maintains temperature control to prevent overheating.

Data centers are approaching and exceeding gigawatt-magnitude loads and demand exceptional reliability—up to 99.999% online service—due to the computational and financial costs of power interruptions. The massive power demands and strain on the energy grid have played a key role in more restrictions being placed on siting. Even Loudon County, Va., home to Data Center Alley, has recently eliminated the "by-right" approval process for new data center development due to the strains the operations are placing on the county's energy infrastructure.

Developers have responded by exploring alternative power sources beyond those sources commonly deployed in regional or national energy grids. Nuclear, solar and other renewable energy sources have shown the most promise, particularly as operators look to balance their energy needs against growing environmental impact concerns that the massive upscale in centers and gigawatt hours attract. As the demand for zero-emission output from international and hyperscaler players increases, the lure of accessing near-constant, unfiltered solar energy makes space an attractive destination for future infrastructure expansion.

The Real World of Off-World Infrastructure

Moving data centers off-planet sounds far-fetched, until one recognizes that the baseline infrastructures needed to establish incremental and scalable networks already exist. Satellites, in-orbit laboratories and lunar modules (such as lunar landers) provide the foundational architecture for data storage and processing in space. These platforms employ a combination of radiofrequency and optical laser links to transmit and receive data both to and from the Earth, as well as intra- and inter-networking in space. By transmitting data among spacecraft, networks can relay data around the globe at rates that cannot be matched by terrestrial telecommunications infrastructure. Pioneering companies currently working on space-based data solutions have started from these first principles and are currently developing and deploying initial data centers in the form of low Earth orbit (LEO) satellite networks (**Starcloud** and **Axiom Space**), and modular lunar storage (**Lonestar Data Holdings**).

Aided by the increasingly lower costs and emerging competition for commercial launch services, companies will soon be able to deploy, scale and iterate on technology platforms at essentially the same rate as terrestrial systems.

Lower costs to orbit have also facilitated the development and advancement of space-based services—known as in-space servicing, assembly and manufacturing (ISAMs)—which seek to enable more complex architectures for the space environment through the launch of raw materials and modular units for in-situ fabrication and construction (for example by companies such as **Redwire** or **Spacedock**).

Unlike terrestrial infrastructure, nearly all space-based platforms are designed to rely on solar power as a principal source of energy over the course of the asset's operational lifetime. In this environment, sunlight is nearly constant for most platforms and without the Earth's protective atmosphere operating as a filter, this energy is delivered with high intensity (measured in watts per square meter). (See also Pillsbury's article on **Microwave Transmission of Solar Energy from Space**.) Spacecraft will orient themselves or their solar arrays towards the sun as the platform or celestial body orbits around the Earth. These arrays are used to charge and recharge onboard battery cells (commonly lithium-ion) that are relied upon to power vehicle operations, including transmissions and maneuvering, and to power the spacecraft while operating in the Earth's shadow.

This dependence on solar energy is also true for space-based platforms that will operate from the surface of a celestial body, such as the Moon. Both countries and companies have keyed in on ideal lunar sites for permanent and semi-permanent infrastructure installations based on the availability of near-constant exposure to sunlight. Illuminated up to 88% and receiving sunlight up to 92% of the time on average, the Peaks of Eternal Light along the Hinshelwood, Peary and Whipple craters at the lunar north pole, and the Shackleton and De Gerlache craters at the lunar south pole, are notable locations that space agencies and companies identified as landing targets for past and upcoming lunar missions.

The expansiveness and remoteness of space has also enabled nuclear-powered technologies to garner political and commercial support for resurgence at levels not mirrored for Earth-bound facilities. This is evidenced by the Trump administration's recently released Directive on Fission Surface Power (FSP) Development, identifying nuclear as "both an essential and sustainable segment of the lunar and Mars power architectures." NASA has already released a draft Announcement for Partnership Proposals seeking to "make power available on the lunar surface and to energize the space industrial base to support a future lunar economy through the deployment of a (100 kW) FSP system on the lunar surface." NASA and DOE's Fission Surface Power System (FSPS) has also previously awarded **three $5 million contracts** for the design, cost estimation and development schedule of an up to four-by-six meter, 6,000-metric-ton fission reactor capable of delivering at least 40 kW of continuous electric power over 10 years by 2030. Though the electrical output of a single FSPS reactor might only satisfy modest data center demands, such reactors could be scaled to meet the growing load demands of data centers and other lunar infrastructure. More importantly, these initiatives signal the government's commitment to partner with industry in the development and advancement of nuclear technologies for sustainable space infrastructure.

Security and Resilience

In-space data centers also provide solutions for players looking for added resiliency, security and data sovereignty. As described in the preceding section, most space-based systems provide a physically diverse pipeline through which customer data can travel, apart from terrestrial telecommunications infrastructure. While some satellite networks intentionally select ground stations that are collocated with data centers (**AWS Ground Station**), the space-based data center paradigm would provide the reverse arrangement for purposes of resilient or redundant capacity.

Physical separation also enhances the security of the assets, making them harder to access and reduces the likelihood of intrusion. For much of the data stored on-orbit, particularly in celestial modules, the data would be in a form of long-term or backup storage, where users are not accessing the information on a constant basis and the modules function as a sort of "vault."

As the nationality of the spacecraft is established based on the licensing authority or the launching state of the object, these platforms also offer a unique solution for data sovereignty. As countries adopt additional requirements for companies and individuals to hold certain information within the country of origin, space-based data centers offer additional locations in which data can be stored within the dictates of the regulations.

High Risk, High Reward

If space-based data centers are essentially an adaptation or extension of existing technologies, why have they not yet been deployed en masse? Well, as the cliché goes, space is hard.

For starters, while the foundational space platforms exist, the data center technology must still be adapted for the harsh environments of space. Adapted or new technologies for performing the computing functions of data centers must be developed and tested before being deployed into the space ecosystem. Once launched, the platform itself can typically only be modified or repaired in minimal ways, such as with software updates. Physical changes to the platform would require the operator to build and launch a replacement spacecraft. While ISAMs may introduce these capabilities in the future, today, many defects will result in the premature end or severe diminishment of the spacecraft.

Costs of accessing space have plummeted over the last decade, but building and sustaining infrastructure projects in space are still massively expensive, capital-intensive endeavors. Beginning with the upfront costs — the minimum-viable network must be designed, built, launched and commissioned before any revenue can be earned on services — the operators must also plan for ongoing operational costs and replacement vehicles. In low Earth orbit, spacecraft are typically designed for operational lifetimes of five to eight years before they are deorbited and a new spacecraft launched as a replacement. On the lunar surface, the assets will likely be hardened to operate for longer periods of time (using geostationary orbit as a benchmark, 15 – 20 years), but its trade-off is that there is less opportunity to iterate on the technology so each future module must also be backwards compatible to interface with existing infrastructure. There are also associated costs for deploying (or leasing) and operating the ground network that communicates with the in-space assets, and which are situated in a multitude of jurisdictions all across the world.

Operating in space carries its own risks, from initial launch through deorbit. While SpaceX's Falcon 9 is beginning to make launch look routine, generally launch still has a distance to go before it becomes as routine and reliable as land and air cargo transport. Competition is also relatively nascent, so ensuring every operator that wants regular launch services can secure the volume they need may be difficult over time unless the launch sector right-sizes at the same pace.

Once on-orbit, operators must navigate additional challenges. As mentioned above, objects in space cannot be repaired in the same way they can on Earth. Space assets must also monitor for and avoid collisions with other objects, including other satellites and stations, as well as debris and meteorites. Spacecraft must also be hardened to protect against radiation and other space weather. The very same solar radiation that powers the solar arrays can cause serious damage. Fluctuations of solar wind and geomagnetic storms can corrupt software, cause short-circuiting and reduce solar cell output.

Despite the extreme cold of space, maintaining a regulated cooling system in the sealed environment of a data center may also introduce its own challenges. The air-based cooling employed in terrestrial data centers is not an option in the vacuum of space. The heat generated by operating centers would have to be ejected via suitably designed large radiators that could still maintain the integrity of the space station's structure. The cooling system would also have to adapt to rapidly changing temperatures based on the center's immediate level of sun exposure. And while liquid-based cooling systems have **increased in popularity** here on Earth, it is unclear whether similar systems would have equal effect in the lunar context, considering the need to account for the potential impacts of lower gravity environments on fluid behaviors and the associated costs of transporting coolants from Earth to the lunar site.

Finally, there is still some regulatory uncertainty around operating data centers in space. While the regulations around the licensing of satellite networks are fairly well established (though different in every country), there is less certainty and consistency around ISAMs and use of lunar and celestial resources. In the United States, the Trump administration recently released an Executive Order instructing the Commerce Secretary to propose new regulations surrounding all novel space activities not currently covered by existing regulations. While orbital data centers arguably rightly fall under the FCC's licensing regime for communications satellites, the case is less clear for modules operating on the lunar surface. Additionally, as we previously covered in Pillsbury's article on **Lunar Natural Resources**, the ability to deploy permanent and semi-permanent structures on the lunar surface is also a matter of international legal debate and the provenance of the first Trump administration's successful Artemis Accords.

A New Space Race?

With all these considerations in mind, at least three separate organizations have started or are nearing the stage where they can demonstrate proof of concept.

- **Lonestar Data**: With its sights set on the Moon, Lonestar is the first to launch and demonstrate an initial proof of concept. Earlier in 2025, the company sent a data-carrying device the size of a shoebox to the Moon as part of Intuitive Machines' Athena Lunar Lander to test its ability to upload, download and transfer data. Despite the failure of the Intuitive Machines lander, the Lonestar device successfully completed several of its test objectives. The company **plans** to launch a data center to orbit the Moon in 2027, and plan to establish data centers on the lunar surface.
- **Starcloud**: Starcloud is working to build a network of low Earth orbit data center satellites, powered by a grid of solar panels that would extend approximately two and a half miles wide. Tests for its data center network will begin with satellite launches in 2025 and 2026. Its goal is to have a network up and running by the early 2030s, deploying 40-megawatt orbital data centers powered by solar energy.
- **Axiom Space**: On top of attempting to create the first commercial space station (which will have a data center onboard), Axiom aims to create its own network of orbital data center nodes in low Earth orbit, beginning with **launches** scheduled for late 2025. These nodes are intended to provide secure, scalable and cloud-enabled data storage and processing, and artificial intelligence/machine learning solutions directly to satellites and other spacecraft with the capability to operate independently of terrestrial infrastructure.

Thus, what may have sounded like science fiction when you started reading this article may actually be coming to an orbital arc near you in just a few short years. ❖

Appendices

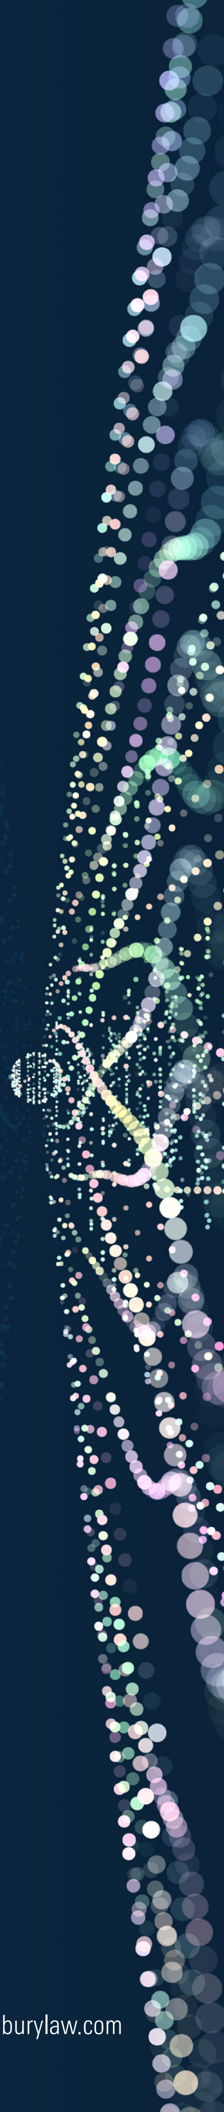

Map of Data Centers and Resources

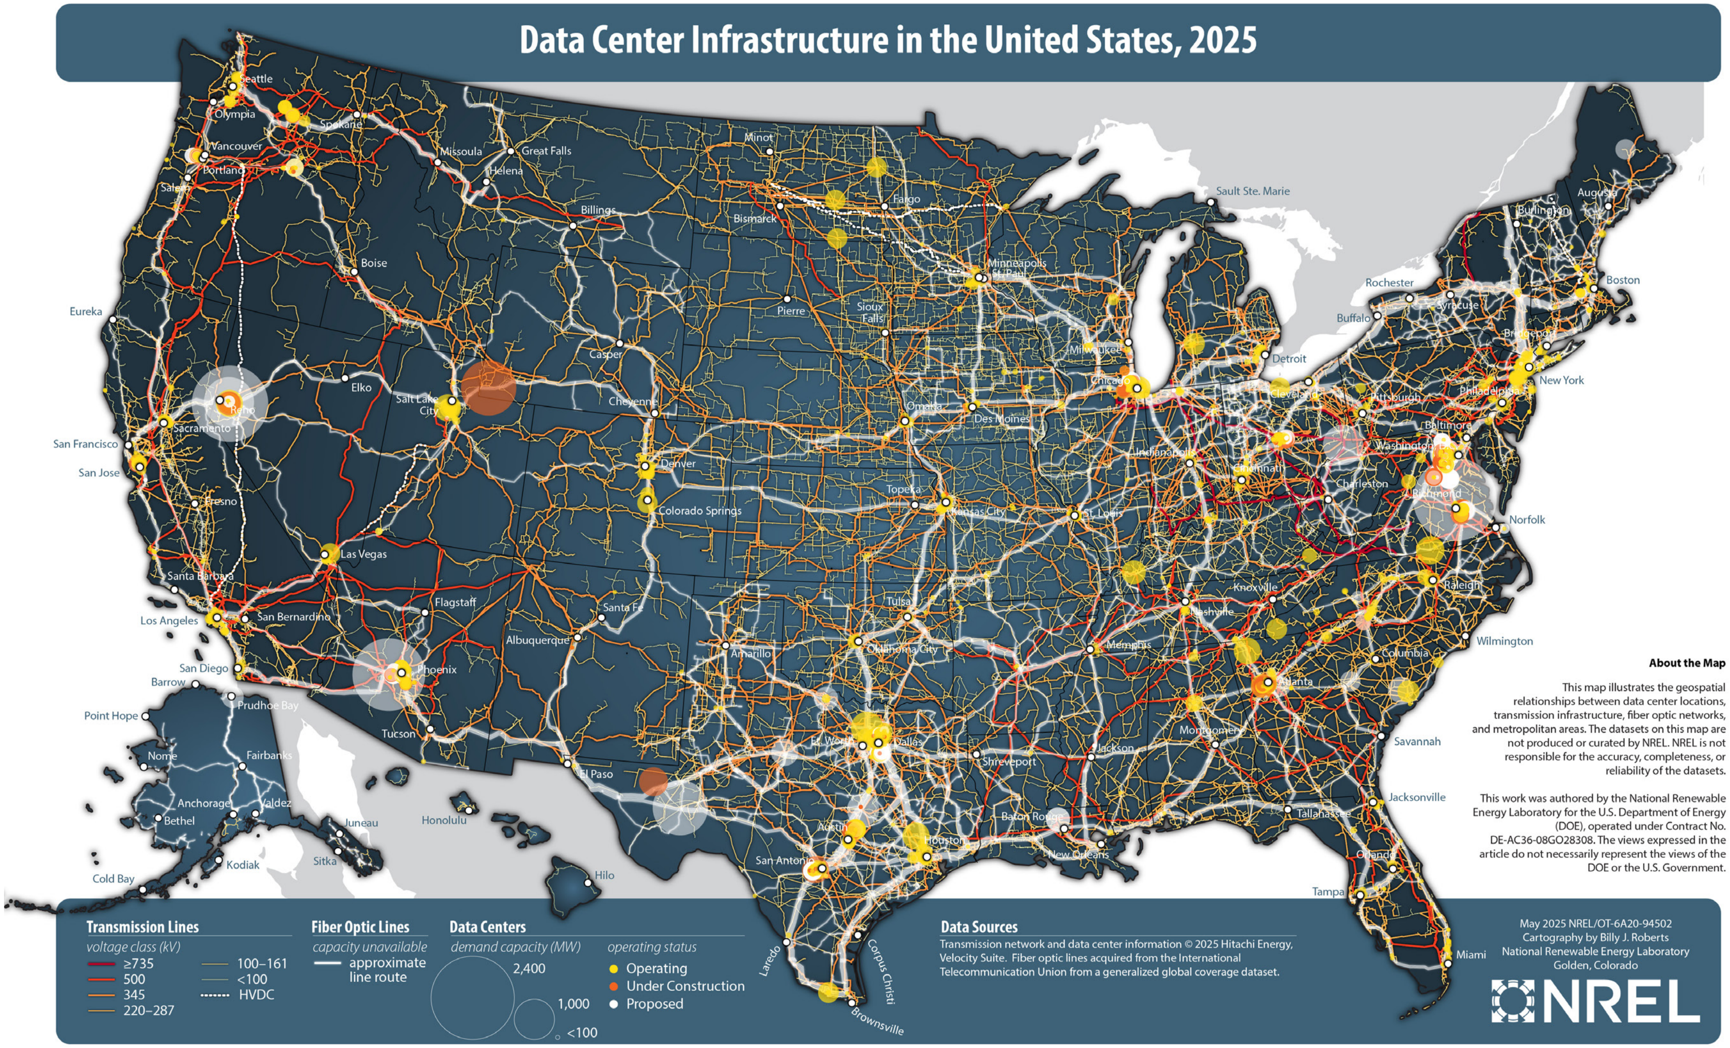

About the Map

This map illustrates the geospatial relationships between data center locations, transmission infrastructure, fiber optic networks, and metropolitan areas. The datasets on this map are not produced or curated by NREL. NREL is not responsible for the accuracy, completeness, or reliability of the datasets.

This work was authored by the National Renewable Energy Laboratory for the U.S. Department of Energy (DOE), operated under Contract No. DE-AC36-08GO28308. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government.

Data Sources

Transmission network and data center information © 2025 Hitachi Energy, Velocity Suite. Fiber optic lines acquired from the International Telecommunication Union from a generalized global coverage dataset.

May 2025 NREL/OT-6A20-94502
Cartography by Billy J. Roberts
National Renewable Energy Laboratory
Golden, Colorado

NREL

Source: National Renewable Energy Laboratory (high-resolution version available by clicking on the map.)

Roster of Data Center Proprietors and Locations

Top 50 Data Center Proprietors (by # of locations)

1. Zenlayer (374)
2. China Telecom (362)
3. Lumen (308)
4. Equinix (270)
5. Digital Realty (255)
6. IPTP Networks (220)
7. Vantage Data Centers (131)
8. Csquare (112)
9. Cogent Communications (109)
10. NTT Global Data Centers (105)
11. CyrusOne (103)
12. MOD Mission Critical (99)
13. DataBank (91)
14. 365 Data Centers (86)
15. Hivelocity (57)
16. QTS Data Centers (46)
17. Cologix (45)
18. Enzu (44)
19. TierPoint (41)
20. Telia Company (40)
21. Flexential (39)
22. Aligned Data Centers (39)
23. Iron Mountain Data Centers (38)
24. Telehouse (34)
25. H5 Data Centers (30)
26. Vultr (30)
27. CoreSite (28)
28. Crown Castle (28)
29. AT&T (26)
30. AtlasEdge Data Centers (24)
31. FDCservers.net (23)
32. NorthC (23)
33. Ascenty (23)
34. Prime Data Centers (23)
35. Colocation America (22)
36. Cirion Technologies (20)
37. Evocative (19)
38. LightEdge (19)
39. Verizon (19)
40. Rackspace (18)
41. iTel Networks Inc. (18)
42. US Signal (17)
43. Limestone Networks (17)
44. Summit (17)
45. Psychz Networks (16)
46. Aptum Technologies (16)
47. Latitude.sh (16)
48. DC BLOX (15)
49. eStruxture (15)
50. FirstLight (14)

Source: DataCenters.com

Roster of Data Center Proprietors and Locations

50-State Survey of Data Centers (United States)

1. Virginia (OPERATIONAL: 320, UNDER CONSTRUCTION: 144)
2. Texas (OPERATIONAL: 212, UNDER CONSTRUCTION: 156)
3. California (OPERATIONAL: 166, UNDER CONSTRUCTION: 7)
4. Ohio (OPERATIONAL: 101, UNDER CONSTRUCTION: 60)
5. Oregon (OPERATIONAL: 98, UNDER CONSTRUCTION: 13)
6. Arizona (OPERATIONAL: 82, UNDER CONSTRUCTION: 36)
7. Illinois (OPERATIONAL: 78, UNDER CONSTRUCTION: 21)
8. Washington (OPERATIONAL: 71, UNDER CONSTRUCTION: 3)
9. Georgia (OPERATIONAL: 61, UNDER CONSTRUCTION: 57)
10. Iowa (OPERATIONAL: 59, UNDER CONSTRUCTION: 15)
11. Florida (OPERATIONAL: 54, UNDER CONSTRUCTION: 1)
12. New York (OPERATIONAL: 50, UNDER CONSTRUCTION: 4)
13. New Jersey (OPERATIONAL: 49, UNDER CONSTRUCTION: 3)
14. North Carolina (OPERATIONAL: 43, UNDER CONSTRUCTION: 12)
15. Pennsylvania (OPERATIONAL: 38, UNDER CONSTRUCTION: 13)
16. Tennessee (OPERATIONAL: 32, UNDER CONSTRUCTION: 2)
17. Michigan (OPERATIONAL: 31, UNDER CONSTRUCTION: 2)
18. Colorado (OPERATIONAL: 31, UNDER CONSTRUCTION: 6)
19. Utah (OPERATIONAL: 30, UNDER CONSTRUCTION: 10)
20. Indiana (OPERATIONAL: 28, UNDER CONSTRUCTION: 28)
21. Minnesota (OPERATIONAL: 27, UNDER CONSTRUCTION: 4)
22. Nevada (OPERATIONAL: 27, UNDER CONSTRUCTION: 29)
23. Nebraska (OPERATIONAL: 26, UNDER CONSTRUCTION: 6)
24. Missouri (OPERATIONAL: 24, UNDER CONSTRUCTION: 10)
25. Massachusetts (OPERATIONAL: 23)
26. Oklahoma (OPERATIONAL: 22, UNDER CONSTRUCTION: 11)
27. South Carolina (OPERATIONAL: 18, UNDER CONSTRUCTION: 13)
28. Wisconsin (OPERATIONAL: 18, UNDER CONSTRUCTION: 14)
29. Alabama (OPERATIONAL: 15, UNDER CONSTRUCTION: 8)
30. New Mexico (OPERATIONAL: 13, UNDER CONSTRUCTION: 5)
31. Wyoming (OPERATIONAL: 12, UNDER CONSTRUCTION: 10)
32. Maryland (OPERATIONAL: 12, UNDER CONSTRUCTION: 12)
33. Mississippi (OPERATIONAL: 10, UNDER CONSTRUCTION: 22)
34. Louisiana (OPERATIONAL: 10, UNDER CONSTRUCTION: 13)
35. Kentucky (OPERATIONAL: 10, UNDER CONSTRUCTION: 2)
36. Kansas (OPERATIONAL: 8, UNDER CONSTRUCTION: 1)
37. Connecticut (OPERATIONAL: 7)
38. Delaware (OPERATIONAL: 6)
39. New Hampshire (OPERATIONAL: 6)
40. District of Columbia (OPERATIONAL: 5)
41. West Virginia (OPERATIONAL: 4)
42. Arkansas (OPERATIONAL: 4, UNDER CONSTRUCTION: 3)
43. Montana (OPERATIONAL: 4)
44. Idaho (OPERATIONAL: 4, UNDER CONSTRUCTION: 2)
45. Maine (OPERATIONAL: 4)
46. Rhode Island (OPERATIONAL: 3)
47. South Dakota (OPERATIONAL: 2)
48. North Dakota (OPERATIONAL: 2, UNDER CONSTRUCTION: 4)
49. Hawaii (OPERATIONAL: 2)
50. Vermont (OPERATIONAL: 1)
51. Alaska (OPERATIONAL: 0)

Source: DataCenters.com

Roster of Data Center Proprietors and Locations

Key Countries with Data Centers

1. United States (3670)
2. China (561)
3. Canada (478)
4. United Kingdom (453)
5. Germany (387)
6. France (345)
7. Japan (253)
8. Australia (241)
9. Brazil (236)
10. India (231)
11. Netherlands (226)
12. Italy (165)
13. Spain (163)
14. Indonesia (157)
15. Singapore (134)
16. Hong Kong (103)
17. Russia (89)
18. Switzerland (82)
19. Sweden (82)
20. Ireland (81)
21. Mexico (79)
22. Malaysia (75)
23. Poland (73)
24. South Korea (71)
25. Austria (68)
26. Reino Unido (65)
27. South Africa (62)
28. Chile (60)
29. United Arab Emirates (56)
30. Belgium (50)
31. Norway (49)
32. Denmark (47)
33. Israel (47)
34. Taiwan (44)
35. Argentina (43)
36. Colombia (43)
37. Finland (42)
38. New Zealand (40)
39. Bulgaria (39)
40. Paises Bajos (38)
41. Francia (36)
42. Vietnam (36)
43. Thailand (34)
44. Portugal (32)
45. Greece (29)
46. Turkey (28)
47. Brasil (27)
48. Romania (26)
49. Nueva Zelanda (25)
50. Polonia (23)

Source:
DataCenters.com keeps an up-to-date roster of **top 50 providers**, **2025 top providers** and **global data center's locations**.

News, Articles and Podcasts

Data Center Trends 2025: Power, AI, and Progress

Data Center Dynamics (DCD) is the world's largest data center publication that serves as a hub for professionals and stakeholders to stay informed about the latest developments, trends, and best practices related to various aspect of data centers, including infrastructure, technology, energy efficiency and security. Content ranges from news articles and updates to case studies and whitepapers.

Additionally, DCD offers DCD>Academy, personalized training programs for companies, and DCD>Intelligence, launching in 2025 an interactive market intelligence platform.

Data Center POST

Datacenter.com is a technology platform for colocation, bare metal, and cloud services within the data center industry providing information on relevant news and trending topic. In addition, they offer various free resources.

Data Center Knowledge

Data Center Knowledge, part of Informa TechTarget, offers in-depth articles on data center technologies, operations, and trend. Covering topics like energy efficiency, cooling solutions, and AI integration.

DatacenterHawk

DatacenterHawk provides the most accurate and useful data center real estate information on the market.

Data Centre & Network News

Data Centre Magazine covers global data center news, focusing on technological advancements, industry leadership, data center operations, and strategic developments.

Consulting Firms

CBRE | Data Center Insights & Research

CBRE maintains a dedicated data center platform offering advisory and transactional services across the data center lifecycle, including site selection, investment sales, valuation, and occupier services. CBRE also publishes regular market reports and outlooks analyzing global data center demand, capacity growth, AI-driven expansion, and power availability trends.

FTI Consulting | Data Centers

FTI Consulting maintains a data centers and digital infrastructure practice advising clients on strategy, transactions, operations, cybersecurity, and disputes across the data center, fiber, telecom, and cloud sectors. FTI also publishes reports and insights addressing AI infrastructure demand, investment trends, and delivery risks associated with next-generation data center growth.

Ernst & Young | Data Centers

Ernst & Young maintains a data centers page offering consulting services and relevant news. Adam Blaylock recently wrote an article for EY addressing "Why there is not silver bullet for data center financing."

JLL | 2026 Global Data Center Outlook

Jones Lang LaSalle (JLL) publishes an annual Global Data Center Outlook, analyzing capacity growth and investment trends.

Schneider Electric USA | EcoConsult for Data Centers

Schneider Electric offers consulting services, EcoConsult, for data center operators looking to boost efficiency and maximize uptime.

Pillsbury's Integrated Data Center Capabilities

Pillsbury's Data Centers Team works with data center developers and operators, data center customers, power providers, investors, real estate developers and others who develop data centers and keep them running without interruption. We are eager to apply our extensive experience and unwavering commitment to deliver not only exceptional legal services but also strategic guidance and insightful recommendations that will support your objectives.

WHY PILLSBURY?

WE ARE A FULL-SERVICE FIRM FOR THE DATA CENTER ECOSYSTEM.

Pillsbury's Data Centers team delivers operators, developers, investors, power suppliers, and users the counsel required to address all aspects of the data center landscape. Our team's experience with real estate, construction, project development and finance, sourcing, cybersecurity, consumer protection and intellectual property issues directly translates to acquiring, developing, financing, powering, and operating data centers, and protecting the data and privacy interests associated with them. See slide 4 for a complete listing of our services.

WE ARE A LEADING FIRM IN THE POWER SECTOR, INCLUDING NATURAL GAS, RENEWABLES AND NUCLEAR ENERGY.

Pillsbury stands at the forefront of the power sector, including renewable and nuclear energy, with decades of experience advising on transformative transactions in the power generation industry. Our team has played a central role in projects that have collectively contributed tens of thousands of megawatts of clean energy to the grid, powering data centers and other energy-intensive facilities. We are consistently recognized as one of the top law firms in the renewable energy and energy transition space, with expertise spanning key regions, including SPP, CAISO, MISO, ERCOT and PJM. Our client roster includes eight of the top ten U.S. renewable energy companies by revenue.

Particularly unique is our expertise in advanced nuclear energy, considered key to future data center development.

WE ARE RECOGNIZED AS A LEADING TECHNOLOGY TRANSACTIONS AND DIGITAL INFRASTRUCTURE FIRM.

Pillsbury's Global Sourcing & Technology Transactions team is consistently recognized among the leading technology and outsourcing practices in the market, including Band 1 rankings and other top-tier industry recognitions. We advise clients ranging from emerging technology companies to global enterprises on complex technology, infrastructure, and commercial transactions, with particular depth in cloud, AI, digital infrastructure, data center, cybersecurity, and managed services matters. Our lawyers regularly structure and negotiate mission-critical commercial arrangements involving hyperscale and colocation facilities, GPU-as-a-service and cloud capacity arrangements, fiber and network infrastructure, technology procurement, outsourcing, SaaS and other strategic services agreements. We have extensive experience advising clients operating in highly regulated and critical infrastructure sectors, including energy, financial services, communications, healthcare, and government/defense.

WE OFFER A FLEXIBLE APPROACH AND EXCEPTIONAL TEAM.

We are committed to adding value to our clients' bottom lines by enabling them to execute transactions more effectively and efficiently. We take a proactive approach to identifying and addressing critical business and legal issues, ensuring that our clients are well-positioned to complete their deals successfully.

DATA CENTERS

The Ecosystem

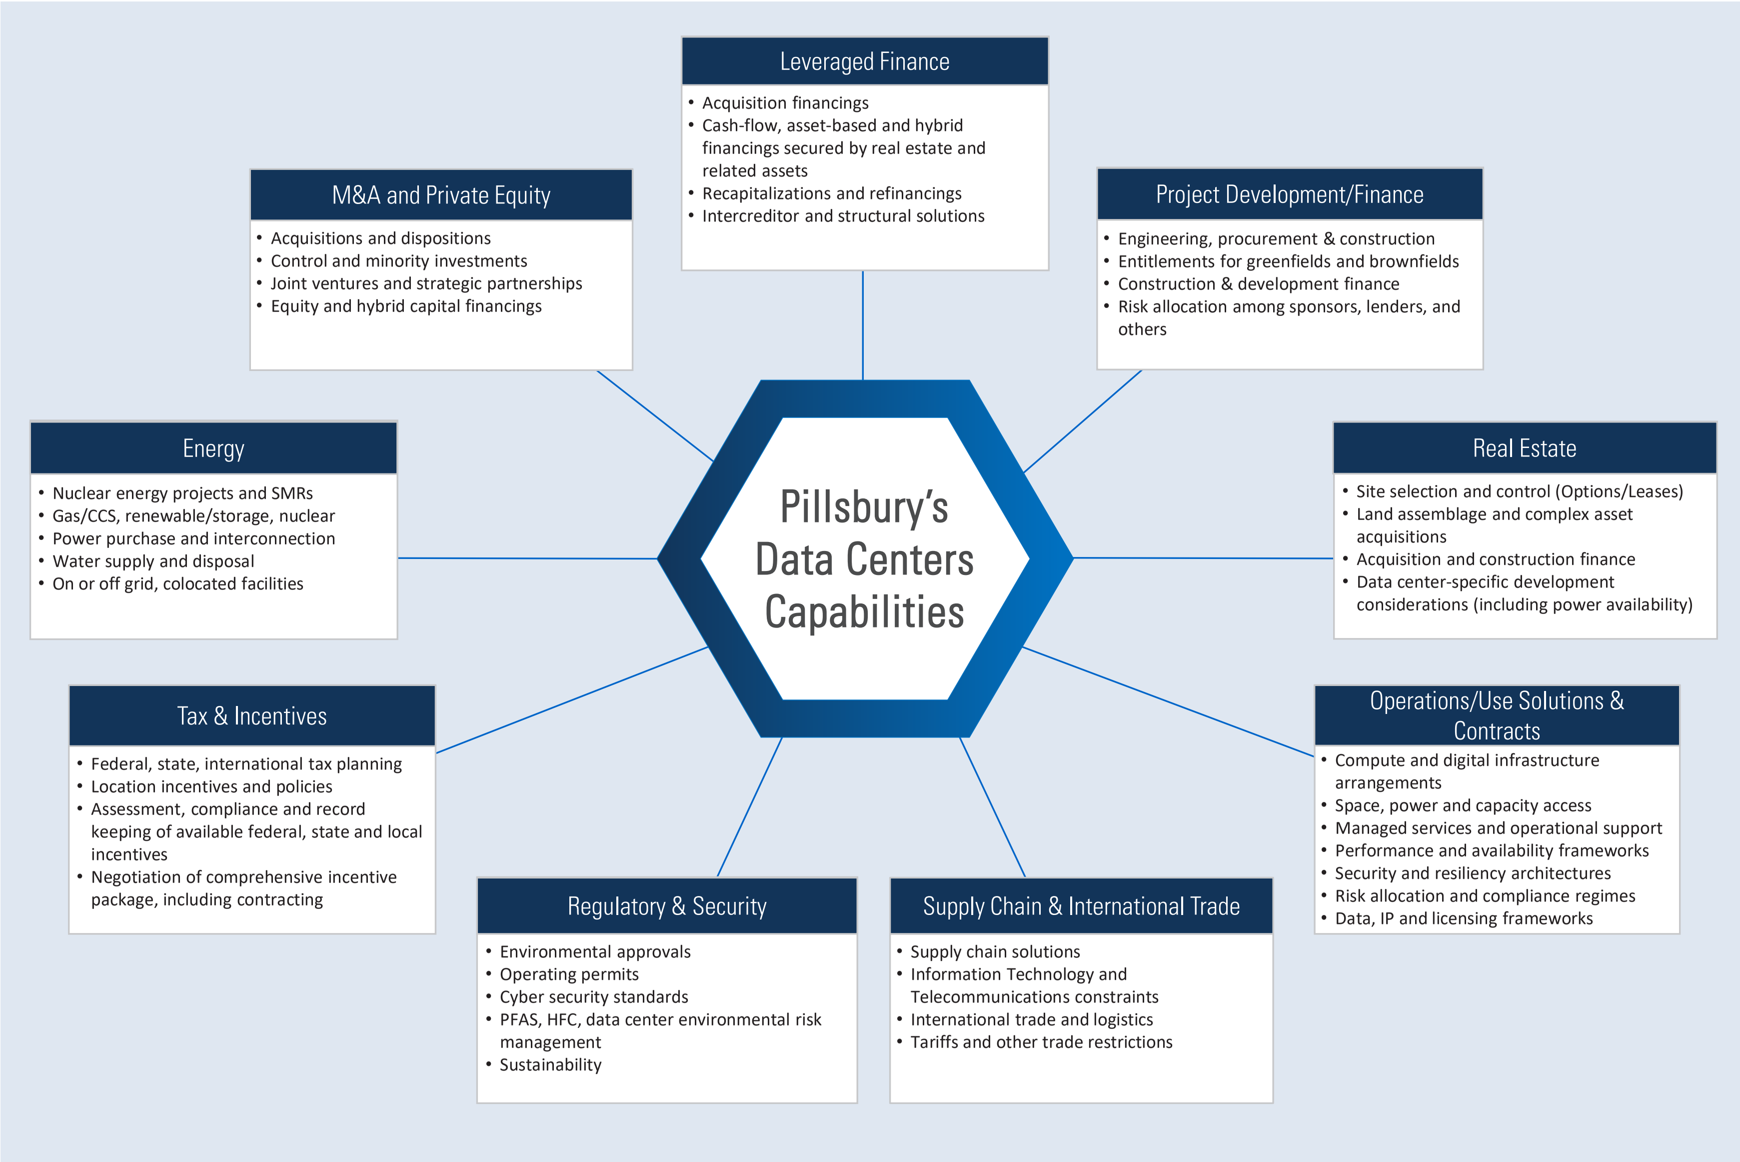

Pillsbury's Data Centers Capabilities

Leveraged Finance

- Acquisition financings
- Cash-flow, asset-based and hybrid financings secured by real estate and related assets
- Recapitalizations and refinancings
- Intercreditor and structural solutions

Project Development/Finance

- Engineering, procurement & construction
- Entitlements for greenfields and brownfields
- Construction & development finance
- Risk allocation among sponsors, lenders, and others

Real Estate

- Site selection and control (Options/Leases)
- Land assemblage and complex asset acquisitions
- Acquisition and construction finance
- Data center-specific development considerations (including power availability)

Operations/Use Solutions & Contracts

- Compute and digital infrastructure arrangements
- Space, power and capacity access
- Managed services and operational support
- Performance and availability frameworks
- Security and resiliency architectures
- Risk allocation and compliance regimes
- Data, IP and licensing frameworks

Supply Chain & International Trade

- Supply chain solutions
- Information Technology and Telecommunications constraints
- International trade and logistics
- Tariffs and other trade restrictions

Regulatory & Security

- Environmental approvals
- Operating permits
- Cyber security standards
- PFAS, HFC, data center environmental risk management
- Sustainability

Tax & Incentives

- Federal, state, international tax planning
- Location incentives and policies
- Assessment, compliance and record keeping of available federal, state and local incentives
- Negotiation of comprehensive incentive package, including contracting

Energy

- Nuclear energy projects and SMRs
- Gas/CCS, renewable/storage, nuclear
- Power purchase and interconnection
- Water supply and disposal
- On or off grid, colocated facilities

M&A and Private Equity

- Acquisitions and dispositions
- Control and minority investments
- Joint ventures and strategic partnerships
- Equity and hybrid capital financings

CASE STUDY

Helping a Client Navigate Complex Power Purchase Agreements with a Leading Hyperscaler

OPPORTUNITY

In a landmark project attracting global attention as a model for the future, a leading hyperscaler chose advanced nuclear technology to supply power to its data center population in rural Washington state, involving several power providers.

Our client, a major public power company, was selected as a provider of that power. The client turned to Pillsbury to navigate the complex multi-party relationships supporting the underwriting of advanced nuclear technology, power purchase agreements and related infrastructure agreements.

Small modular nuclear reactors (SMRs) and fuel are at the heart of these transactions. Our client in Washington state and utilities in Virginia are engaged on these projects. Four reactors will initially generate 320 megawatts of capacity for our client, but the project could increase to 960 megawatts, enough to power approximately 770,000 homes.

The small footprint and modular construction of SMRs allow them to be constructed more quickly and in more places than traditional nuclear reactors.

The announcements come as tech companies scour the nation for electricity to power their increasingly power-hungry data centers. The U.S. government and private firms have been trying to advance SMR technology for years.

Tech industry power needs, however, are giving SMRs new momentum. AI developers see the plants as a potential major source of power that can scale quickly and provide energy around-the-clock. Data centers—and tech companies' willingness to pay above-market prices for reliable and carbon-free energy—could be what SMRs need to establish themselves as viable power generation resources.

Data Centers Experience: Energy

Pillsbury's Data Centers Team works with data center developers and operators, data center customers, power providers, investors, real estate developers and others who develop data centers and keep them running without interruption. We are eager to apply our extensive experience and unwavering commitment to deliver not only exceptional legal services but also strategic guidance and insightful recommendations that will support your objectives.

Energy for Power & Cooling: Bridge, Conventional and Zero-Emission PPAs and VPPAs, On and Off the Grid

- Assisting on long-term strategic planning for advanced large-scale sources of zero-emission power for data centers. Next-generation sources include advanced nuclear fission (small modular reactors (SMR) and Generation IV larger reactors), advanced geothermal, nuclear fusion, and advanced sodium-based storage solutions.
- Advising on a broad range of physical power purchase agreements (PPAs) and virtual power purchase agreements (VPPAs) for supply of electricity for data centers.
- Representing a utility in connection with contemplated power purchase agreement for data center needs in connection with development, construction, financing, and operation of SMRs.
- Review of numerous data center power purchase agreements (virtual and involving energy delivery) in connection with acquisitions and contemplated acquisitions of renewable energy generation asset portfolios.
- Representing a large renewable energy producer in connection with contemplated development, construction, and operation of small modular nuclear reactor (SMR) designed to meet data center energy needs.
- Representing a joint venture in proposing zero emission power sources for hyperscale facilities, including colocation of data centers and modular reactor projects.
- Representing a joint venture in developing a data center and crypto mining campus colocated behind the meter with a power generation facility converting from conventional fuel sources to hydrogen.
- Working with an international environmental organization on its Nuclear Scaling Initiative, which seeks to aggregate demand for SMRs among energy intensive end-users, including data centers.
- Working with a separate international environmental organization on a contracting methodology to spread the risks for first-of-a-kind (FOAK) nuclear reactor design, construction and operation amongst a group of users, including data centers.
- Advising a startup company seeking to develop and deploy nuclear assets for generating the power needs of data centers and other industrial users.
- Advising several SMR and advanced reactor nuclear vendors on regulatory and generation issues associated with power development for data centers.
- Advising a national energy association on regulatory and policy matters supporting the deployment of SMRs and advanced reactors for a variety of commercial and industrial applications.

CASE STUDY

Deploying Data Centers at Scale Across the Country

OPPORTUNITY

Our client, a global data centers operator and services provider, has been actively expanding its U.S. operations aligning with key trends in the data center market, particularly the growing demand for AI infrastructure, sustainability, and strategic geographic diversification. It operates nearly one million square feet of data center space across seven U.S. campuses, including locations in Virginia, California, Texas, Illinois, Oregon and Arizona. The surge in AI workloads is transforming the data center industry, driving unprecedented demand for power-intensive infrastructure. This client's expansion and investments position it to meet this growing need.

This client turned to Pillsbury to standardize and streamline its acquisition and development processes to ensure that it can quickly and efficiently capitalize on opportunities regardless of which state or jurisdiction the property is located. This client's presence in diverse U.S. markets reflects a strategic approach to site selection, considering factors like power availability, connectivity and regional demand, which are critical in the current data center landscape.

Our work includes a standardized bespoke set of documents tailor-made to address data center development and site selection issues (instead of relying on generic forms or the counterparty's preferred forms). Not having to start over from scratch each transaction helps this client execute much more efficiently. Finally, as developers have caught on and are now creating "data center ready" sites that include access to utilities, completed power studies, and in some cases are even pre-entitled, these sites are flying off the shelves—so being able to act quickly is critical.

Developers aiming to attract major data center companies are increasingly tailoring their site development strategies to meet the unique and evolving needs of hyperscale operators. Developers prioritize sites with access to robust and scalable utility power, especially in areas with redundant grid infrastructure. With operators committed to sustainability, proximity to wind, solar or hydroelectric power sources with storage is also a major factor in the site selection process. Developers work closely with utilities to ensure quick deployment of substations and reduced time to energization, while also working with state and local governments to pre-entitle land to remove permitting friction.

As demand for "data center ready" sites in key markets across the country continues to rise, institutional players who have a tested strategies and a legal framework tailored to address specific data center needs and risks are well positioned to capitalize on those opportunities.

Data Centers Experience: Development & Investment

Acquisitions, Leases and Project Development

- Represented a major energy company on the transition of its data center requirements from a proprietary physical facility in Texas to a cloud-based system of a major technology company. We advised on the sale, lease, transfer or termination of the data center real and personal property assets, leases, contracts, and associated technology and the negotiation of the long-term data services contracts and licenses.
- Counseled the owners of a data center business in connection with an investment by a leading merchant bank and buy-out of existing investors.
- Represented a Japanese technology company in the sale and long-term leaseback of its one million-square-foot Texas data center campus, allowing the client the opportunity to monetize their asset and secure below market lease rates for operational efficiencies.
- Represented the owner in the sale and leaseback of a 200,000-square-foot data center in Quincy, WA.
- Representing a large renewable energy producer in connection with contemplated sale of interest in data center, together with related loan and power purchase agreement restructurings.
- Represented a real estate investment firm on its $212 million divestiture of a Paris data center.
- Represented a major data center developer on its $874 million acquisition of "Project Cheetah" assets, a divestiture directed by the European Commission.
- Represented the acquirer of a nationwide provider of data center and cloud infrastructure in the UK.

CASE STUDY

Financial Technology Data Center Solutions for a G-SIFI Client

OPPORTUNITY

A Global Systemically Important Financial Institution (G-SIFI) recently embarked on a multibillion-dollar technology modernization program to drive innovation, reduce technical debt and improve operational resilience across its global infrastructure.

A key pillar of this effort involved consolidating the company's 79 data centers down to 13 critical locations (Japan, Singapore, Hong Kong, China, Germany, Ireland, Canada and the United States) and migrating core workloads to public cloud platforms, including both AWS and Microsoft Azure, with cloud spend commitments approaching $1 billion. To help execute this ambitious transformation, this G-SIFI sought the advice of Pillsbury to lead the legal and commercial negotiations across numerous high-impact transactions.

Our work includes:

- A 7-year, $500 million AWS agreement
- 5-year, $250 million Microsoft Azure commitment
- $250 million managed services contract covering colocation data centers and "smart-hands" services globally
- Negotiating complex colocation data center lease and buildout terms for all data centers in the reimagined portfolio
- Advising on multiple sales, leases and leasebacks of data center properties
- Providing ongoing hypercare services to both client and its vendors, including restructuring key migration agreements to drive efficiency and accountability with vendors

Data Centers Experience: Customer-Facing Solutions

The Core Services Relationship Between Data Centers & Users

- Advised numerous Fortune 100 and emerging growth companies, including in regulated industries like financial services and healthcare, in connection with data centers.
- Activity includes consolidation initiatives, data center migrations, colocation deals, and information technology outsourcing deals (ITOs) that include data center support on a remote and hosted basis. Additionally, our team has negotiated data center leases and agreements for space and services for our clients with nearly all of the leading data center providers, which has given us deep familiarity with their market terms, structures, and other essential aspects of these types of deals.
- Negotiated a large outsourcing that includes data center hosting, end-user support, help desk, application management) and claims processing.
- Advised a public entity on sourcing of data center services (including colocation and hosting), shifting from older mainframe-centric data centers to more modern high-density server-centric data centers.
- Represented a retail company in a full scope IT outsourcing, including data center, network (voice and data), end user computing, store support, web hosting, and applications development and maintenance services.
- Representing a life insurance company in the negotiation of multiple colocation agreements in various data centers.
- Represented multiple U.S. corporations establishing and operating IT capability centers in India.

CASE STUDY

IP Protection for Critical Data Center Infrastructure

OPPORTUNITY

A client engaged in providing data center solutions wanted to protect its innovative solutions, which included:

- Semi-freight container sized air-handling unit capable of handling many different environmental conditions throughout each 24- hour period in the desert ('782 patent)
- Integrated wiring system and thermal shield system within aisles of the cabinets, capable of handling wiring for telecommunications, power, and internal data communication ('780 patent)
- Overhead-air air handling with multi-level roof ('538 patent)
- Control system for air handling in a data center ('495 patent)
- Various efficient facility configurations ('331 and '643 patents)
- And numerous others

Pillsbury worked directly with the client in putting together a broad patent portfolio that has been licensed, asserted, and recognized for its value, representing a key asset of the company.

Data Centers Experience: Counseling

Mission-Critical Advice on Operations, Physical and Cybersecurity, and Compliance

- Developing due diligence program and action plan for all permits and approvals required for buildout of data centers in both established and key emerging jurisdictions.
- Assisting several nuclear utilities on physical security and cybersecurity issues associated with current and advanced reactors.
- Advising on artificial intelligence applications for data centers, including operation, data center infrastructure management (DCIM), power, cooling, and cybersecurity threat detection and response.
- Assisted a data center developer in obtaining patents for technologies related to the efficient cooling of data centers.
- Advising on complete range of information technology and telecommunications technology solutions, including patent prosecution, licensing, enforcement
- Advising on regulatory compliance advice germane to data centers, including:
    - environmental compliance standards
    - sustainability standards for electrical usage and end-of-life management of e-waste
    - Advice on HFC and other regulated chemicals in HVAC systems
    - general OSHA provisions pertaining to high voltage electrical usage
    - fire protection and PFAS regulation for fire suppression equipment
    - facility and contractor safety policies
    - data protection and cybersecurity policies and processes
    - trade issues associated with the import and export of "dual-use" IT hardware, software and technology
    - tenant insurance coverage
    - waste management

CASE STUDY

Data Center Alliance

OPPORTUNITY

Pillsbury advised Client X, one of Japan's largest energy and industrial companies, on a strategic investment in Investor Y's flagship Texas data center.

Pillsbury evaluated the strength of the operation and customer contracts justifying the investment–lease economics, term and extension rights, commencement conditions, and tenant termination and downsizing rights.

Pillsbury negotiated definitive documents to secure Client X's rights and protections.

This deal illustrates:

- How investors and data center developers can partner to expedite project development grounded on low-emission power.
- The value add of Pillsbury's cross-group, full lifecycle data centers team. The team advising consisted of corporate and securities, real estate, project finance, energy, and technology transactions attorneys.

PILLSBURY AT A GLANCE

Pillsbury is a forward-thinking law firm thriving in an environment of intense competition and change.

Entrenched in elite markets, we do cutting-edge work for market leaders, innovators and disruptors throughout the United States and globally. Whether advising entrepreneurs working out of a garage or the world's largest public and private companies, we deliver exemplary client service and results, with highly regarded lawyers and business professionals operating across a diversity of leading practices.

We are one of the world's foremost advisors to energy companies and their lenders and investors, as well as clients in the technology, financial, life sciences and digital health, and real estate and construction industries.

A VALUES-DRIVEN LAW FIRM

Opportunity, community, and belonging are bedrock values at Pillsbury. We reward high performance, collegiality, collaboration, and inclusivity, with a prized culture that sets us apart.

Pillsbury 360 is a firmwide commitment to leverage the unique qualities of individuals across our global law firm. We believe this is essential to helping clients and each other effectively pursue opportunities, address challenges and thrive in an ever-changing world. Our approach fosters innovation, encourages the development of creative solutions, and engenders strong external and internal relationships. To make this possible, we seek out and recruit the industry's top talent with varied backgrounds and experiences. We identify future leaders and invest in professional development to provide opportunities for all to reach their full potential.

GLOBAL REACH

Pillsbury operates from key global positions and across an integrated platform. More than 30% of our work crosses borders, and Pillsbury lawyers are fluent in 35 different languages. These capabilities are complemented by a bespoke network of 100 law and advisory firms in more than 70 countries, providing clients with access to one of the largest and most sophisticated legal platforms in the industry.

BY THE NUMBERS

1868	22	700+	1,300+
Year Founded	Global Offices in the U.S., Europe and Asia	Total Attorneys	Employees Worldwide

OUR OFFICES

UNITED STATES

AUSTIN
401 West 4th Street, Ste. 3200
Austin, TX 78701
+1.512.580.9600

BOSTON
10 Post Office Square, Floor 8
Boston, MA 02109
+1.857.323.5420

HOUSTON
609 Main Street, Ste. 2000
Houston, TX 77002
+1.713.276.7600

LOS ANGELES
725 S. Figueroa Street, 36th Floor
Los Angeles, CA 90017
+1.213.488.7100

MIAMI
600 Brickell Avenue, Ste. 2600
Miami, FL 33131
+1.786.913.4900

NEW YORK
31 West 52nd Street
New York, NY 10019
+1.212.858.1000

NORTHERN VIRGINIA
7900 Tysons One Place, Ste. 500
Tysons, VA 22102
+1.703.770.7900

PALM BEACH
324 Royal Palm Way, Ste. 220
Palm Beach, FL 33480
T: +1.561.232.3300

SACRAMENTO
500 Capitol Mall, Ste. 1800
Sacramento, CA 95816
+1.916.329.4700

SAN DIEGO
11682 El Camino Real, Ste. 200
San Diego, CA 92130
+1.858.509.4000

SAN FRANCISCO
Four Embarcadero Center, 22nd Floor
San Francisco, CA 94111
+1.415.983.1000

SILICON VALLEY
2400 Hanover Street
Palo Alto, CA 94304
+1.650.233.4500

WASHINGTON, DC
1200 Seventeenth Street, NW
Washington, DC 20036
+1.202.663.8000

ASIA

BEIJING
Suite 3106, China World Office 2
No. 1 Jian Guo Men Wai Avenue
Beijing 100004
+86.10.8572.1188

HONG KONG
18 Chater Road, Ste. 1704
Central, Hong Kong
+852.3959.7500

SHANGHAI
Suite 3001, 30th floor
Jing An Kerry Center Tower 2
1539 Nanjing Road W.
Shanghai 200041
+86.21.6137.7999

TAIPEI
Level 34, Taipei Nanshan Plaza
No. 100, Songren Road, Xinyi District
Taipei 100, Taiwan
+88.62.8729.5858

TOKYO
Marunouchi Eiraku Building 20th floor
4-1 Marunouchi 1-chome
Chiyoda, Tokyo 100-0005
+81.3.6268.6710

EUROPE

LONDON
Level 34
100 Bishopsgate
London, EC2N 4AG
+44.207.847.9500

MIDDLE EAST

DOHA
Office 855, Tower 2
Globex Business Center
Doha, Qatar
+974.4020.6655

RIYADH
3272 Al Imam Saud Ibn Faisal Rd.
Al Malqa
Riyadh, 13522
Saudi Arabia
+966.112.173.477

GLOBAL OPERATIONS

NASHVILLE
333 Commerce Street, Ste. 1300
Nashville, TN 37201
+1.615.622.3400