



Beyond available power: the new logic of data center siting

[Communications](#) / Article

July 27, 2026

A practical framework for choosing data-center locations in the age of power, water, and permitting constraints

Data-center siting is no longer just a search for land and power. It is becoming a more complex, multidimensional investment decision: matching the requirements of a workload to the constraints and advantages of a location.

The core question is changing: where can this workload remain valuable over the long term rather than be hindered by grid saturation, water scarcity, network limitations, climate risk, public opposition, or sovereignty constraints?

That question matters across the value chain, from operators and hyperscalers to developers and investors, because a data center is a long-duration claim on local infrastructure. If a site cannot expand, refinance, or adapt to future customer requirements, the investment case weakens even if the first phase gets built.

The numbers explain why the issue has moved from real estate hunt to capital allocation run. The International Energy Agency estimates that [data centers consumed roughly 485 TWh of electricity in 2025 and projects about 950 TWh by 2030](#). In the United States, Lawrence Berkeley National Laboratory's 2025 update [estimates a 2030 reference case of 649 TWh, with a range of 521–843 TWh](#). This is no longer a niche real estate segment. It is a load-growth category large enough to reshape local grids.


Capital is following that load. JLL’s 2026 Global Data Center Outlook expects [nearly 100 GW of new data-center capacity globally between 2026 and 2030](#), requiring up to \$3 trillion across real estate, fit-out, and IT infrastructure.

Location is becoming a strategic choice, not just a real estate input. The investment question is whether capital is being deployed into locations that can still absorb, connect, cool, permit, and monetize the load in the 2030s and beyond. Of course, there is no perfect data-center location; the goal is to identify the most resilient trade-off between demand and constraints.

How should companies approach this? We recommend a three-step process based on Kearney’s Data Center Site Selection Framework (see figure 1).

Figure 1
Kearney recommends a three-step process for data-center siting Data-center site selection

From workload characterization to weighted site suitability



Source: Kearney analysis

First comes the gate: does the target workload belong in this location, and is demand deep enough to support the investment case? Second, apply Kearney’s six data-center siting lenses to test whether the location can support that demand under real-world constraints. Third, score the site against those lenses to determine whether it can support a scalable, sustainable, long-term investment.

Step 1: Start with demand—not every workload needs the same location

One of the most important changes in siting is that demand can no longer be treated as a generic market overlay. Different workloads create different requirements, and those requirements should shape the site search before teams start scoring land parcels.


A practical segmentation can start with six AI-era data-center siting archetypes. The point is not to create a rigid taxonomy. It is to make sure the first question answers both commercial and technical fit: what job does this location need to do (see figure 2)?

Figure 2

There are 6 data-center siting archetypes in the AI era

Framework

A diagnostic framework for evaluating site suitability by workload type



What matters most Siting implication

Source: Kearney analysis

This distinction changes the siting conversation. A location with abundant power may be ideal for AI training and unsuitable for a latency-sensitive inference workload. A market with strong customer demand

may command better pricing, yet still fail if grid congestion, water scrutiny, or local opposition make delivery uncertain. A low-cost site may also disappoint if the demand it was built to serve never follows.

Distance is not always a problem. Once a site is connected to the right communications network, distance mainly shows up through latency, resilience, and data-sovereignty requirements. That is why some workloads can move toward Scandinavia for cooling and low-carbon power, while others stay near Frankfurt, London, Northern Virginia, or other dense network nodes.

For investors, the demand analysis should stay simple: identify the target workload and customer segment; test whether that workload can actually operate in the location; and confirm that the demand case is strong enough to support the investment horizon and the next phase of expansion.

Demand and workload fit determine whether the capacity can be monetized. If the workload does not belong in the location, the site should not advance to the constraint screen.

Step 2: Apply six siting lenses—can the site support the workload?

The six siting lenses are not a stand-alone checklist. They test the same commercial question from different angles: can this site serve the intended demand better than the alternatives, and can it keep doing so as constraints tighten (see figure 3)?

Figure 3
Six siting lenses can help choose the best data-center location

①	Power path	Can the site secure enough deliverable, scalable, and cost-competitive power on the timeline customers need?
②	Low-carbon power optionality	Will the site's energy strategy remain credible as customers and regulators ask harder questions about clean power?
③	Connectivity, latency, and sovereignty	Can the site connect to the right networks, customers, and jurisdictions for the workload it is meant to serve?
④	Heat reuse and local energy integration	Can the site turn waste heat or energy infrastructure into a local benefit rather than a future liability?
⑤	Water and cooling resilience	Can the site operate reliably if water availability, cooling needs, regulation, or local scrutiny become more constrained?
⑥	Social license and permitting	Can the project keep moving once communities, regulators, and local attention increase?

Source: Kearney analysis

1. Power path: can the site secure power that is deliverable, scalable, and commercially credible?

Power remains the first feasibility screen, but the definition of available power has changed. It's not enough to ask whether a utility has indicated that power may be available. Investors need to know whether the power is deliverable, scalable, reliable, cost-competitive, sustainable, and politically durable.

In many markets, interconnection queues and grid upgrades are becoming the gating issue. A site with 200 MW theoretically available in 2031 is not the same as a site with 80 MW firm in 2027 and a credible path to expand. Power that depends on major network upgrades, exposes customers to volatile costs, or lacks a credible long-term supply strategy should be underwritten differently from power that is firm, phased, and commercially credible.

It is important to verify that the data center can be part of the power solution rather than only part of the problem. Some workloads, especially AI training and offline batch compute, may be able to shift timing, ramp up when renewable supply is abundant, or participate in grid-support arrangements. On-site or proximate generation, storage, and demand-response capabilities can therefore become siting advantages rather than engineering add-ons.

Texas illustrates the opportunity. Renewable generation in West Texas can exceed the transmission capacity available to move power to major demand centers. Flexible AI-training and batch workloads can absorb power that might otherwise be curtailed and reduce demand when the grid is tight. They are not batteries, and the model depends on software controls, interconnection terms, and adequate network access, but it can make flexible compute part of the power solution.

Ireland shows how quickly power access can become conditional. [Data centers accounted for 22 percent of Ireland's metered electricity consumption in 2024](#), according to the Central Statistics Office. As system pressure increased, the regulator moved toward a connection policy that asks new data centers to bring system support, such as on-site or proximate generation, storage, market participation, or additional renewable supply. The lesson is broader than Ireland: grid access is shifting from "can you connect?" to "can your load fit into, or help support, the local power system?"

Former power-plant sites can also offer a head start through industrial land, transmission access, and established zoning. Google's Alabama data center has operated since 2019 on the former Widows Creek coal-plant site and is now being expanded. The example is promising, but legacy infrastructure does not remove the need to test remediation, water, clean-power strategy, and community support.

2. Low-carbon power optionality: will the energy strategy remain credible over time?

Low-carbon power should not be treated as a sustainability sidebar. It is increasingly part of the commercial power strategy, because hyperscalers and large enterprise customers are asking more detailed questions about how facilities are powered.

The plain-English question is whether the site's power strategy will still be acceptable as expectations rise. Where does the electricity come from? Does it support new clean supply? Is it reliable enough for the intended use? And can the facility manage the cost and delivery risks that come with different energy options?

Renewable power can reduce exposure to fuel-price volatility and support customer requirements, but it can also introduce delivery risk and balancing needs. Not all clean-power claims are equal. A site that buys

certificates while consuming the prevailing local grid mix has a different risk profile from one that contracts for new supply on the relevant grid, matches consumption more closely in time, and uses storage or load flexibility to manage intermittency. Investors should ask what is physically available, what is additional, and what is primarily an annual accounting match.

3. Connectivity, latency, and sovereignty: can the site serve the workload it is meant to host?

Connectivity used to be treated as a box to tick after power and land. That is no longer enough. A site can have abundant power and still be wrong if it cannot meet the network, latency, or jurisdictional requirements of the target workload.

The question is not simply whether fiber is nearby. It is whether the site connects to the right network corridors, Internet exchanges, cloud on-ramps, customers, and sovereign boundaries. For some workloads, proximity to users is the key issue. For others, proximity to a high-speed Internet node matters more than proximity to a city center. For sovereign and regulated workloads, the relevant boundary may be the EU, a country, a state, a sector, or a specific customer-controlled environment.

Frankfurt shows why network density can outweigh geographic distance. DE-CIX Frankfurt, Europe's largest Internet exchange, recorded a 19.6 Tbit/s traffic peak in July 2026. Its concentration of carriers, cloud on-ramps, financial institutions, and regulated enterprise demand makes it a natural hub for real-time AI inference, interconnection-heavy, and cloud workloads, even as power and land constraints make it less suitable for highly flexible training workloads.

This is also where smaller deployments matter. Hyperscalers and cloud providers increasingly support customer-site or dedicated deployments where enterprises do not want certain data or workloads to reside in public-cloud regions. That does not replace large campuses, but it does remind investors that data-center demand can arrive in more than one physical format.

4. Heat reuse and local energy integration: can the site turn waste heat into local value?

Heat reuse is a separate siting question from low-carbon power. Low-carbon power is about where electricity comes from. Heat reuse is about whether some of the energy used by the facility can create value after it leaves the data center.

Heat reuse is most valuable where there is a credible local off-taker, such as district heating, industrial users, greenhouses, or public buildings. A rural site with abundant power may have little practical heat-reuse potential. A denser location may face more permitting scrutiny, but it may also have more ways to turn waste heat into a local benefit.

Meta's Odense data center offers a practical example. Its heat-recovery system was designed to recover 100,000 MWh annually, enough to heat about 6,900 homes through the local district-heating network. The underwriting lesson is not to assume meaningful heat revenue; it is to confirm the off-taker, required temperature lift, infrastructure ownership, seasonality, and economics before assigning value to heat reuse.

Policy is moving in this direction. In the European Union, the Energy Efficiency Directive requires data centers above certain thresholds to report energy-performance information, including indicators related to waste heat and heat reuse. Germany has gone further through its Energy Efficiency Act, which introduces data-center-specific requirements around energy efficiency, renewable power, and waste-heat reuse. In the United States, state-level policy attention is also emerging: Virginia's HB323 directs the state Department of Energy to identify opportunities and policy recommendations for beneficial use of waste heat from data centers.

5. Water and cooling resilience: can the site operate if water becomes constrained?

Water is local, seasonal, emotional, and politically visible. A corporate water-positive commitment may not satisfy a community worried about local aquifers, drought, or ratepayer impacts.

The siting question should be concrete. What water source is being used? Is it potable, reclaimed, or industrial? How do withdrawals change by season? Is the basin already stressed? Can the cooling design adapt if politics, regulation, or climate conditions change?

Water-positive strategies should also be tested locally. Replenishing water in the same basin can be powerful when the timing, seasonality, and hydrology make sense. A commitment to restore water somewhere else may not resolve the local concern if the project is seen as adding pressure to an already stressed system.

Arizona illustrates the point. In Tucson, the proposed Project Blue data-center campus was presented as a major economic-development opportunity, with a reclaimed-water infrastructure commitment and a promise to be water positive. Public concern over water, power demand, and transparency pushed the project into local politics. Tucson ultimately rejected the development path and adopted rules requiring large water users to provide more detailed applications and conservation plans before connecting to the city water system. The lesson is not that Arizona is a bad market. It's that in water-stressed, fast-growing regions, a water strategy and a trust strategy have become part of site selection, not mitigation after the fact.

6. Social license and permitting: can the project keep moving once people are paying attention?

Community acceptance is no longer a communications issue that starts after site control. It is a development-velocity, permitting, incentive, and valuation issue.

Large data centers can bring tax revenue, infrastructure investment, and construction activity. They also raise local questions about power use, utility costs, water demand, permanent jobs, farmland, traffic, noise, transmission lines, and local control. These questions need to be addressed early, because a technically feasible site can still become politically fragile.

[Microsoft's 2025 decision in Caledonia, Wisconsin](#) is a useful example. The company had planned a data-center campus on a 244-acre site in Racine County, but after strong local opposition and public concern, Microsoft said it would not move forward with that specific site, citing community feedback. Importantly, the company did not abandon Southeast Wisconsin; it said it remained committed to investing in the

region and would work with local leaders to identify a site that better aligned with community priorities and long-term development goals. That is the point: social license does not necessarily mean a market is closed, but the exact site, process, and local value proposition matter.

For investors, this should be part of underwriting. A site that depends on communities accepting large load growth while receiving limited visible benefits is carrying underpriced risk. A stronger site is one where local benefits (tax treatment, water use, power strategy, and so on) can be explained in terms residents and officials understand.


Step 3: Score the site—weight each lens for the selected workload

The goal is not to give every factor equal weight. It is to make sure the weighting reflects the job the site is supposed to do.

The weighting should change by workload. A hyperscale AI training campus will likely place more emphasis on scalable power and cooling resilience, while an inference deployment will likely prioritize latency, network depth, sovereignty, and reliability (see figure 4).

Figure 4
Siting factors should be weighted to reflect the job the site is supposed to do Kearney – data-center siting framework

Relevance of each siting constraint per siting archetype



Source: Kearney analysis

- **Power-path security and low-carbon energy access** are near-universal imperatives: critical for compute-intensive workloads such as AI training and cloud hyperscale, and still highly relevant across the board.
- **Connectivity and digital sovereignty** are more differentiated. They are defining constraints for real-time inference, interconnection-heavy, and sovereign workloads, but less central for AI training clusters that can operate remotely.
- **Water and cooling resilience** follow a similar pattern: decisive for large-scale AI and cloud deployments, but less material for edge, on-premises, or outpost deployments.
- **Heat reuse and local energy integration** are more situational, becoming most relevant where co-location with industrial consumers or district-heating networks is feasible.
- **Social license and permitting risk**, often treated as secondary, remain material across every workload type; community acceptance and regulatory speed can make or break even the most technically optimized site.

In every case, once credible demand is established, investors should apply a weighted scoring model to test whether the site suits the selected workload and eliminate sites that do not support confidence in a profitable, sustainable, long-term investment.

Conclusion: plan for scalability; location optionality is the next advantage

For investors and data-center developers, much of the value lies beyond phase one: in the ability to scale and ultimately exit successfully. A project can clear its initial hurdles and still become a weak platform if any of the six siting constraints limit expansion. The highest-return site protects near-term revenue while preserving long-term resilience.

Data-center siting has entered a new phase. A more diverse set of AI-era siting archetypes and data-center workloads must now be matched to a broader set of constraints. Sustainability matters because it is increasingly the operating language of resource resilience.

The next advantage is location optionality: the ability to place the right workload in the right geography, secure power that can actually arrive, connect to the right network, manage water and climate risk, create local value where possible, and preserve expansion value for the next owner.


As we noted in our previous article on social license to operate, this is not simply corporate responsibility; it is capital discipline. The next phase of digital infrastructure will reward those who look beyond the fastest available megawatts and underwrite the long-term value creation of the location. Making megawatts more sustainable is not only good for the planet, it is increasingly a strategic decision for the bottom line.

Authors



Axel Freyberg

Partner



**Christoph
Neunkirchen**

Partner



Arianna Molino

Principal



Fred Copeman

Consultant

About Kearney

For 100 years, Kearney has been a leading management consulting firm and trusted partner to three-quarters of the Fortune Global 500 and governments around the world. With a presence across more than 40 countries, our people make us who we are. We work impact first, tackling your toughest challenges with original thinking and a commitment to making change happen together. By your side, we deliver—value, results, impact.

For more information, permission to reprint or translate this work, and all other correspondence, please email insight@kearney.com. A.T. Kearney Korea LLC is a separate and independent legal entity operating under the Kearney name in Korea. A.T. Kearney operates in India as A.T. Kearney Limited (Branch Office), a branch office of A.T. Kearney Limited, a company organized under the laws of England and Wales.