

2026 Data Center Power Report

When Power Defines Growth:

How Power Availability is Reshaping
the Data Center Industry





Executive Summary

One year ago, Bloom Energy's inaugural **Data Center Power Report** documented an emerging reality: AI-driven compute demand was beginning to outpace the grid's ability to deliver power at scale. Six months later, the **Mid-Year Pulse** showed clear signs of stress—widening interconnection timelines, rising uncertainty, and growing interest in alternative power strategies.

This year's survey findings indicate that the pressures have only intensified. Power availability has moved beyond being a planning consideration to become a defining boundary on data center growth.

Our research with hyperscalers, colocation providers, utilities, and equipment suppliers reveals there is a structural shift already underway.¹ Capital is concentrating in power-advantaged regions, long-standing assumptions about grid delivery timelines are being challenged, and onsite generation is increasingly being incorporated into long-term power strategies rather than treated as a temporary bridge. At the same time, data center campuses are scaling toward gigawatt-scale AI factories, accelerating the shift of next-generation electrical architectures from roadmap concepts to near-term designs.

Together, the findings from the **2026 Data Center Power Report** show that the determinants of data center growth have changed in the AI era.

¹Surveys were commissioned to enable a double-blind process between Bloom Energy and survey respondents; surveys were conducted in April 2024 (n=30), November 2024 (n=64), April 2025 (n=100), and November 2025 (n=152), and the latest survey included hyperscalers, colocation providers, utilities and IPPs, and equipment providers, with 84% of respondents U.S.-based. Interviews were also conducted with industry leaders to pressure test findings and to understand real-world implications

Key findings

The following findings summarize how reported shifts are translating into real-world changes across markets, development strategies, and technology choices.

- 1. Power availability is driving a geographic reallocation of data center growth, creating clear winners and losers.** Within three years, Texas is poised to become the nation's leading data center market. Legacy markets, such as California and Oregon, are expected to lose more than half of their relative market share. Despite Tier 1 markets largely driving capacity growth, emerging markets are collectively expanding their share by over 20% in aggregate.
- 2. An increasing share of data center campuses are expected to exceed gigawatt scale.** By 2030, about one in five data center campuses are expected to exceed gigawatt scale, rising to one in three by 2035. At this scale, power delivery remains the primary constraint, but cooling capacity, water availability, permitting complexity, and network infrastructure will emerge as critical challenges.
- 3. Developers and utilities remain misaligned on time-to-power, with gaps widening in key hubs.** Utility respondents say time-to-power will take ~1.5-2 years longer on average than expected by hyperscalers and colocation providers. Over the last six months, the power expectation gap has widened in three critical hubs—Northern Virginia, the Bay Area, and Atlanta.
- 4. With grid constraints, onsite power is increasingly entrenched in data center power strategies.** Over the past six months, expectations among hyperscalers and colocation providers for fully onsite-powered data centers have increased by 22%. Over one-third of data centers are expected to use 100% onsite power by 2030 and developers increasingly expect permanent onsite generation to emerge as a preferred long-term approach.
- 5. Next-generation data center architectures are arriving within the next two years.** Rising rack densities are accelerating the adoption of next-generation data center architectures. Higher-voltage central busways and direct current (DC) distribution architectures are expected to lead this shift: by year-end 2028, 60% of respondents expect to adopt high-voltage central busways, and 45% expect to adopt DC architectures. These designs are likely to be incorporated into data centers entering development this year.

1. Power availability is driving a geographic reallocation of data center growth, creating clear winners and losers

Data center demand is accelerating rapidly, with estimates suggesting that total U.S. IT load capacity could roughly double over the next three years from ~80 GW in 2025 to ~150 GW in 2028—more than twice the level projected in 2024 forecasts.^{2,3}

Growth remains concentrated in a small number of top hubs, where developers are expanding existing campuses and adding new phases, preferring to avoid unfamiliar geographies to mitigate stranded-asset risk as AI workloads shift toward inference.⁴

Within this group, Texas is expected to stand out as a winner. By 2028, Texas is projected to exceed 40 GW of capacity—nearly 30% of total U.S. demand—representing a 142% increase in market share relative to today.

Georgia is also gaining momentum, with market share projected to rise 75% as developers expand deeper into the Southeast. While Virginia continues to grow in absolute terms, its relative share is expected to decline as grid constraints intensify in one of the industry's most mature markets.

“We’re seeing a geographic shift as certain regions become more power-friendly and therefore more attractive for data center construction. While conditions still vary widely, more areas are beginning to recognize the opportunity and are working with operators to make development easier.”

VP of Energy at hyperscaler

By contrast, several legacy and secondary markets are losing ground. California, Iowa, Oregon, and Nebraska are each projected to see declines of more than 50% in relative market share, reflecting tighter power availability, permitting complexity, and longer interconnection timelines that limit their ability to support large-scale AI-driven expansion.

At the same time, **the rest of the U.S. is emerging as a meaningful aggregate winner**, with market share projected to increase 21% by 2028. This increase reflects a shift toward a broader set of power-advantaged states, as developers pursue regions where large blocks of power can be secured more quickly. Though individual markets like Mississippi and Oklahoma remain smaller in absolute terms, their growth highlights how far developers are willing to extend siting flexibility to meet AI-driven compute demand—even when that requires trade-offs in network proximity and workload placement.⁵

² McKinsey & Company: “Scaling bigger, faster, cheaper data centers with smarter designs”

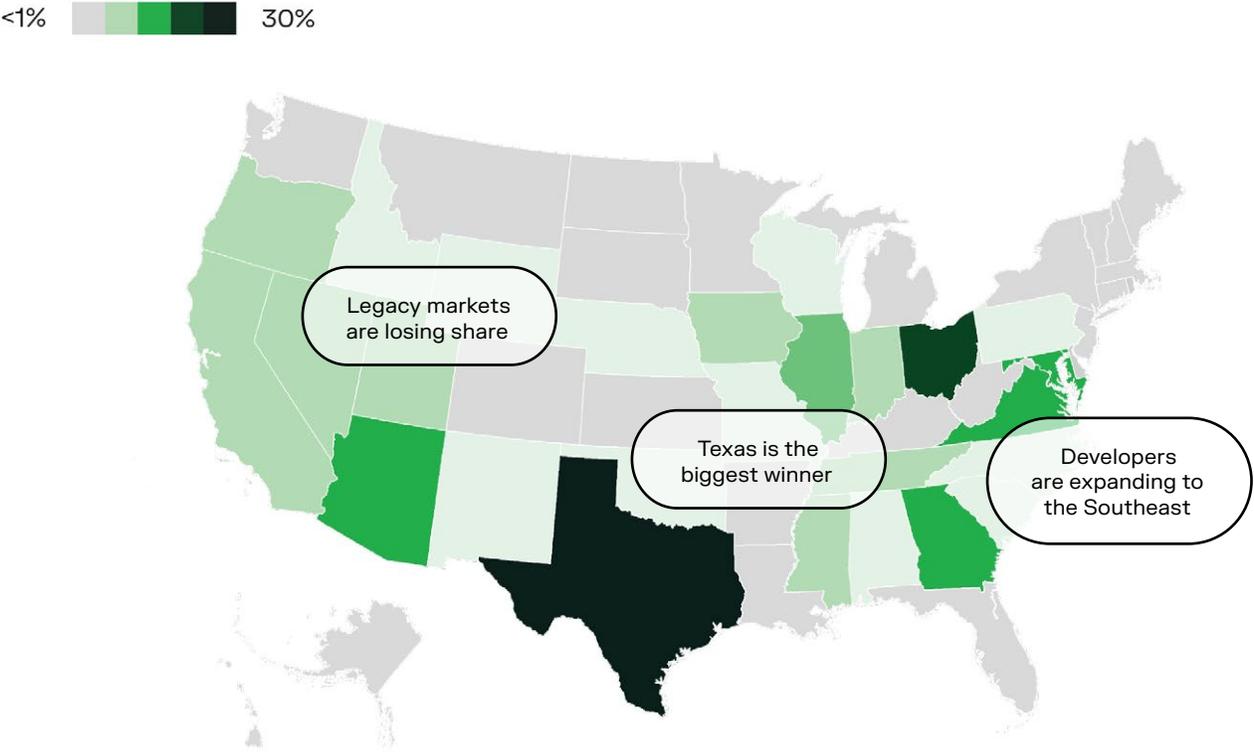
³ McKinsey & Company: “How data centers and the energy sector can satiate AI’s hunger for power”

⁴ Bloomberg Intelligence: “AI inference, productivity tools offset Gen-AI margin headwinds”

⁵ For instance, while Oklahoma is expected to only capture ~2% market share by 2028, it is expected to grow deployed capacity 6x between 2025 and 2028

As developers increasingly focus on power-advantaged regions, the geographic distribution of data centers will shift

Expected data center market share by state in 2028



Expected change in market share in top 2025 markets

Texas	142%	California	-50%
Georgia	75%	Iowa	-60%
Arizona	0%	Oregon	-67%
Ohio	-17%	Nebraska	-75%
Illinois	-25%	Rest of U.S.	21%
Virginia	-35%		

Largest gains in rest of U.S.		
Idaho	Wisconsin	S. Carolina
Mississippi	N. Dakota	Indiana
Maryland	Louisiana	S. Dakota
Kansas	Kentucky	

Source: DC Byte analysis as of Oct. '25, including 50 states and the District of Columbia

2. An increasing share of data center campuses are expected to exceed gigawatt scale

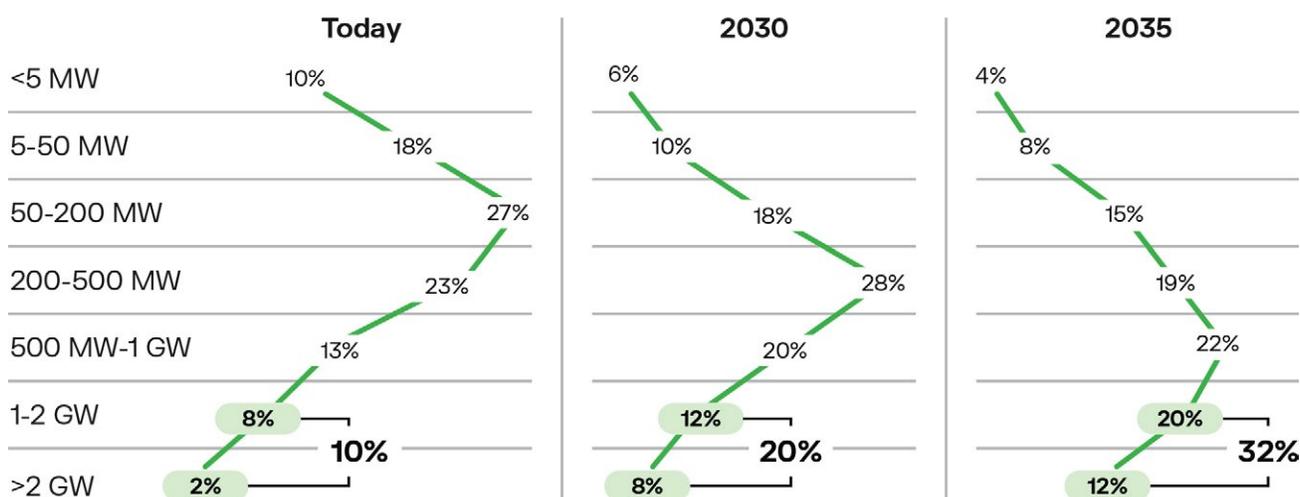
The surge in demand is being fueled not only by more data centers but by significantly larger ones. Even as developers continue building smaller, low-latency edge sites near metropolitan centers, hyperscaler and colocation provider survey respondents indicate that the **share of new campuses expected to exceed 1 GW rises from roughly one in five in 2030 to nearly one in three by 2035**. To contextualize this scale, each 1 GW campus would consume as much as roughly 20% of New York City’s entire electricity load.⁶

Upward revisions of tens of gigawatts in long-term load forecasts from grid operators between 2024 and 2025 underscore that developers’ expectations are anchored in observed shifts in grid planning assumptions rather than speculative optimism. ERCOT, for example, increased 2030 estimates of data center growth from 29 GW to 77 GW, while PJM revised 2030 peak demand estimates upwards by 10%.^{7,8}

As facilities approach gigawatt scale, developers face challenges beyond power availability, and developers expect cooling capacity, water access, permitting complexity, and network infrastructure to become significant bottlenecks.

Developers expect almost one third of data center sites to exceed 1 GW in scale by 2035⁹

Distribution of data center sizes expected to be a part of developers’ portfolios



Source: Bloom Energy Data Center Survey (Nov. 2025)

⁶ New York City Mayor’s Office of Climate & Environmental Justice, 2024

⁷ ERCOT 2025 TSP Long-Term Load Forecast Report

⁸ PJM 2025 Long-Term Load Forecast Report

⁹ N=92; Survey question: Which of these data center formats are expected within your organization’s portfolio of data centers?

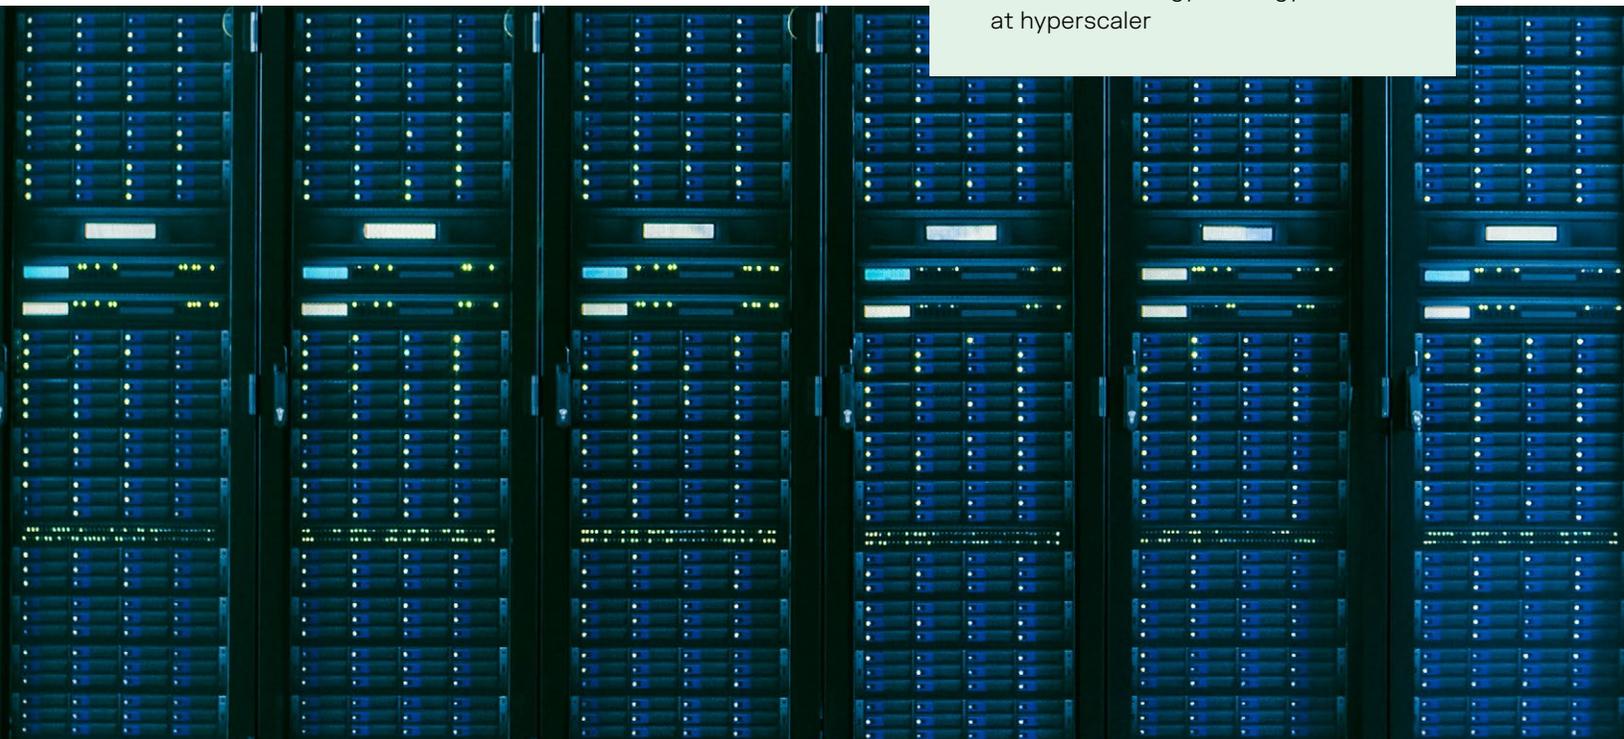
3. Developers and utilities remain misaligned on time-to-power, with gaps widening in key hubs

Power access remains the single most important determinant of site selection, driving developers to prioritize markets where interconnection can be achieved quickly and predictably. Even so, **more than half of developers report that securing power has become more difficult over the past twelve months.**¹⁰

The fundamental disconnect in time-to-power expectations identified in the June 2025 Data Center Power Report remains firmly in place. **Survey data show that hyperscalers and colocation providers consistently expect power to be available up to two years earlier than utilities and IPPs believe they can deliver it.** In three geographies—Northern Virginia, the Bay Area, and Atlanta—this gap has widened in the past six months, reflecting the extended interconnection queues, permitting delays, and aging transmission infrastructure in major hubs.

“Grid capacity continues to fall behind demand, and I’ve never seen conditions improve—only deteriorate. Building the necessary infrastructure is extraordinarily difficult, which is why project lead times are becoming longer and far more uncertain.”

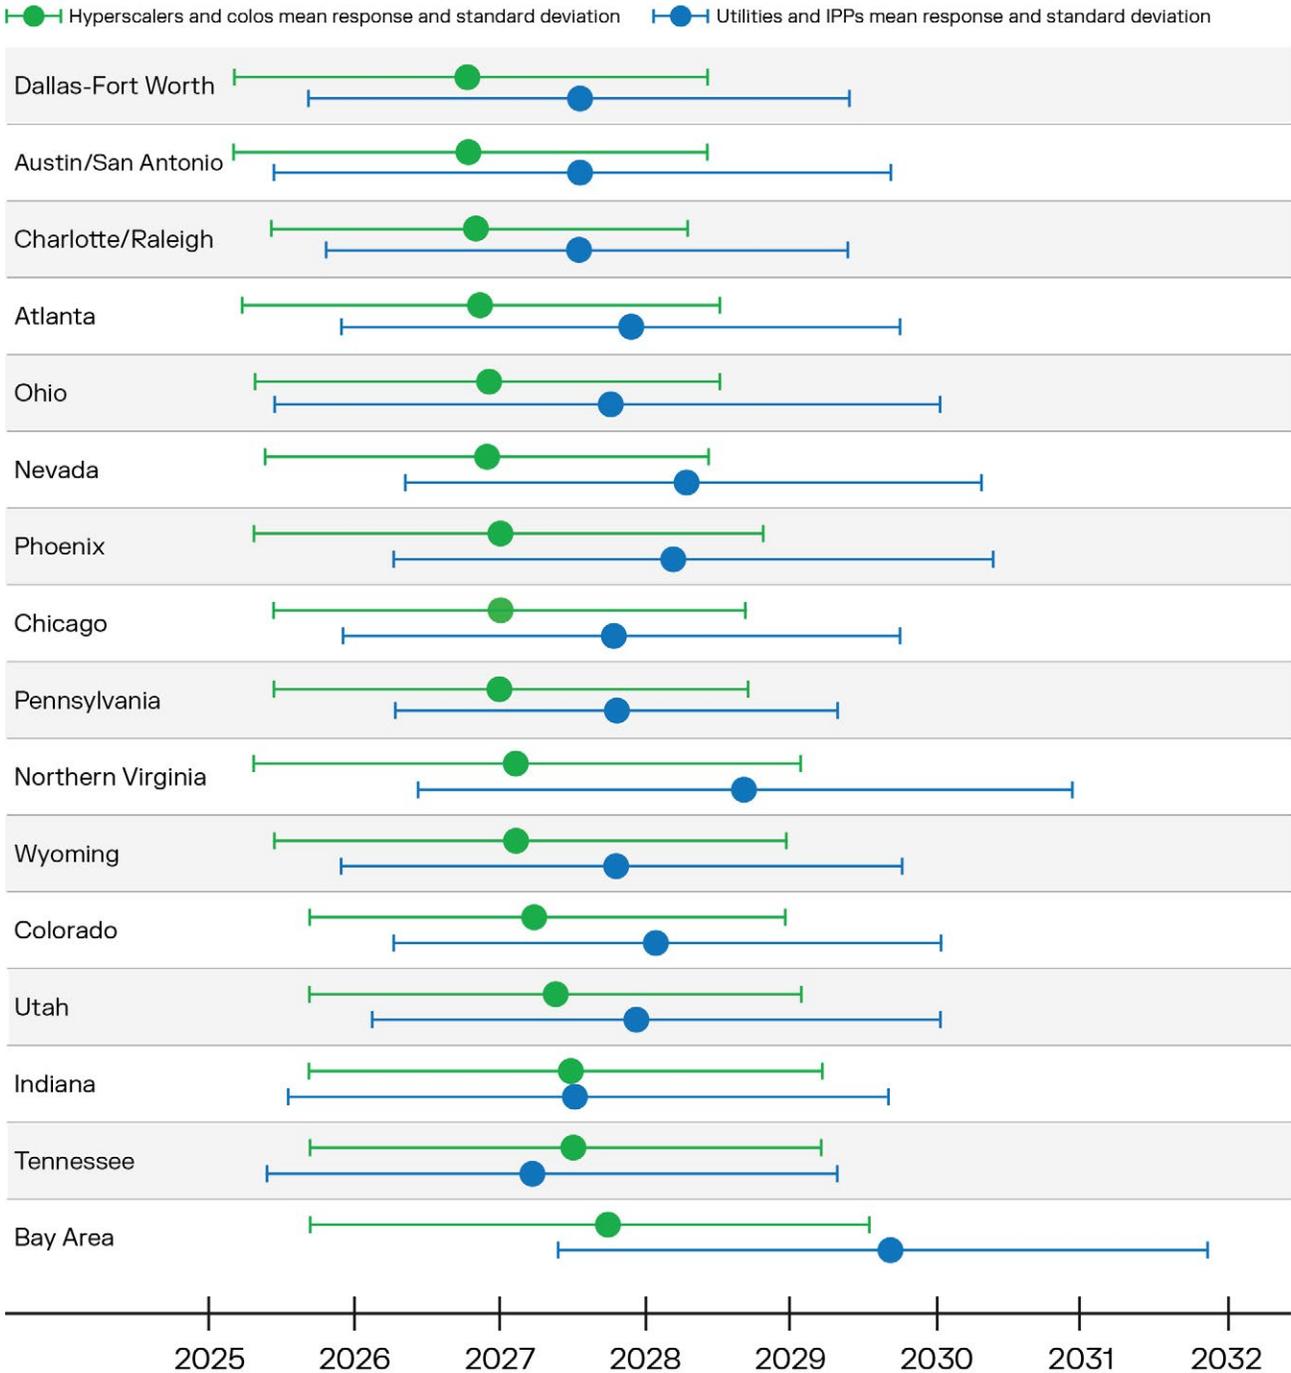
Director of Energy Strategy
at hyperscaler



¹⁰ N=92; Survey question: In your experience, how has the ability to access power at data center sites changed over the past twelve months?

Developers systematically underestimate lead times for utility power across markets¹¹

Data center market expectations of the earliest a utility can provide power



Source: Bloom Energy Data Center Survey (Nov. 2025)

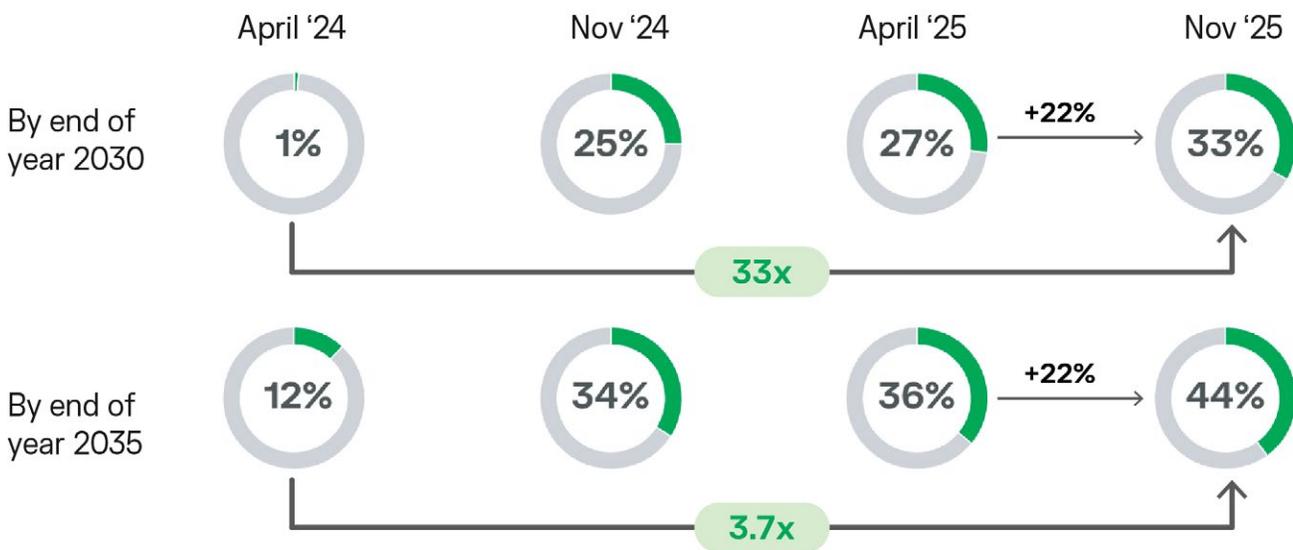
¹¹ N=113; Survey question: For each of these data center markets, when is the earliest that a utility would likely be able to give you power?

4. With grid constraints, onsite power is increasingly entrenched in data center power strategies

In this survey, average response from hyperscalers and colocation providers **suggests that roughly one-third of data centers in 2030 will use 100% onsite power**, a 22% increase from the previous report. Expectations for onsite power solutions are rising amid weakening confidence in grid delivery timelines. In parallel, 73% of respondents report actively evaluating or selecting onsite power providers, underscoring the growing emphasis on fast, predictable time-to-power.

Developers' expectations of onsite power deployment continue to grow, signaling a need for fast, predictable time to power¹²

Expected share of data centers with 100% onsite generation



Source: Bloom Energy Data Center Survey (Nov. 2025)

This shift is also evident in how developers rank solutions to accelerate timelines and control costs. Although developers would prefer to rely on grid power where and when available, onsite generation has increasingly become a necessary tactic. Temporary bridge-to-grid power remains the most common near-term deployment model, but **by 2030, permanent onsite power emerges as the top-ranked solution for minimizing development timelines and costs**—reflecting a familiar pattern in other power-intensive industries such as cement, refining, and steel.¹³

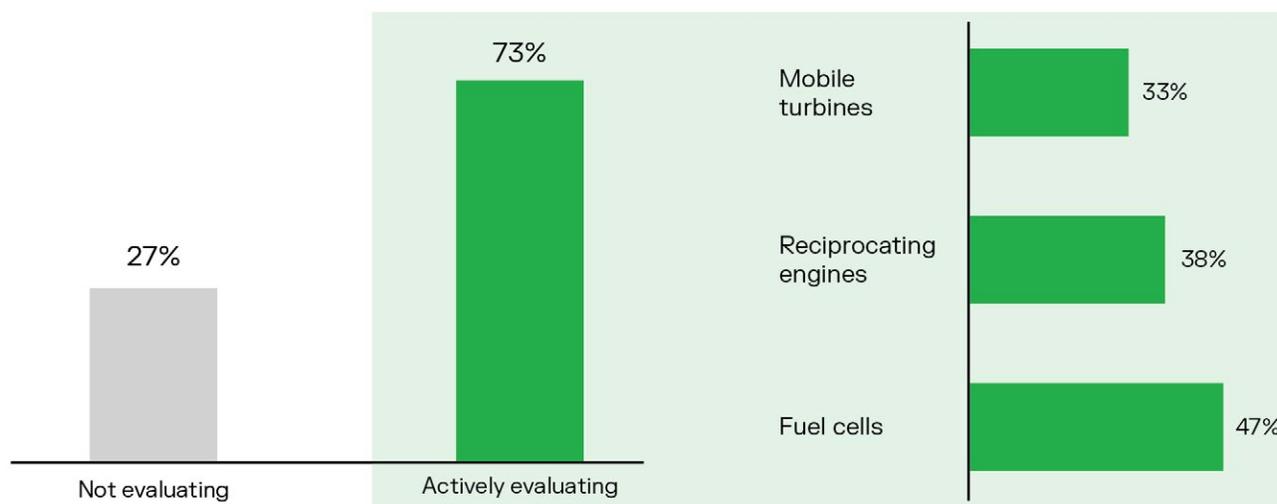
¹² N=30 in April 2024, n=64 in November 2024, n=80 in April 2025, and n=92 in November 2025; Survey question: What percentage of data centers sites in the U.S. do you expect to be powered fully by onsite power?

¹³ N=92; Survey question: Which solutions are you most likely to deploy to help accelerate your data center execution timeline and/or save money on power costs?

Within this context, **fuel cells are emerging as a leading onsite technology under consideration**, with developers recognizing four practical advantages: (1) Shorter lead times that reduce time-to-power risk, (2) Lower local emissions, which support faster permitting and greater community acceptance, (3) Scalability and modularity, including the ability to ramp output, and (4) Alignment with long-term sustainability objectives such as achieving 24/7 carbon-free energy compliance.

The majority of developers are actively evaluating onsite generation technologies, and fuel cells are a leading technology under consideration

Share of developers evaluating onsite power solutions¹⁴ and corresponding technologies under evaluation or already deployed¹⁵



Source: Bloom Energy Data Center Survey (Nov. 2025)

¹⁴ N=92; Survey question: What is your likelihood of choosing onsite power for some or all of your generation needs for data center expansions? Rate from have not considered to have decided to purchase.

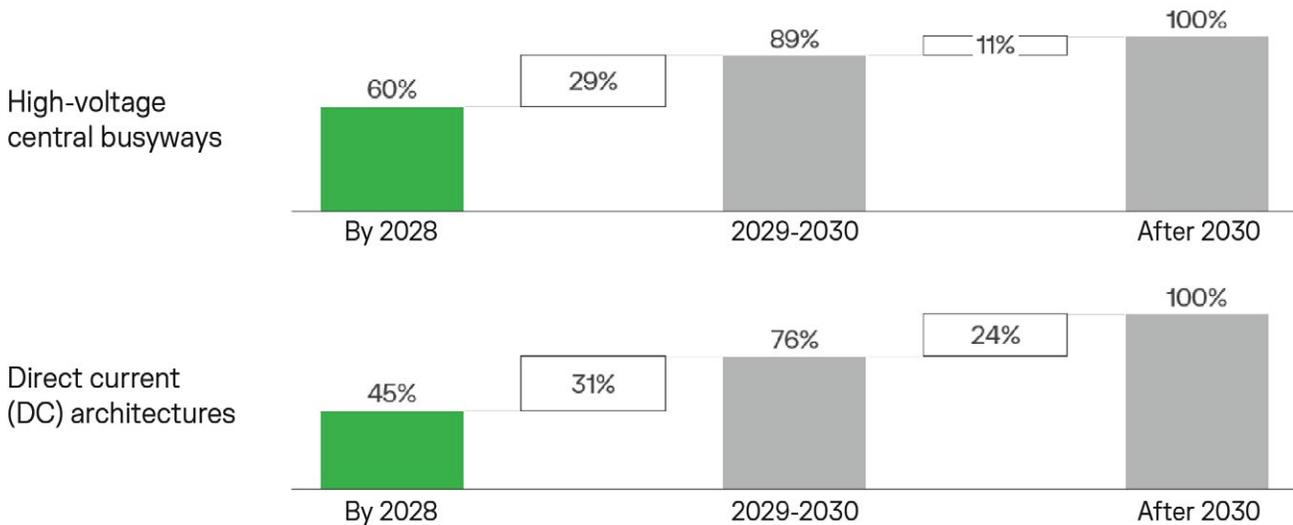
¹⁵ N=92; Survey question: Which technologies are you seriously evaluating for deployment in power constrained markets? Rate from 1-10, where 1 = not evaluating and 10 = actively deploying. Scores of 8 or higher are classified as “under evaluation” or “already deployed.”

5. Next-generation data center architectures are arriving within the next two years

As data centers scale toward gigawatt-scale development, traditional electrical architectures are expected to become increasingly inadequate. **Rising rack densities are driving operators to adopt next-generation electrical designs that deliver higher power density, greater efficiency, and faster deployment.** Higher-voltage central busways and direct current (DC) distribution architectures are expected to be among the first next-generation designs prioritized: by year-end 2028, 60% of respondents expect to adopt higher-voltage central busways and 45% expect to implement DC architectures. These designs are likely to be incorporated into the crop of data centers entering development this year.

High-voltage and direct current architecture are expected to be among the first next-generation designs adopted¹⁶

Developers prioritizing or planning to adopt next-generation architectures, %



Source: Bloom Energy Data Center Survey (Nov. 2025)

¹⁶ N=92; Survey question: Which next-generation data center technologies or design advances is your organization prioritizing or planning to adopt, and when will you adopt them?

Next-generation architectures offer meaningful system-level advantages, particularly as data centers scale towards AI-dense workloads:

- **More efficient campus-scale power distribution.** High-voltage busways move large blocks of power more efficiently across expansive campuses through simpler, modular distribution paths, allowing for greater redundancy as facilities scale.
- **Greater efficiency and improved thermal performance.** Direct current systems reduce power-conversion stages, limiting associated losses and heat while delivering more stable power for high-density AI workloads.
- **Cleaner integration with onsite power sources.** Onsite generation integrates more directly with these architectures, minimizing conversion steps and enabling shorter, more efficient distribution paths from generation to load.

Realizing these advances will require addressing practical implementation barriers, including updating legacy AC-centric designs, establishing clear safety and permitting standards, strengthening supply chains for emerging architectures, and expanding design standards, integration practices, and technician training.



Conclusion

The realities of grid constraints, long interconnection timelines, and rapidly scaling AI loads mean that power can no longer be treated as a downstream consideration. Operators that continue to plan projects around traditional grid assumptions risk falling behind on both schedule and scale. In this environment, power strategy must move to the front of the development process.

The next chapter of development will belong to operators who act earlier and plan more aggressively—securing flexible power paths, aligning electrical architectures with AI-scale requirements, and expanding where and how they build. Those that do not will increasingly find that power is the limiting factor on growth.

“Challenges and obstacles certainly exist, but the industry’s trajectory is clear—people will keep scaling data centers as far and as fast as they can.”

Executive at the Data Center Coalition





About Bloom Energy

Bloom Energy empowers enterprises to meet soaring energy demands and responsibly take charge of their power needs. The company's solid oxide fuel cell systems provide ultra-resilient, highly scalable onsite electricity for Fortune 500 customers around the world, including data centers, semiconductor manufacturing, large utilities, and other commercial and industrial sectors. Headquartered in Silicon Valley, Bloom Energy employs more than 2,000 people worldwide and manufactures its systems in the United States. For more information, visit bloomenergy.com/industries/data-center-power

**Bloom
energy**

bloomenergy.com

