



A Primer on Investing in Data Centers

WHITEPAPER | FEBRUARY 2026

CONTRIBUTORS

Alison Adams, PhD
Christy Fields
Lisa Bacon, CAIA

In 2025, the Artificial Intelligence (“AI”) buildout turned data centers into one of the biggest stories in the economy. This wave was large enough to show up in macro data, including estimates that AI-related investment (including data centers) accounted for a very large share of GDP growth, and that without this spending overall GDP expansion would have been nearly flat.¹ Forecasts point to even more spending on the horizon, including an estimate that the “big five” hyperscalers will spend roughly \$450 billion on capital expenditures (capex) related to AI infrastructure such as data centers and related equipment in 2026.² The projected capital spending by Meta, Amazon, Microsoft, and Alphabet combined as a percentage of GDP in 2026 is estimated to be on par with that for the US railroad expansion of the 1850s.³

Investors may find data centers to be a fast-growing segment of income-oriented real assets that can be accessed through familiar vehicles. Exposure is typically gained through listed Real Estate Investment Trusts (“REITs”), private real estate funds, or infrastructure strategies that own or finance data center facilities, power and cooling upgrades, and related development. Owners generally seek returns by developing or acquiring facilities, leasing capacity under multi-year contracts, and operating the assets to high availability and reliability standards. Because these leases and revenue arrangements are often long term and counterparties are typically large, creditworthy users, data centers can potentially offer a combination of current income and growth, with risk and return drivers that may differ from those of more traditional physical assets.

Key Takeaways

- › AI is turning data centers into a macro-relevant asset class. Capital spending on AI-related data center infrastructure is large enough to influence economic growth, but the returns to investors typically hinge on asset-level execution rather than industry demand alone.
- › Power, not space, is the binding constraint. The economic value of a data center is driven by its ability to deliver reliable, usable power on schedule, particularly where grid capacity, equipment availability, and permitting are constrained.
- › Cash flows can look infrastructure-like but still vary by project. Long-term contracts with tenants often support stable income, yet outcomes can differ widely based on location, contract terms, and operational performance.
- › Geography matters because constraints are local. Power availability, permitting timelines, grid reliability, and cooling requirements often differ materially by region and can outweigh proximity to end users for large AI workloads.
- › There are multiple ways to invest, each with trade-offs. Public REITs offer liquidity but higher volatility, while private real estate and infrastructure strategies typically provide more direct exposure with greater dispersion of outcomes and reduced liquidity.

What Are Data Centers?

A data center is a purpose-built facility that holds computers and equipment used to store, process, and send digital information. It contains many powerful machines that run software and keep data available for people and organizations. These machines handle tasks like running websites, saving files, managing email, and supporting business systems that people use every day.

Data centers tend to house massive server farms, but they are more than just a collection of computers, storage devices, and networking equipment. In addition to computing power, data centers need steady electrical power delivery, cooling to prevent overheating, strong physical security to limit access, fire suppression, and multiple network connections.

The “product” a data center sells is not square footage in the ordinary real estate sense. Rather, the product is essentially the power used by the computers and related equipment that do the actual work in a data center (see Figure 1). The power needs to be reliably available so services can be provided without interruption. This requirement for reliability leads to engineering redundancy (e.g., backup power, multiple network paths) and operational discipline.⁴

Because of the massive amount of power that is needed not just to run the servers but the cooling and other equipment, data centers can also be understood as massive energy demand hubs. Small data centers (e.g., a single company site) often draw hundreds of kilowatts, while large data centers (e.g., major cloud providers) often draw tens of megawatts,⁵ and some campuses can be 100+ megawatts across multiple buildings.⁶

The emphasis on contracted powered capacity reflects a fundamental constraint in the sector: reliable grid power cannot be expanded quickly or without significant coordination and capital investment. Utilities must evaluate and approve new large-load interconnections, which can require substation upgrades, transmission reinforcement, specialized transformers, and long-lead electrical equipment. In many markets, these processes take multiple years and involve queue positions and regulatory approvals. As a result, the economic value of a project often hinges less on the availability of land or building space and more on the certainty, timing, and cost of securing deliverable power. Contracts and underwriting therefore tend to center on committed megawatts, because those megawatts represent the scarce input that enables the facility to generate revenue.

Data center customers typically comprise three categories: hyperscalers; enterprises; and network and digital service providers. Hyperscalers are large cloud and AI platforms that lease big blocks of capacity or build their own data center campuses for cloud services and AI workloads.⁷ Some enterprises (often in the financial services, healthcare, media, or industrial sectors) may prefer not to build and operate their own facilities but still need high compute capacity and storage, data security and reliable performance. Network and digital service providers (e.g., telecoms, content delivery, Software as a Service or “SaaS” firms) may value the connectivity and proximity to customers provided by data centers. These customers may act as large individual tenants or co-location tenants (i.e., multiple tenants leasing within a shared facility).⁸

Figure 1
Illustrative Sources
of Data Center
Power Demand

Source: 2024 US Data Center Energy Usage Report, Lawrence Berkeley National Laboratory, December 2024.

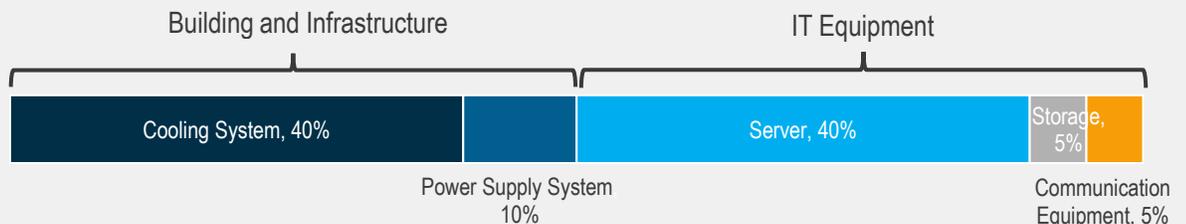
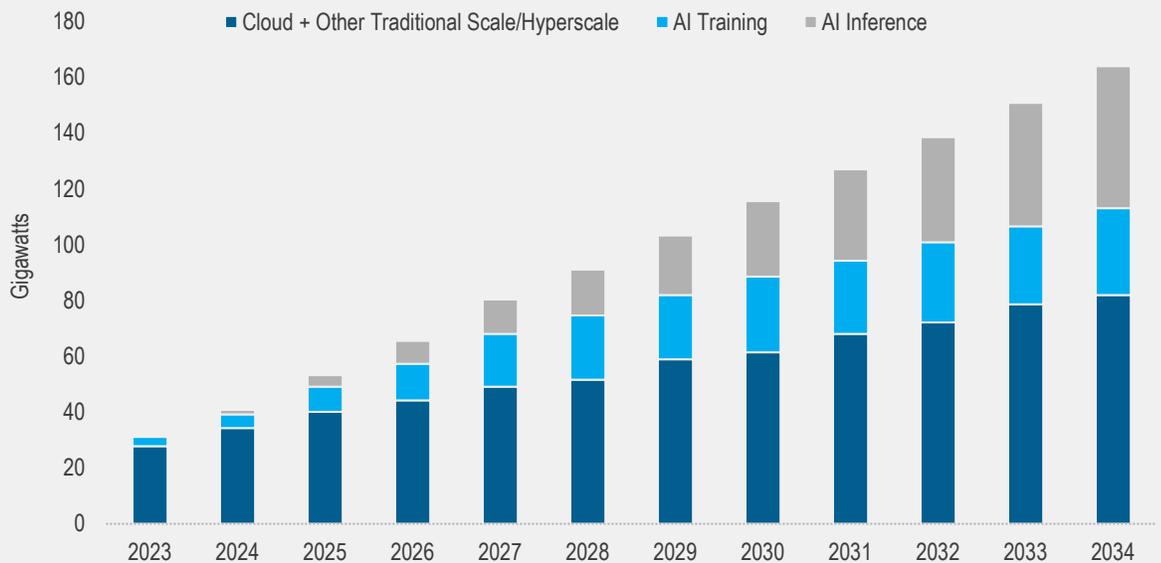


Figure 2
Total Global
Installed and
Projected Data
Center Base
Capacity

Source: Brookfield. "Building the Backbone of AI." August 2025.



Building a modern data center is a substantial undertaking, and one that is only getting more complex. Development starts with getting a site permitted specifically for a data center that has access to appropriate power and can support heavy electrical and mechanical systems. On-site infrastructure often includes utility interconnection, substations, transformers, backup generators, batteries, and either air-based or increasingly liquid-enabled cooling designs. Construction typically includes the shell, electrical rooms, cooling plant, and redundant systems. The building shell is only part of the work; a meaningful portion of cost and time goes into the electrical and mechanical backbone, permitting, and commissioning (i.e., testing redundancy and reliability).⁹ While it is not uncommon for any construction problem to face delays, for data centers the pace is often constrained by long lead times for electrical equipment and the timeline to obtain and energize grid interconnections.¹⁰

As a result of the complexity of the development cycle, most large data center development today, particularly for AI workloads, is build-to-suit for hyperscalers, with capacity pre-leased or contractually committed before construction begins. The combination of high capital costs, power uncertainty, and long lead times has sharply limited speculative development.

The lifespan of the building and its infrastructure is long, often 20 years or more, especially when the shell, land, and major mechanical/electrical systems are appropriately maintained and upgraded.¹¹ However, the computing equipment inside refreshes on much shorter cycles (often three to five years for hardware and one to three years for Graphics Processing Units (“GPUs”) in AI data centers), and design requirements evolve (e.g., higher rack density, more cooling, more power distribution).¹² Hence, a facility can remain valuable if it can be upgraded economically, but it can lose competitiveness if power and/or cooling constraints make retrofits too costly.

The shift toward AI workloads has amplified these design pressures. Earlier generations of cloud data centers were primarily built around lower-density Central Processing Unit (“CPU”)–based servers, with relatively moderate power and cooling requirements per rack. AI training and inference clusters rely more heavily on high-density GPU servers that draw substantially more power and often require more advanced cooling solutions. **As a result**, newer facilities are being designed for higher power density and greater electrical and mechanical capacity, which increases capital intensity and can make some older facilities less adaptable without significant reinvestment.

Data Centers and Revenue Generation

A data center’s revenues usually come from multi-year leases or contracts and recurring service charges. Data center owners/operators can earn income from capacity rent (i.e., space and/or power reserved), power and cooling charges (either passed through to the tenant or bundled), and connectivity fees (e.g., cross-connections to carriers and other tenants).¹³ Contract structures vary, but many are designed to be predictable: long-term leases with escalation clauses, plus mechanisms to allocate electricity costs. In competitive markets, “powered capacity” often comes with higher pricing because tenants value the immediate reliability of running their computing needs, not just the amount of space available.¹⁴

Revenue streams often resemble those of typical infrastructure assets, with long contracts and counterparties that are generally investment grade; however, they may differ materially by location and deal structure. Operators typically earn most of their revenue through leases priced per kilowatt or megawatt (see Figure 3).¹⁵ In most arrangements, tenants contract for a defined amount of powered capacity over a multi-year term (e.g., 20 MW for 10 years), with recurring payments tied to that reserved capacity once it is energized and available. In many co-location arrangements, customers pay for both capacity and power consumption (i.e., both how much power is available to them and how much they

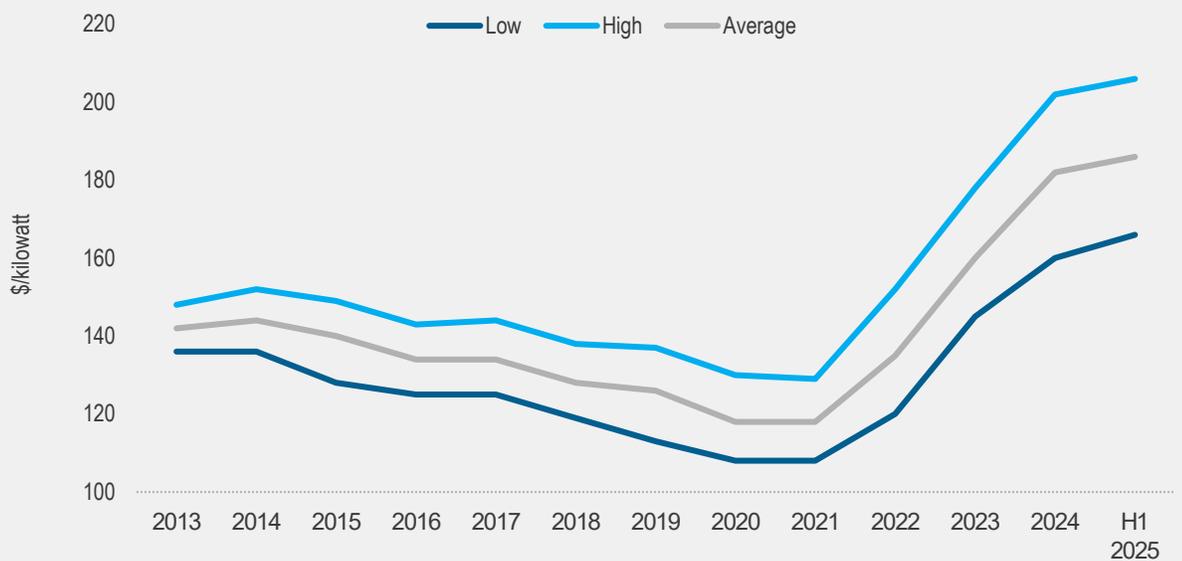
actually use). In single-tenant hyperscaler arrangements, contracts can resemble long-dated offtake agreements, where the economic value lies in the ability to provide usable, reliable power (i.e., “powered capacity”).

Geography matters for data centers because they are physical assets that depend on local conditions and limitations. In practice, location affects a number of factors, first among which is usually power availability and cost. Some electricity grids can add large new loads quickly and cheaply while others cannot. Permitting and construction timelines can vary based on local rules and can create delays. Dispersion in investor outcomes is often driven by deliverability of a project, especially power access and the ability to bring capacity online when promised, because delays can push revenue out and raise costs.¹⁶

Network performance can be affected by proximity to users or major network hubs. Network performance matters substantially for certain workloads, but it is less likely to be the thing that stops a project. Most major data center markets already have strong fiber options, and many applications can tolerate being farther away. Where it is most important is in “latency-sensitive” use cases, such as real-time trading, gaming, or interactive apps where user experience depends on response time. Finally, climate and water constraints influence cooling needs and operating cost.¹⁷

Figure 3
Rental Rates for Data Center Capacity in the US

Source: CBRE Research, CBRE Data Center Solutions 1H25. These are the Asking Rents for US Powered-Shell Data-Center Capacity for \$/kW-month. Illustrative market information only; not a prediction of future rental terms and not indicative of investment returns.



Ways To Invest in Data Centers

For investors that require liquidity, the most accessible investment options are often publicly listed vehicles, typically structured as REITs, but also available through sector-specific mutual funds and exchange-traded funds (“ETFs”). These vehicles can offer exposure to data centers while providing liquidity and transparency. But listed ownership also introduces public-market sensitivities, including substantial volatility and valuation swings driven by market sentiment about the assets, as well as interest rate sensitivity, given that the portfolio is composed of long-duration assets.

For investors who are able to commit capital to private market funds, there are additional choices available for more direct exposure to data centers. Private ownership sits primarily in private real estate and private infrastructure funds. Many private real estate strategies have added data centers to their diversified real estate portfolios in recent years. Likewise, private infrastructure strategies have for some time included digital infrastructure exposure in diversified portfolios and are increasingly adding to it. This includes data centers, fiber, towers, spectrum, and other related services. More recently, there has also been a meaningful increase in both real estate and infrastructure funds with a dedicated digital infrastructure mandate, particularly with a focus on data centers, to meet the growing bandwidth demand for digital networks and AI deployment. Private vehicles may provide smoother reported returns but come with liquidity constraints and greater dispersion by manager/fund.

Investment Considerations

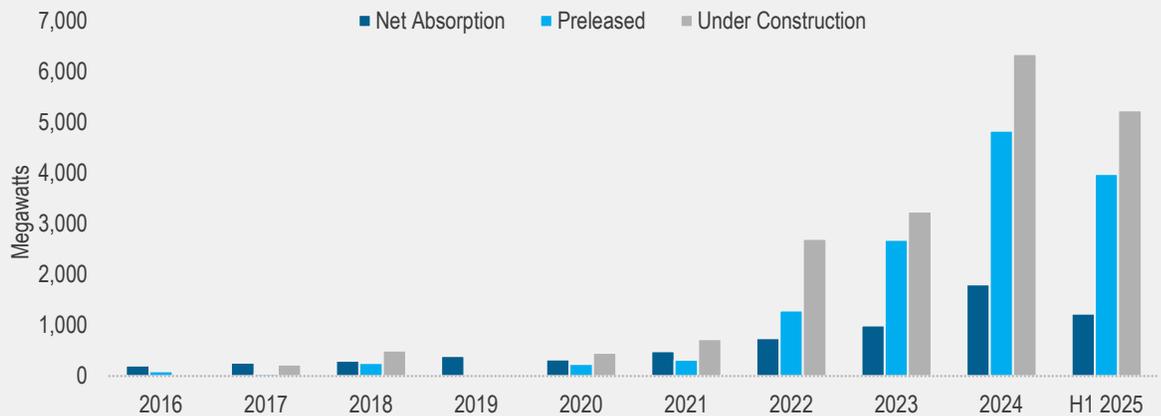
In practice, the return profile for a data center can vary depending on whether it is a brownfield asset or greenfield project (i.e., stabilized vs development), where it is located (e.g., core demand hubs vs fringe markets), and how cash flows are contracted and financed. In this regard, data centers are like many other real estate and infrastructure assets. For example, in major data center markets like northern Virginia and Dallas-Fort Worth, preleasing rates of data centers under construction were reported at above 75% and extend to 2027 on average.¹⁸

A key long-term driver of data center success is power certainty. The project’s timeline expectations, leasing agreements, and the contingency plan for grid delays will also impact returns. Power uncertainty may add timing risks and higher costs to cash flows.

An additional consideration is growing public and regulatory scrutiny around data center development, particularly in regions where large facilities coincide with rising electricity prices, grid congestion, or water usage concerns. In some markets, this has led to slower permitting, additional requirements related to grid upgrades or community impact, and greater uncertainty around development timelines and operating costs. While these dynamics vary significantly by jurisdiction, they can affect both the timing and cost of development, and they reinforce the importance of local market knowledge and proactive engagement with utilities and regulators.

Figure 4
Data Center Construction, Prelease and Absorption in the US

Source: CBRE Research, CBRE Data Center Solutions 1H25. Prior to 2020, data shown represents Meketa estimates based on CBRE and other data center reports shown for scale.



Contract durability is a central factor in determining whether a data center asset generates stable, infrastructure-style cash flows or aligns more closely with cyclical lease-up dynamics. The tenant profile, including credit quality, concentration, and contract duration, plays a vital role, along with the protections embedded in agreements such as escalation clauses, renewal options, and remedies for reduced usage. The treatment of power costs further influences margins: stable pass-through arrangements can help insulate earnings, while owner-borne power cost volatility may introduce regional risks. In competitive markets, long-term, secure contracts can help distinguish assets and support reliable performance.

Market quality and execution capability also shape outcomes for investors. Dense urban environments with robust network connectivity often attract a greater number of creditworthy tenants, offer deeper leasing markets, and facilitate re-tenanting when contracts expire, which contrasts with more speculative locations that might present attractive land or power costs but carry heightened leasing and residual-value risks. The ability to deliver on schedule and maintain operational competency becomes increasingly important as technical demands for cooling and power intensify. Financial resilience, reflected in funding strategies and balance sheet flexibility, is often important for weathering project delays or shifts in market conditions, especially as the sector's reliance on external capital and structured financing continues to grow.

Another consideration for closed-end fund investors is uncertainty around terminal value. Terminal value refers to the estimated value of an investment at the end of a fund's life or holding period, often based on projected cash flows or sale price. Data centers are long-lived physical assets, but their economic relevance depends on evolving power density, cooling requirements, and tenant demand. As a result, exit assumptions, particularly for large, single-tenant or hyperscaler-oriented facilities, can be difficult to underwrite. Investors should consider who the likely buyers will be at the end of a fund's life, whether the asset will remain competitive without significant reinvestment, and how concentrated demand may affect liquidity. Some underwriting frameworks may address this

uncertainty by placing limited or even zero terminal value on certain assets, shifting the investment case toward contracted cash flows rather than resale value.

Summary

Data centers have moved from a niche investment opportunity to a critical piece of economic and investment infrastructure. The scale of AI-related capital spending has pushed the sector into the macro conversation, while the underlying assets remain grounded in practical requirements like power delivery, reliability, and long-lived physical systems. For investors, this combination means data centers are influenced both by top-down forces such as technology adoption and economic growth, and by bottom-up execution factors that resemble traditional infrastructure and real estate.

At the asset level, outcomes depend less on abstract demand for "AI" and more on concrete questions: Can power be secured on acceptable terms? Can capacity be delivered on time? Are contracts structured to support durable cash flows? Facilities that can reliably provide usable, powered capacity in constrained markets often capture greater value, while delays or design mismatches can materially impair returns. This creates meaningful dispersion across projects, markets, and managers.

Investors can gain exposure to data centers through multiple channels, primarily including listed REITs, private real estate funds, and private infrastructure vehicles. Each comes with different trade-offs between liquidity, volatility, and control over development and operating risk. As the sector continues to grow, disciplined underwriting around power certainty, contract durability, and execution capability may matter more than broad enthusiasm for AI-driven demand.

Ultimately, data centers are best understood not as a speculative technology bet, but as enabling infrastructure for a digital and AI-intensive economy. Investors who approach the space with a clear view of its physical constraints, cash flow mechanics, and risk drivers are better positioned to separate durable opportunities from those that may rely too heavily on optimistic assumptions.

End Notes

- ¹ Harvard economist Jason Furman found that investment in information-processing equipment and software (much of it tied to AI and data centers) was about 4% of US GDP in the first half of 2025 but accounted for roughly 92% of GDP growth in that period.
- ² Source: MUFG Americas, The AI Chart Weekly, "Hyperscalers' Capex Above \$600 Bn in 2026." The big five hyperscalers are Amazon, Microsoft, Alphabet, Meta, and Oracle.
- ³ Source: "Big Tech's AI Push Is Costing a Lot More Than the Moon Landing", Wall Street Journal, Meghan Bobrowsky, Drew An-Pham, and Alana Pipe, February 7, 2026.
- ⁴ Source: Congressional Research Service, M. et al., "Data Centers and Their Energy Consumption: Frequently Asked Questions." updated January 23, 2026.
- ⁵ Power is measured in kilowatts, megawatts, and increasingly gigawatts (1,000 kilowatts equal one megawatt, and 1,000 megawatts equal one gigawatt). Ten kilowatts is roughly equivalent to the continuous electricity demand of 6 to 10 average US homes. Ten megawatts can approximate the average load of 7,000 to 10,000 homes, and a 100-megawatt campus can approach the electricity demand of a small city. At the largest scale, 10 gigawatts would represent enough capacity to serve several million households and is comparable to the output of multiple utility-scale power plants.
- ⁶ For example, reports describe xAI's Colossus project, based in the Memphis, Tennessee area, as a major AI-focused data center that has expanded quickly and attracted scrutiny over power use and environmental concerns. It is being built to support large-scale AI training workloads and could reach multiple gigawatts of power capacity, putting it among the most energy-intensive AI facilities in the world.
- ⁷ Source: KKR, Waldemar Szlezak, et al., "Beyond the Bubble: Why We Think AI Infrastructure Will Compound Long after the Hype." November 2025.
- ⁸ Source: McKinsey & Company. "Data Centers and AI" (Private Capital Practice). September 2024.
- ⁹ Source: Brookfield. "Building the Backbone of AI." August 2025.
- ¹⁰ Source: JLL. "North America Data Center Report: Midyear 2025." Accessed January 2026.
- ¹¹ Source: KKR, Waldemar Szlezak, et al., "Beyond the Bubble: Why We Think AI Infrastructure Will Compound Long after the Hype." November 2025.
- ¹² Source: Brookings Institution, N. Turner Lee et al., "The Future of Data Centers." Brookings Institution, November 5, 2025. A Graphics Processing Unit is a specialized electronic circuit originally designed to speed up the creation of images and videos. Its ability to perform vast numbers of calculations rapidly has led to its adoption in diverse fields, including artificial intelligence and scientific computing.
- ¹³ Source: McKinsey & Company. "How Data Centers and the Energy Sector Can Sate AI's Hunger for Power." September 17, 2024.
- ¹⁴ Source: Brookings Institution, N. Turner Lee et al., "The Future of Data Centers." Brookings Institution, November 5, 2025.
- ¹⁵ Source: McKinsey & Company. "Data Centers and AI" (Private Capital Practice). September 2024. Operators may also charge rent for space, include escalation clauses, and seek various structures for power cost recovery.
- ¹⁶ Source: Congressional Research Service, M. et al., "Data Centers and Their Energy Consumption: Frequently Asked Questions." updated January 23, 2026.
- ¹⁷ As an interesting aside, SpaceX, Google, and others are exploring the idea of putting data centers in space to reduce land, water, and local power constraints. The concept of orbital data centers is still highly experimental because launch costs, maintenance, reliability, and data connection limits currently make it far more difficult and expensive than building on Earth.
- ¹⁸ Source: CBRE Research, "North America Data Center Trends H12025," September 8, 2025.

Important Information

This report (the “report”) has been prepared for the sole benefit of the intended recipient (the “recipient”).

Significant events may occur (or have occurred) after the date of this report, and it is not our function or responsibility to update this report. The information contained herein, including any opinions or recommendations, represents our good faith views as of the date of this report and is subject to change at any time. All investments involve risk, and there can be no guarantee that the strategies, tactics, and methods discussed here will be successful.

The information used to prepare this report may have been obtained from investment managers, custodians, and other external sources. Some of this report may have been produced with the assistance of artificial intelligence (“AI”) technology. While we have exercised reasonable care in preparing this report, we cannot guarantee the accuracy, adequacy, validity, reliability, availability, or completeness of any information contained herein, whether obtained externally or produced by the AI.

The recipient should be aware that this report may include AI-generated content that may not have considered all risk factors. The recipient is advised to consult with their Meketa advisor or another professional advisor before making any financial decisions or taking any action based on the content of this report. We believe the information to be factual and up to date but do not assume any responsibility for errors or omissions in the content produced. Under no circumstances shall we be liable for any special, direct, indirect, consequential, or incidental damages or any damages whatsoever, whether in an action of contract, negligence, or other tort, arising out of or in connection with the use of this content. It is important for the recipient to critically evaluate the information provided.

Certain information contained in this report may constitute “forward- looking statements,” which can be identified by the use of terminology such as “may,” “will,” “should,” “expect,” “aim,” “anticipate,” “target,” “project,” “estimate,” “intend,” “continue,” or “believe,” or the negatives thereof or other variations thereon or comparable terminology. Any forward-looking statements, forecasts, projections, valuations, or results in this report are based upon current assumptions. Changes to any assumptions may have a material impact on forward-looking statements, forecasts, projections, valuations, or results. Actual results may therefore be materially different from any forecasts, projections, valuations, or results in this report.

Performance data contained herein represent past performance. Past performance is no guarantee of future results.