Plugged In & Parched

STRATEGIC DATA CENTER SITING IN A RESOURCE-CONSTRAINED WORLD

Artificial intelligence (AI) is revolutionizing nearly all sectors and industries by enhancing efficiency, improving decision-making, and creating new opportunities.

As AI technology continues to advance, its impact is expected to grow, bringing both challenges and opportunities that will shape the future. The data centers needed to train and run AI are growing exponentially, too. Siting them presents its own challenges, particularly as data centers compete for massive quantities of power and water across the United States.

Studies have shown that a ChatGPT query uses about 10 times more electricity than a Google search, and that an AI bot requires about 18 ounces of water to generate one 100-word email. This voracious demand has sparked a race to locate and develop sites with enough power and water for the data centers necessary to train and run AI, among other computing uses such as cloud computing and crypto mining. The resulting power and water crunch has become more acute as individual data centers grow exponentially in size. These data centers will require billions of dollars of investments and well-sited tracts of land to support their infrastructure.

AI data centers require substantial amounts of power due to several factors related to the nature of AI workloads and the infrastructure needed to support them. Training complex AI models, especially deep learning models, involves processing vast amounts of data through numerous iterations. This training requires significant computational power, often provided by specialized hardware like graphics processing units (GPUs) and tensor processing units (TPUs).

Moreover, AI models need to perform "inference," which is the process of making predictions or decisions based on data. Inference can also be computationally intensive, especially for real-time applications. The need to store and manage all that data used to train AI adds to the power needs of data centers. Moreover, data centers typically operate continuously, 24 hours a day, seven days a week, to support real-time AI applications and services. This constant operation leads to high cumulative power consumption.

Recent analysis projects that data centers could account for up to 44% of U.S. electricity load growth from 2023 to 2028. In 2023, data centers consumed 176 terawatt hours (TWh) of energy, 4.4% of total U.S. electricity usage.[1] By 2028, these data centers are projected to double or even triple their energy consumption to consume between 325 to 580 TWh, roughly 6.7% to 12% of total U.S. electricity.

Data centers also require a lot of water, primarily for cooling purposes. The high computational power of AI hardware generates significant heat.

Total Data Center Electricity Consumption (TWh)

Figure 1 - Total U.S. data center electricity use from 2014 through 2028

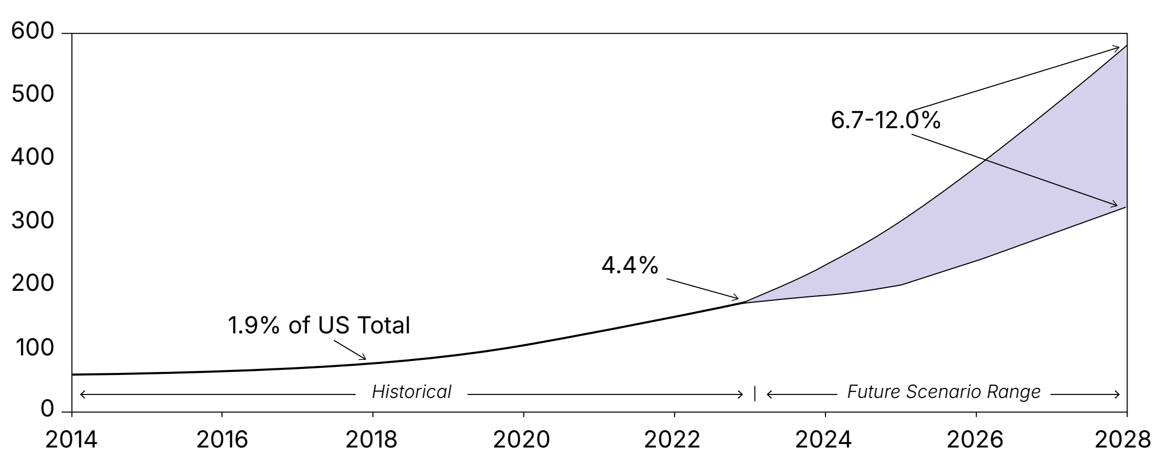

Figure 1 Source: Lawrence Berkeley Lab, 2024 United States Data Center Energy Usage Report

[1]Arman Shehabi Et Al., Lawrence Berkeley Nat'l Lab., 2024 United States Data Center Energy Usage Report 5–6 (2024).

Effective cooling systems are essential to maintain optimal operating temperatures and prevent overheating. Water-based cooling systems are effective in managing the heat generated by the high-performance computing equipment used for AI, but require significant quantities of water to function. A hyperscale data center can use north of 500,000 gallons of water a day. GPUs tend to run hotter than traditional servers, thus requiring more power and water to cool them.

In the coming decade, the United States hopes to be at the forefront of data center construction and operation, driven largely by private investments in data center infrastructure. President Trump has signaled his support for co-location of AI infrastructure and generation and expressed his intent to use emergency powers to fast-track such projects.[2] This builds off former President Biden's efforts, including Executive Order 14141, that remains in effect as of this writing, directing agencies to accelerate large-scale AI infrastructure development according to five guiding principles aimed to promote national security, economic competitiveness, and preventing increased energy costs to other consumers and businesses, among other things.

To achieve these goals, however, the United States and data center developers will need to address many critical challenges in the months and years ahead.

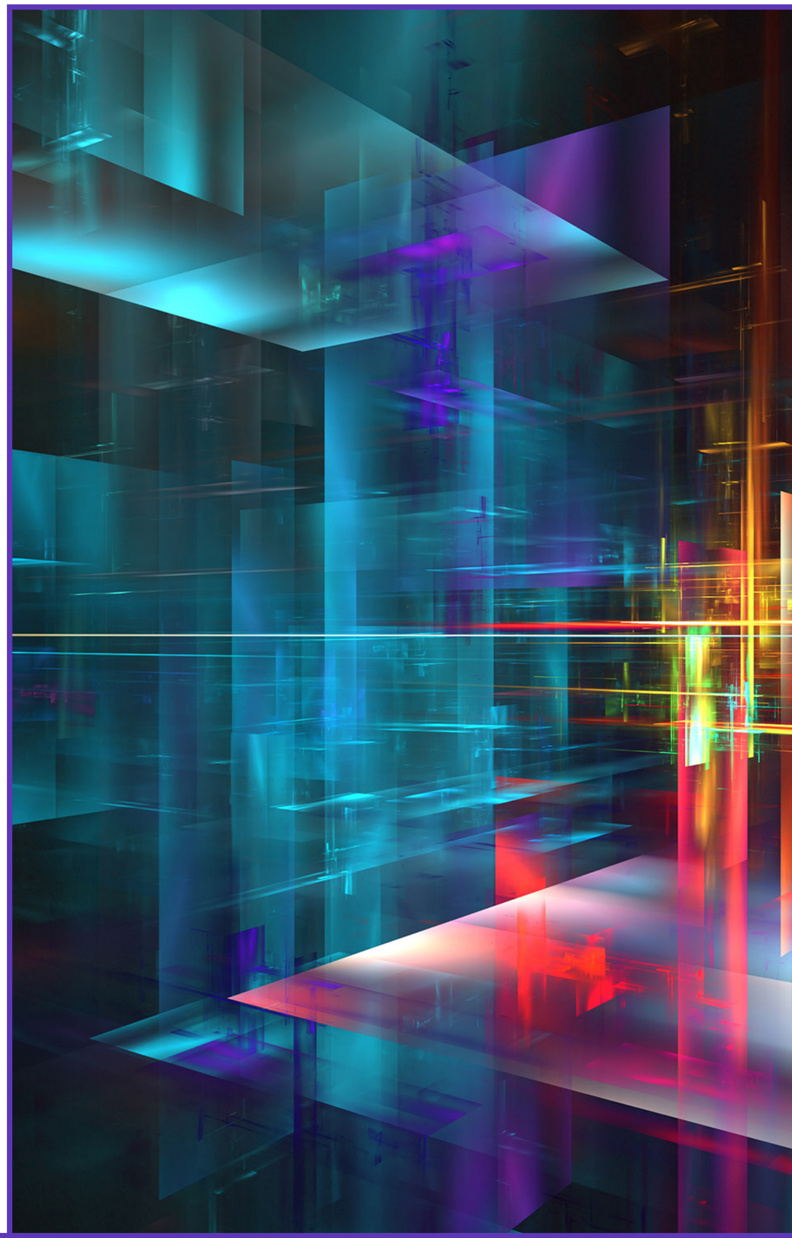

A "Hyperscale" data center can house **tens or hundreds of thousands** of servers, use well over **50,000 square feet** of floor space, and consume over **1,000 MW** of peak power – about the same as a small city.

[2]Executive Order, "Removing Barriers to American Leadership in Artificial Intelligence" (Jan. 23, 2025), available at https://www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/.

Finding the "Goldilocks" Site

Data centers, as physical central hubs for housing and managing large amounts of information, are critical in today's digital world. In particular, the rapid advancement of AI technology presents significant growth opportunities for data center real estate in the United States. As such, demand for land that is suitable for data center use has quickly risen. When selecting land for data centers, stakeholders should consider the following to ensure optimal performance, security, and scalability.

The first, and arguably most important, consideration is location. As suggested by Executive Order 14141, the ideal site will be near a reliable—and major—power source and renewable energy sources to enhance sustainability and reduce operational costs. In addition, the site should be easily accessible for construction of the data center and the appropriate workforce to run the data center. For water-cooled data centers, access to a reliable water source is key. In addition, the ideal site would be near existing fiber optic networks and telecommunications infrastructure to optimize connectivity. In terms of size and scalability, the site must be large enough to accommodate the data center, all related infrastructure, and potential future expansion.

The site must also comply with all applicable environmental regulations and the site's zoning must permit the construction and operation of a data center. One challenge in data center development is the availability of sites that meet all the criteria of the developer and operator.

A common theme when selecting and entitling data center sites is the effect that data centers have, or are perceived to have, on the environment and local communities surrounding the data centers. On the one hand, data centers can drive economic development by creating jobs for the local community and providing additional tax base, but on the other hand, the local communities affected by the data center may be racked with concerns about the environmental and social impacts. Former President Biden was particularly concerned about AI infrastructure resulting in increases in energy costs for consumers and businesses and noted in Executive Order 14141 that "the development of AI infrastructure should benefit those working to build it." Just as community members grapple with the opportunities and challenges that data centers bring, we weigh how those unique opportunities and challenges affect the search for the "perfect" data center site herein.

Rise of Buyer-Side Brokers and Local Developers

In our experience, it is more likely that data centers are constructed in secondary and tertiary real estate markets, which are popular for a variety of reasons. One, there are plenty of large, vacant parcels, or agricultural land, available at reasonable prices. Two, construction costs are lower, both in terms of construction materials and labor costs. Third, it is often easier to entitle a data center in a smaller municipality without competing sophisticated projects lining up for approvals. It is not uncommon for data center developers to target remote sites, particularly agricultural land that is already cleared, as it has a low property tax rate.

The owners and sellers of agricultural or vacant land often turn to local brokers who specialize in real estate transactions involving agricultural or vacant land. In recent years, buyer-side brokers have developed data center specializations, and extensive networks and relationships with others in the data center industry. A well-informed buyer-side broker may be helpful, particularly in educating sellers unfamiliar with data center buyers. Understanding the unique requirements associated with data center leases and sales is instrumental in site selection and transaction negotiations.

There are well **over 5,000 data centers in the United States**, far more than any other country.

In addition to brokers, some data center developers are thinking outside the (big) box, and they are working with local developers, who have sophisticated teams that can assist with site selection, entitlements, and site preparation. These local developers, who know both the complexities of building data centers as well as the local community, can save valuable time for data center operators. Often, local developers can secure site control of many parcels necessary for a large data center and related interconnection without the name of the data center developer being used. This acquisition strategy usually allows the local developer to acquire the land at a lower cost and then sell the bundled package to the data center developer at a premium. There must be a great deal of trust in the local developer and its counsel, as the data center developer will have to stand by during public hearings and discussions with the municipality, as the local developer plays its role as the face of the project at the early stages.

Next Frontier? Siting on Public Lands

As the hunt for data center sites heats up, there is growing interest in exploring federal lands for data center development, and in particular for siting large generating facilities necessary to power them. In the past year, recommendations for siting data centers on federal lands have accelerated, including most recently in Executive Order 14141. Executive Order 14141 directed the U.S. Department of Defense (DOD) and the U.S. Department of Energy (DOE) to select sites on which to construct large-scale data centers based on accessibility to high-capacity transmission infrastructure and ability to minimize effects on communities, natural resources, and the environment. It further directed the U.S. Department of the Interior (DOI) to identify priority sites suitable for clean energy that can support data centers on DOE and DOD sites, with a focus on priority zones suitable for geothermal power and thermal storage. The breadth of this order and requirements surrounding identification of the priority sites, permitting and environmental review, and interconnection are indicative of the multiple legal and compliance challenges that developers building and powering data centers will face.

Constructing data centers and renewable energy facilities on DOD or DOE lands or constructing supporting energy facilities on lands managed by DOI could provide multiple benefits, such as increasing security and potentially accelerating development approvals. Further, constructing data centers on public lands could be more cost efficient, especially if the federal government is the user of that data center, as the cost of land is no longer an issue and property tax payments are no longer a factor.

The federal government manages large swaths of land in the United States, particularly in the West, but pursuing sites on federal lands involves navigating a dizzying list of statutes that govern different federal agencies' ability to transfer or "dispose" of property by sale or by lease. The laws governing management of those lands and leasing for infrastructure projects vary widely.

Case Study

In July 2024, DOE entered into negotiations with Hecate Energy, LLC to construct a 1 gigawatt (GW) solar project on DOE-owned land. The site in Washington state was previously established as part of the Manhattan Project in 1943 and is subject to DOE's authority under the AEA. The negotiations were part of the broader Cleanup to Clean Energy initiative, which aimed to repurpose parts of DOE-owned lands to support clean energy. As this initiative arose under the auspices of an Executive Order that has since been rescinded, the future of the project is unclear; however, it remains a good example of the kind of public-private partnership that could facilitate data center development.

Western Federal Lands Managed by Five Agencies

Figure 2

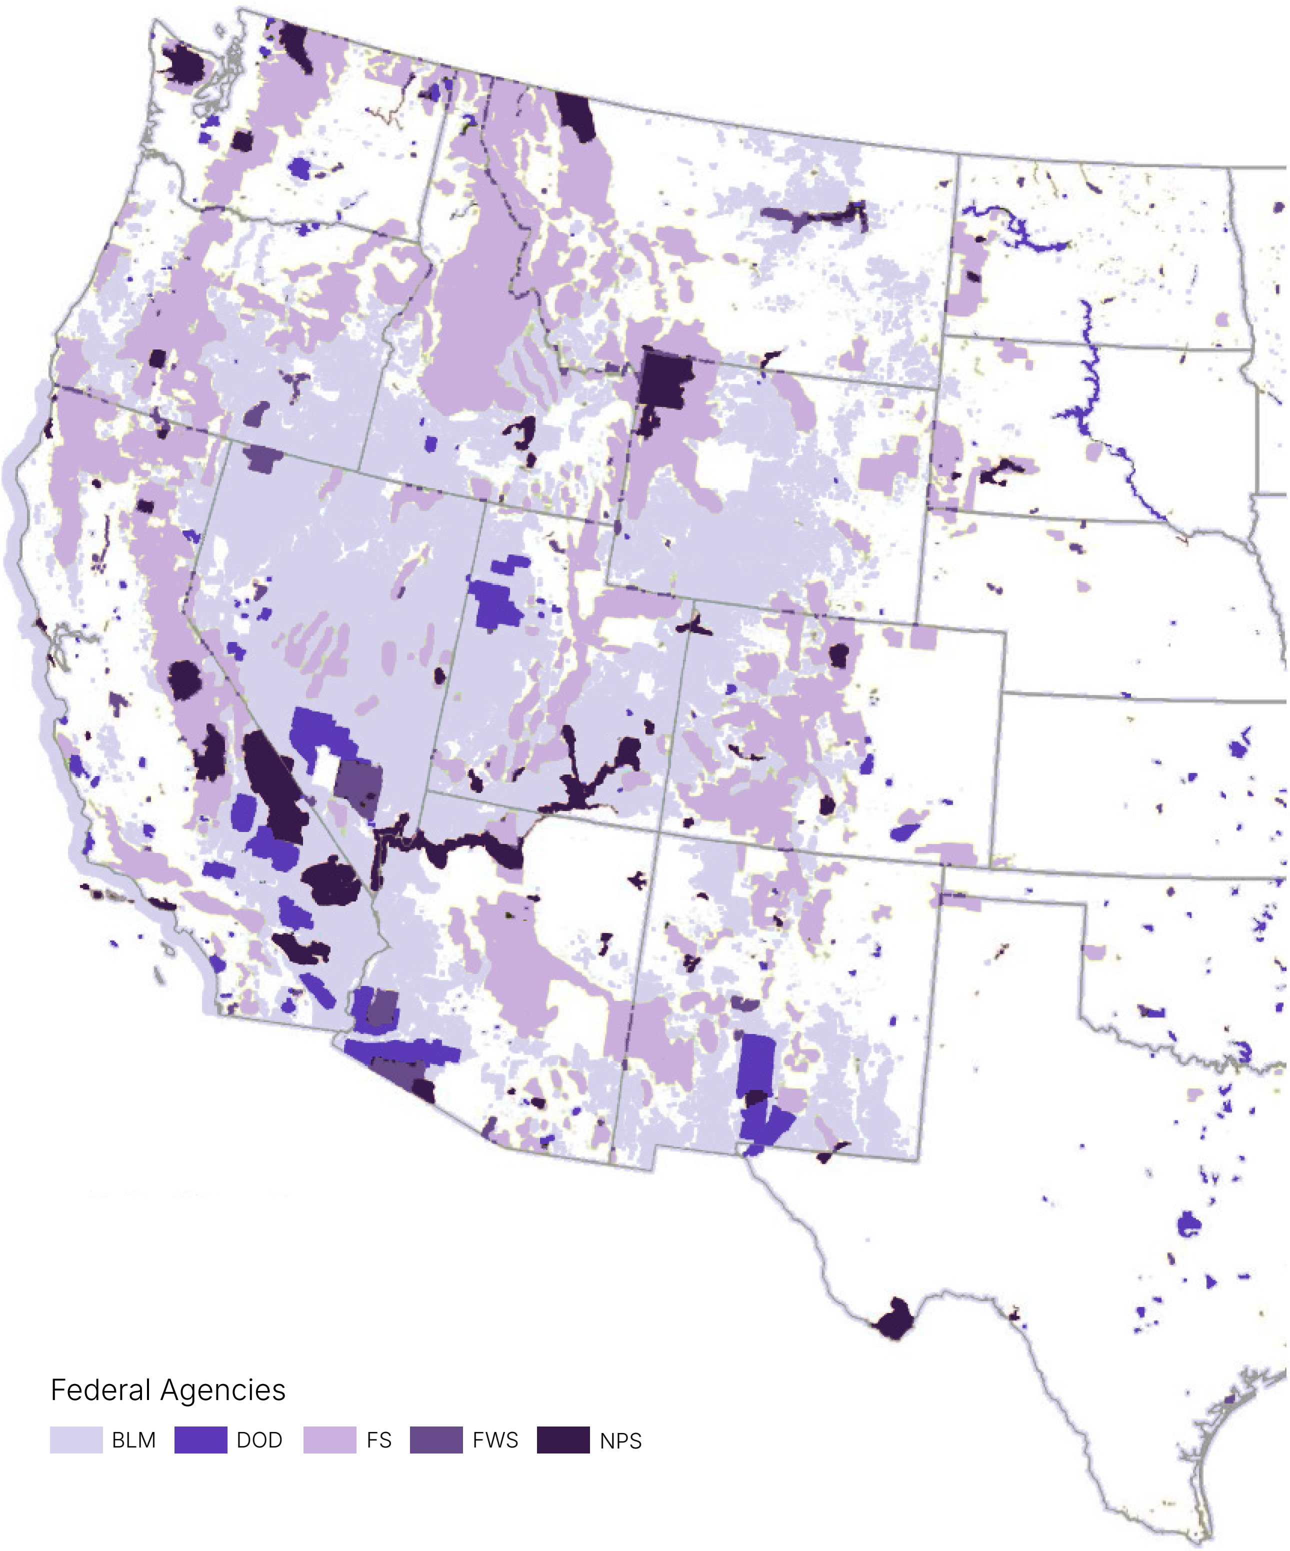

Figure 2 Source: Federal Land Ownership: Overview and Data, R42346 (Feb. 21, 2020), Map boundaries and information generated by CRS using federal lands GIS data from the National Atlas, 2005, and an ESRI USA Base Map. BLM refers to Bureau of Land Management. FS refers to Forest Service. FWS refers to Fish & Wildlife Service. NPS refers to National Park Service.

Eastern Federal Lands Managed by Five Agencies

Figure 3

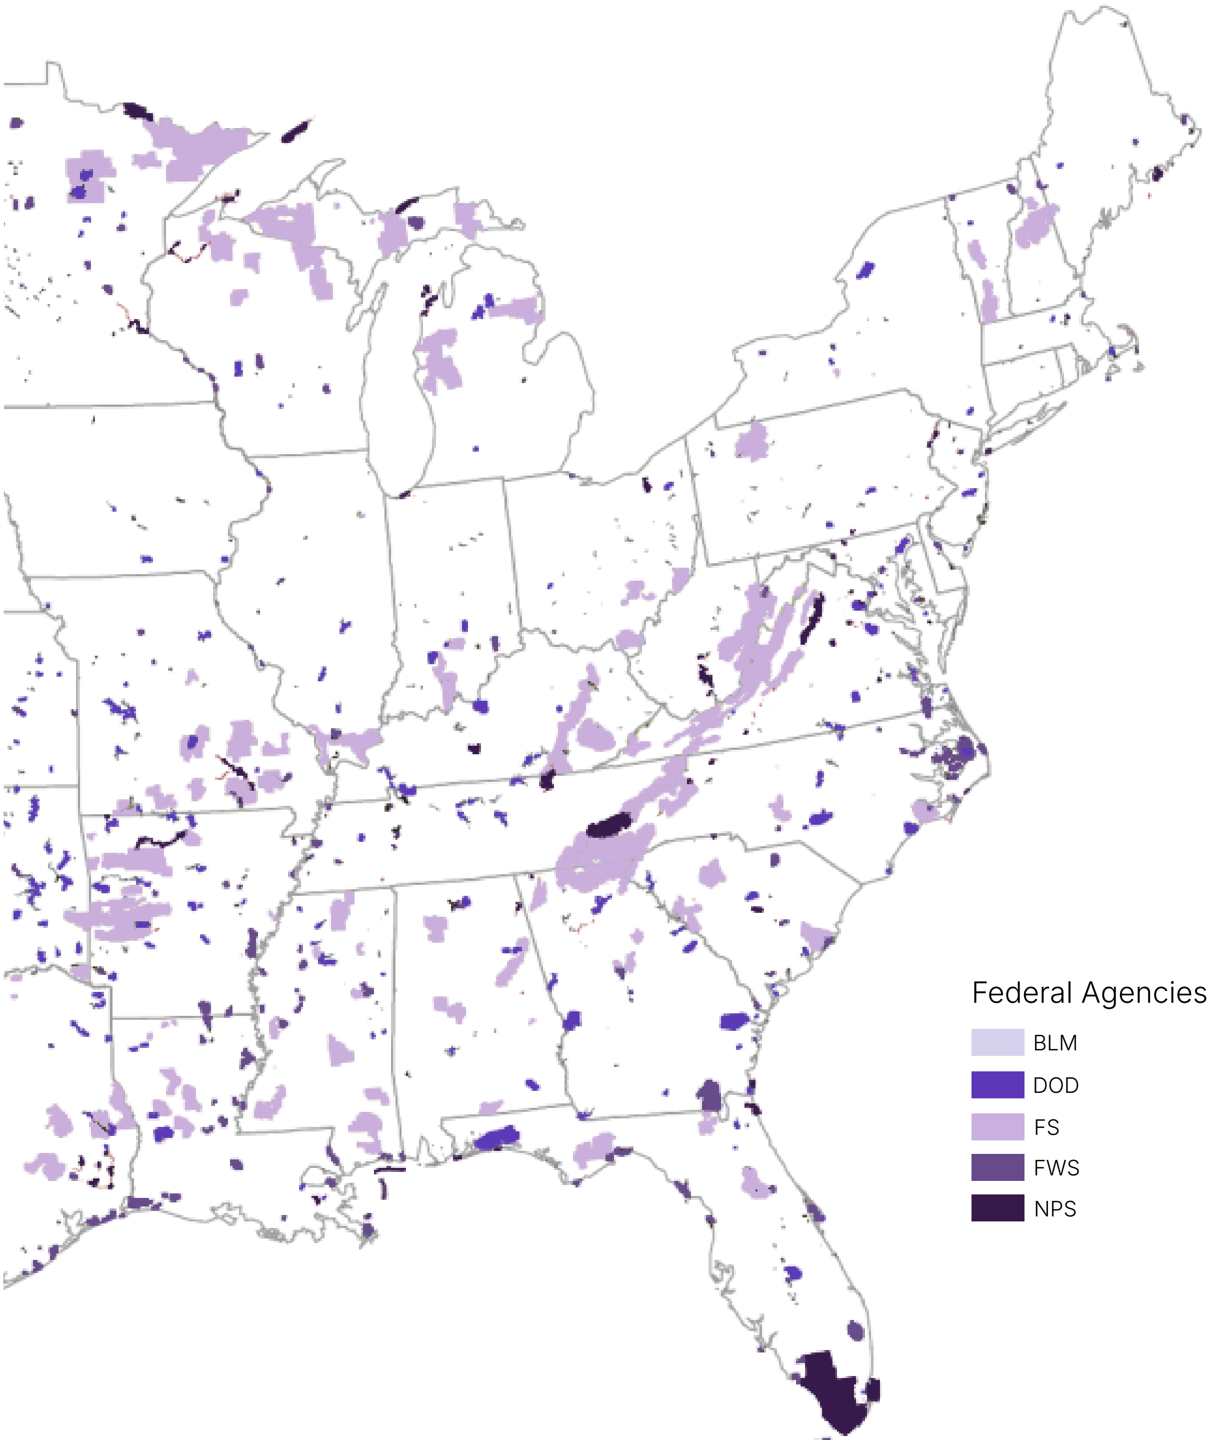

Figure 3 Source: Federal Land Ownership: Overview and Data, R42346 (Feb. 21, 2020), Map boundaries and information generated by CRS using federal lands GIS data from the National Atlas, 2005, and an ESRI USA Base Map

Federal Lands in Alaska and Hawaii Managed by Five Agencies

Figure 4

Figure 4 Source: Federal Land Ownership: Overview and Data, R42346 (Feb. 21, 2020), Map boundaries and information generated by CRS using federal lands GIS data from the National Atlas, 2005, and an ESRI USA Base Map

The DOD, DOE, and DOI manage lands under a variety of statutory authorities and implementing regulations, which determine the type of infrastructure that can be sited on those lands and what type of lease or transfer authority can be used. The primary authorities for DOE to transfer property, by sale or lease, include, among others, the Atomic Energy Act (AEA), the DOE Organization Act, and Section 3158 of the National Defense Authorization Act for Fiscal Year 1998. DOE, as the successor to the Atomic Energy Commission, has authority to lease property originally acquired in connection with AEA objectives, which include the promotion of nuclear energy for peaceful purposes while ensuring public health and safety, national security, and environmental protection.[3] The DOE Organization Act authorizes DOE to transfer certain real property that is "excess" to public or private entities.[4] Such a lease is limited to a term of 10 years or less, with the option to renew where such a renewal would promote national security or the public interest.[5] In light of the national security and public interest considerations involved in maintaining the United States' dominance in AI, this path may show promise for siting data centers.

DOE may also transfer, by sale or lease, real property at defense nuclear facilities for economic development purposes.[6] Some DOE legacy-managed sites and assets that no longer serve a DOE mission after remediation have been used for "beneficial reuse," including for energy-related purposes or commercial or industrial use. These could present another opportunity for data centers or generation.

Leases of DOD non-excess real property are authorized pursuant to 10 U.S.C. § 2667. Leases under § 2667 are typically for a term less than five years, unless the Secretary determines a longer term is in the public interest, which could present a pathway for data centers. Complex ground

[3]See 42 U.S.C. § 2011.
[4]42 U.S.C. § 7256(c), (e).
[5]42 U.S.C. § 7256(d).
[6]50 U.S.C. § 2811(a)(1); see also 10 C.F.R. § 770.2.

leases are typically issued under the Enhanced Use Lease Program (EUL), which provides a long-term use lease (up to 55 years) of property to a private developer in exchange for cash or in-kind services. This process has been used for both energy and commercial redevelopment. DOD is further provided the authority under 10 U.S.C. § 2668 to grant easements for rights-of-way over, in, and upon public lands under the Secretary's control, provided that the casements will not be against the public interest. Such easements could be used for transmission lines, substations, and pumping stations, among other required associated infrastructure.

Real property disposition decisions for both DOD and DOE will be based, in large part, on the availability of the property following evaluations of mission need, environmental conditions/status, potential environmental impacts, and the interests of the local community.

Approximately 16 million acres of DOD-managed land is public land withdrawn for military use. The Act of February 28, 1958 (43 U.S.C. 155-158), sometimes referred to as the Engle Act, places on the Secretary of the Interior the responsibility to process DOD applications for national defense withdrawals, reservations or restrictions aggregating 5,000 acres or more for any one project or facility. These withdrawals, reservations, or restrictions may only be made by an act of Congress, except in time of war or national emergency declared by the President. The analysis of the purposes of a withdrawal and areas of continuing DOI jurisdiction is a facility-specific issue that requires study of the purposes and limitations of the withdrawal.

Multiple laws govern the leasing of public lands by the DOI for the exploration and production of oil and gas and for renewable energy projects, including geothermal, wind, and solar energy,

Case Study

In February 2023, a new solar and storage facility was built by Terra Gen LLC on Edwards Air Force Base (Edwards) in California under the Air Force (EUL) program. The EUL area is composed of up to 4,000 acres of non-excess, underutilized property at Edwards and surrounding private property.[7] The solar and storage facility, which has almost two million solar panels installed, can supply up to 1,300 megawatts (MW) of power to the California Independent System Operator grid and is one of the largest solar and storage projects in North America. This project is estimated to power over 238,000 homes, displacing more than 320,000 tons of CO2 emissions annually. Moreover, the Air Force estimates the project could yield cash rent consideration of over $75.8 million throughout the expected 35-year lease.

Figure 5 - Edwards Sanborn

Figure 5 https://terra-gen.com/edwards-sanborn/
[7]412th Test Wing Public Affairs, Largest Private-Public Collaboration in Department of Defense History Reflects Commitment to Clean Energy, Edwards Air Force Base (Feb. 6, 2023), https://www.edwards.af.mil/News/Article/3289170/largest-private-public-collaboration-in-department-of-defense-history-reflects/.

along with those governing the permits that a lessee must obtain. The Bureau of Land Management (BLM) manages public lands under the Federal Land Policy and Management Act (FLPMA), and issues rights-of-way for transmission and renewable energy generation. BLM also leases public lands for oil and gas under the Mineral Leasing Act of 1920 and for geothermal energy pursuant to the Geothermal Steam Act of 1970. Each of those statutory authorities are implemented through multiple regulatory authorities and BLM directives. Data center developers seeking to use enhanced geothermal resources to power their sites may focus on opportunities with the DOI, particularly given the recent efforts to streamline permitting for developing that resource.

The permitting and environmental review process for the construction of infrastructure projects on federal lands is also multilayered and technology- and project-specific and governed by a litany of statutory and regulatory requirements. For example, at the federal level, projects on federal lands will require: compliance with the National Environmental Policy Act (NEPA); compliance with Section 106 of the National Historic Preservation Act; tribal government-to-government

Case Study

The Falcon Hill Project Area is a public-private partnership between Sunset Ridge Development Partners (SRDP) and the Military Installation Development Authority (MIDA), an "authority" created by the Utah Legislature in 2007 to facilitate the development of military land in Utah. MIDA is an innovative nonprofit that has authority to act as the taxing and land use authority, and it serves as the trustee of the Air Force's payment-in-kind account for the Falcon Hill Aerospace Research Park (Falcon Hill). While Falcon Hill crosses multiple jurisdictions, MIDA serves as a single authority over the project, which helps streamline the permitting process and increases efficiency for SRDP when securing zoning and building permits and other necessary approvals. Thus far, there is over 1.3 million square feet of commercial development at Falcon Hill. Falcon Hill is the largest Air Force EUL in the country.[8]

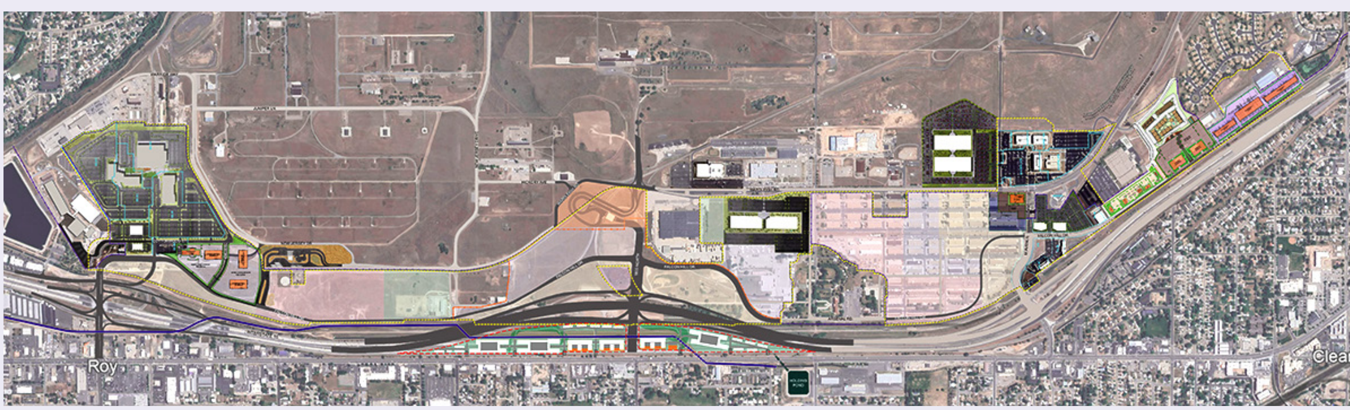

Figure 6 - Falcon Hill Masterplan

Figure 6 https://www.falconhill.com/leasing#ExistingDevelopment
[8]Falcon Hill National Aerospace Research Park, Project Partners: Military Installation Development Authority, https://www.falconhill.com/project-partners.

consultation; consultations with the U.S. Fish and Wildlife Service under the Endangered Species Act and the Migratory Bird Treaty Act; permits issued by the Environmental Protection Agency under the Clean Air Act; and permits issued by the Army Corps under the Clean Water Act and the Rivers and Harbors Act. Another potential hurdle depending on transaction structure is that transfers of property from the federal government come with a major caveat: Section 120(h)(3)(A) of CERCLA requires that a federal agency transferring real property to a nonfederal entity include a covenant in the deed of transfer warranting that all remedial action necessary to protect human health and the environment has been taken prior to the date of transfer with respect to any hazardous substances remaining on the property. Depending on the uses to which the federal property has been put in the past, remediation can be a significant undertaking.

Challenges surrounding development of data centers on federal lands include, among others, consistency with land management plans; impacts to environmental, cultural, and water resources; historic properties; noise; remediation (particularly for DOE properties); impacts to local communities; interconnection and grid infrastructure; and differing timelines between developing and completing a data center project and a renewable energy generation project. It remains to be seen how the high-energy demand for data centers can be met with clean energy projects on federal lands, given the complexity of the leasing and

In the last year, the BLM has taken a number of actions to **streamline permitting for geothermal energy.** Not only did it finalize a new categorical exclusion to accelerate the discovery of new geothermal resources on public lands, it also proposed a separate categorical exclusion that will include activities related to the search for indirect evidence of geothermal resources. Moreover, the **BLM has adopted additional categorical exclusions from the U.S. Department of the Navy and the U.S. Forest Service** for geothermal exploration operations. 89 Fed. Reg. 28797 (Apr. 19, 2024).]

permitting processes and the length of time needed to plan, construct, and commence operations of these facilities.

Executive actions to expedite permitting and streamline environmental reviews have been underway over the last decade and will certainly continue to be a priority under the second Trump administration, at least for certain types of energy. Permitting reform also has been a constant refrain in Congress, including for multiple types of energy and transmission, as set forth in the Energy Permitting Reform Act of 2024, introduced in the last Congress.

However, when balanced against protection of the environment, water resources, cultural resources, and the interests of multiple stakeholders, including surrounding communities and tribes, developing on federal lands may be more difficult and time-consuming than necessary for the rapid growth of data centers and their power sources. What effect will executive actions and statutory and regulatory changes, including judicial decisions surrounding NEPA review, have upon the desire to site AI infrastructure on federal lands? To what extent will the political winds surrounding changing priorities for energy resources that would power data centers affect investor decisions and decisions surrounding the grid? What risk, if any, do the initiatives in Executive Order 14141 and its relatively nondescript and ambiguous directive to "minimize adverse effects" pose to the preservation of cultural and natural resources for generations to come? These are all challenges data center developers must navigate.

Building a Reliable Power Supply

Data centers have increased the nation's need for reliable energy and capacity, and this need is expected to increase exponentially for the foreseeable future. This demand will not be met in a vacuum—data center owners, developers, and other actors in data center supply chains will have to contend with a myriad of competing interests throughout the energy industry. Co-location of generating resources and data centers is one promising solution, but co-location by itself will not erase factors such as reliability risks, administrative burdens, and concerns of other stakeholders.

Meeting data center power needs requires the reconciliation of multiple competing interests. For instance, data center owners often seek carbon-free energy solutions to mitigate their data centers' climate impact and require a reliable flow of power to be installed quickly. State regulators considering new generation and transmission applications must balance the interests of utilities seeking to construct or acquire these facilities with the interests of the new data center customers, the utilities' existing customer base, and other relevant stakeholders. The risk of constructing generation or transmission that becomes a stranded asset while balancing the needs for all customers to

pay their "fair share" can often lead to slow-downs or even paralysis. Regional grid operators face similar concerns, as do federal regulators who are concerned with grid reliability and national security risks.

How Fast Can New Generation Be Built and New Load Connected?

Against the backdrop of backlogged interconnection queues for both generation and new loads, rapidly increasing data center load poses acute physical feasibility and administrative coordination issues for grid operators. Data centers can ramp up power demand quickly as servers are added to the site. But generating facilities and other grid infrastructure like transmission lines require years, if not decades to develop and construct. A primary bottleneck lies in the queues to connect new generation and load to the grid. A recent study estimated that 2,600 GW of largely renewable generation and storage capacity, more than twice the total installed capacity of the existing U.S. power plant fleet, were in generator interconnection queues at the end of 2023.[9] The queues to connect new large loads to the grid are equally bloated, and in a growing number of regions, data centers make up a significant share of the new loads trying to connect. In some regions, data center service requests are facing interconnection wait times as long as seven years. The Federal Energy Regulatory Commission (FERC) approved reforms to speed the generator interconnection process last year, including a first-ready, first-served cluster study process, faster interconnection queue processing, and leveraging technological advancements in the interconnection process, but it remains to be seen how and whether these reforms will suffice to address grid operators' massive backlogs. Moreover, load connection queues are primarily regulated at the state level, making a national solution nearly impossible.

There also remains uncertainty about how much energy data centers will need, and how long they will need it. Recent AI advancements abroad, such as DeepSeek AI, appear to use significantly less energy and have led the industry to question whether data centers will continue to need energy at projected scales over the useful life of

As of August 2024, **59 of the 80** GW of Oncor's service requests from large commercial and industrial customers in Texas **came from data centers, and data center requests** are rising year-over-year.

[9]Berkeley Lab Energy Mkts., & Pol'y, *Grid Connection Backlog Grows by 30% in 2023, Dominated by Requests for Solar, Wind, and Energy Storage* (April 10, 2024), https://emp.lbl.gov/news/grid-connection-backlog-grows-30-2023-dominated-requests-solar-wind-and-energy-storage.

grid resources built to serve them. Utilities that overbuild generating resources and other grid infrastructure to meet data center demand risk findings of imprudence or overcharging other ratepayers if data center demand decreases.

State integrated resource planning processes attempt to deal with these competing interests by requiring utilities to define their energy and capacity needs over time along with their plans to install or acquire resources to meet those needs. However, these proceedings are often lengthy and contentious, and information can be deemed outdated before the process is even complete resulting in delays or dissatisfaction among stakeholders. Sophisticated navigation of these proceedings, including consideration of novel alternatives to the process, is required to ensure that they do not impede the pace of necessary facilities.

Who Pays for Building Out the Grid?

Building transmission, distribution, and generation infrastructure is incredibly expensive. This means the question of who pays for that infrastructure is of critical importance. The question is complicated by the way the grid is regulated in the United States—in the mainland United States, outside Electric Reliability Council of Texas (ERCOT), transmission service and wholesale energy markets are primarily regulated by the federal government, while distribution service and retail supply are primarily regulated by the individual states.

At the state level, utilities and their regulators have taken different approaches to charging data center owners for the energy they use while accommodating data center owners' climate goals. For example, the Public Utilities Commission of Nevada is currently considering a Clean Transition Tariff (CTT) created by NV Energy in collaboration with Google, which allows Google's data centers to draw from the grid while also paying the premium for intermittent renewable resources to dispatch into the same grid.[10] This arrangement could streamline Google's procurement of renewable energy on an hour-for-hour basis without changes in state law that would otherwise prevent Google from purchasing electricity from renewable energy generators directly. Duke Energy has similar tariffs waiting for North Carolina's and South Carolina's respective state utility regulators' approval.

Other utilities and regulators have taken a less accommodating approach to data center owners. Last May, American Electric Power Ohio (AEP Ohio) proposed a tariff to the Public Utilities Commission of Ohio (PUCO) that would have required data center customers to enter 10-year electric service contract commitments with minimum demand charges based on 90-95% of their contract capacity. Several technology companies, in turn, filed a settlement agreement with PUCO that would require large-scale data center customers to pay for up to 85% of their projected demand, even if they don't use all of it. This restriction is meant to cover infrastructure costs associated with serving data centers' loads.

At the federal level, rapid data center development—frequently concentrated in regions that historically were not major load centers—raises acute issues related to transmission planning and necessary investment signals to procure the generating capacity needed to serve

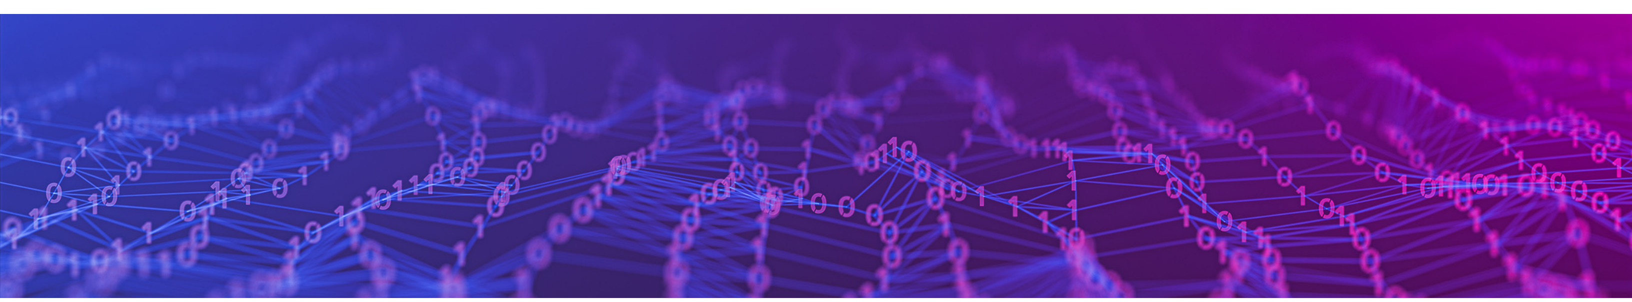

[10] https://blog.google/outreach-initiatives/sustainability/google-clean-energy-partnership/.

these new loads. It can take years, and sometimes decades, to build new energy infrastructure. Grid operators and utilities use load projections as one input into the transmission planning process to forecast what transmission upgrades the grid will need to accommodate expected future load. Likewise, load forecasts underpin grid operator decisions about how much generating capacity to procure, either through FERC-regulated capacity markets or through state resource planning processes. Rapid load growth that outstrips the pace of energy infrastructure development—as we are seeing with data center development—severely stresses the ability of the grid to provide reliable service. Moreover, projections of future data center load can be volatile and uncertain over the time frames necessary to build energy infrastructure—real dollars have to be spent on energy infrastructure before grid operators have certainty about the load materializing. This raises challenges in determining appropriate allocation of those costs among customers that use the transmission network.

Some data center owners have proposed several load forecasting policies to address these challenges. For example, Google recently proposed that data center load could be included in load forecasts used for transmission planning and capacity procurement when it is "commitment-backed" by "material, up-front financial commitments" to mitigate and appropriately allocate the risk that transmission and network upgrades will be overbuilt.[11]

Co-Locating Data Centers and Generating Resources

Another solution currently trending is co-location of generation and data centers, rather than relying solely on grid supply developed by utility providers. Co-location offers data center owners a path toward faster operation and carbon reduction goals. But co-location can pose reliability trade-offs for data center owners and system integrity concerns for grid operators, regulators, and other energy customers. For example, co-locating data center load with *existing* generating facilities draws away generating capacity otherwise available to the broader market. In markets where generating capacity is scarce, pulling existing capacity off the grid has the potential to raise capacity prices across the entire market.

On the other hand, co-location of *new* data centers with new generating facilities can enhance the reliability of the grid by netting out the individual impacts of the generation and load on the grid. Co-location with a new generating resource also helps to demonstrate that the data center operator is shouldering at least some of the costs incurred to serve its load. However, existing interconnection practices are rarely optimized to assess grid impacts from load and generation together. Nor are they typically structured to ensure that customers receive the benefit of netting co-located generation and load where both are connected directly to the grid at the same location. In the ERCOT market, the grid operator's metering rules permit the creation of a Private Use Network (PUN), which allows ERCOT to net the generation and load at a shared point of interconnection with the grid. This structure represents significant savings for customers using the PUN, as it avoids grid charges for load that is served by the co-located generator. However, structures like the PUN are not available in all jurisdictions. Moreover, the structure may raise anew the question of who pays for grid infrastructure of communal benefit, particularly where grid infrastructure is typically paid for through network charges based on quantity of demand served from the grid.

Co-location also raises esoteric and novel regulatory questions. Since the regulation of the grid can vary substantially state by state, a structure that works in one state like Texas may not work elsewhere. For example, it is often

[11] Post-Technical Conference Comments of Google LLC at 7, *Large Loads Co-Located at Generating Facilities* (No. AD24-11-000), https://elibrary.ferc.gov/eLibrary/filelist?accession_number=20241210-5090&optimized=false.

technologically and economically more efficient for the co-located generation and load to share undivided ownership interests in interconnection facilities. Indeed, in some jurisdictions, shared ownership may be important to ensuring that behind-the-meter generating resources can serve loads without running afoul of state utility laws. However, sharing ownership of interconnection facilities in regions regulated by FERC can lead to uncertainty regarding the regulatory status of the data center owner. Generator interconnection facilities are considered to be transmission facilities by FERC, and owners of such facilities are deemed to be public utilities regulated under the Federal Power Act (FPA). While generators can minimize the regulatory burdens of being a public utility under the FPA by taking steps such as applying to FERC for market-based rate authority and self-certifying exempt wholesale generator (EWG) status, those routes may not be readily available to load-owning entities that also own an interest in the generator interconnection facilities in order to draw power from the co-located generator and grid. These and other regulatory considerations remain to be resolved as developers explore new ways to co-locate loads with generating resources.

Case Study

In November 2024, FERC rejected an amended interconnection service agreement (ISA) between PJM, Susquehanna Nuclear, LLC, and the interconnected transmission owner. *In re PJM Interconnection, L.L.C.*, 189 FERC ¶ 61,078 at P 1 (2024). The ISA contemplated a co-location configuration between an existing nuclear generating facility and a data center. The amended ISA sought to increase the amount of co-located load from 300 MW to 480 MW, and contained other bespoke terms intended to support reliable system operations and clarify expectations of how the data center would operate, especially when the nuclear plant was offline. FERC rejected the ISA for PJM's failure to adequately support that the bespoke terms were just and reasonable deviations from PJM's standard form ISA.

As a result, for now, data center co-location is in a federal regulatory twilight zone: On the one hand, data center co-location configurations are becoming too popular for FERC to make one-time exceptions. On the other hand, FERC has not caught up with the break-neck speed of data center development to accommodate co-location in grid operators with generally applicable rules reflected in grid operators' *pro forma* agreements. FERC has since held a technical conference regarding the connection of co-located loads and sought comment from the industry on ways to ensure reliable connection of such loads. At the time of writing, FERC appears poised to initiate a proceeding to further develop its policies on these issues.

Finding a Reliable Water Supply

Water supply issues associated with cooling data centers have become increasingly critical as the demand for data centers continues to grow, driven by such factors as AI and crypto mining. The main concerns relate to (1) high water consumption, (2) stress on local water supplies, (3) the impact of climate change, (4) regulatory and community pressures, (5) sustainability challenges, and (6) operational risks.

1. **High water consumption.** Many data centers use water-based evaporative cooling systems because they are energy-efficient, as compared to traditional air cooling. However, these systems consume significant amounts of water. For example, a single large data center can use millions of gallons of water per day, which is comparable to the water used by a small city.

2. **Stress on local water supplies.** Data centers are often located in regions with low energy costs or favorable tax incentives, which may not align with areas with abundant water resources. Arizona is one such example. In these water-scarce regions, data center cooling demands may compete with agricultural, industrial, and municipal water needs.

3. **Impact of climate change.** Higher ambient temperatures increase cooling demands, which lead to greater water usage. Areas are also at a greater risk for water shortages, which exacerbates supply challenges.

4. **Regulatory and community pressures.** Increasing scrutiny from governments and communities regarding water usage may limit data center operations or require more sustainable practices. Heavy water use by data centers can also lead to tensions with local communities, who may perceive corporate needs as outweighing public needs.

5. **Sustainability challenges.** Some data centers attempt to mitigate water use by recycling or using gray water, but implementing these systems can be costly and logistically challenging. Liquid immersion cooling systems can also reduce water usage, but they may require specialized infrastructure.

6. **Operational risks.** Depending on local water supplies can make data centers vulnerable to interruptions due to drought, infrastructure failures, or regulatory restrictions. Water scarcity and associated costs can also impact operational budgets and long-term planning.

Mitigation strategies can include transitioning to closed-loop cooling systems or liquid immersion cooling to minimize water usage; leveraging renewable resources by using nonpotable water (*e.g.*, gray water or harvested rainwater); prioritizing locations with abundant and sustainable water resources; and improving monitoring by implementing real-time water usage monitoring systems.

Balancing the need for efficient cooling with water resource sustainability remains a critical challenge for data center operators, especially as demand for these centers continues to rise.

Authors

Kira Loehr
PARTNER
Environment, Energy & Resources
KLoehr@perkinscoie.com

Camarin Madigan
PARTNER
Real Estate & Land Use
CMadigan@perkinscoie.com

Laura Morton
PARTNER
Environment, Energy & Resources
LMorton@perkinscoie.com

Jane Rueger
PARTNER
Environment, Energy & Resources
JRueger@perkinscoie.com

Valerie Johnston
COUNSEL
Real Estate & Land Use
VJohnston@perkinscoie.com

Margaret McCallister
ASSOCIATE
Environment, Energy & Resources
MMcCallister@perkinscoie.com

Jacob Neeley
ASSOCIATE
Environment, Energy & Resources
JNeeley@perkinscoie.com

perkinscoie.com/DataCenters