Data Centres: The Leading Questions

Knight Frank

2026

Tackling the principal topics defining the global data centre landscape

knightfrank.co.uk/research

Foreword

STEPHEN BEARD
GLOBAL HEAD OF DATA CENTRES DEVELOPMENT & INVESTMENT

The global data centre market has entered its most aggressive investment and build-out phase in its history. Demand is being reshaped by the AI super cycle, traditional hyperscale cloud providers, and escalating power constraints that increasingly dictate market behaviour. Investors are navigating an environment of unprecedented capital influx, rising tenant concentration risk, and a widening divergence between power-rich and power-constrained regions.

The global acceleration of artificial intelligence is reshaping the digital and physical landscape at a pace, driving profound change across the data centre industry and the wider energy ecosystem. AI is no longer a theoretical technology, it is a present and immediate force, catalysing a historic build out of digital infrastructure. Over the next five years, AI related demand will require as much as $1.6 trillion in global investment, transforming data centres into one of the most capital-intensive asset classes in the world. What was once a sector defined by location, connectivity and occupier demand is now being led by a single overriding variable: power.

Global data centre capacity is projected to nearly double by 2028, rising from 62GW to more than 110GW, with AI specific workloads expanding even faster. Yet this rapid growth is running headlong into structural constraints. Power grids across key markets are stretched to their limits, with connection queues extending close to a decade in some regions. These pressures have intensified vacancy shortages, pushed rents upward and elevated power availability to the defining factor in site selection. In markets such as Frankfurt, Ashburn, Singapore and London, vacancy rates have fallen to near historic lows, signalling an environment where deliverable megawatts, not land or capital, determine growth potential.

At the same time, hyperscale cloud operators are investing on a scale unmatched in corporate history. By 2026, combined capital expenditure from Microsoft, Amazon Web Services, Google and Meta is expected to exceed $650 billion as they race to meet surging AI compute requirements. Hyperscaler dominance of global take up, now approaching 70%, is not only reshaping market dynamics but also increasing exposure to the capital cycles and strategic priorities of a small cohort of tenants. This concentration is prompting both investors and developers to rethink risk, delivery timelines and long-term planning.

The power challenge is also accelerating a shift toward alternative energy strategies. Major technology companies are increasingly turning to nuclear and private power procurement, while governments are tightening regulatory frameworks to balance grid pressures with economic opportunity. Emerging AI focused infrastructure operators are scaling rapidly in response, but often through highly leveraged models that introduce new layers of risk to the ecosystem.

Ultimately, the period ahead will be defined by two fundamental questions: whether AI workloads can generate sustainable long-term revenues, and whether global power infrastructure can expand quickly enough to support growth. Failure on either front carries significant implications, from overbuilding to systemic supply shortages. Yet amid this uncertainty lies one clear conclusion: the AI era is driving an irreversible transformation of global data centre development and the energy systems that support it.

This report provides a comprehensive assessment of these dynamics, offering clarity at a moment when the intersection of technology, infrastructure and energy is becoming one of the defining economic challenges of our time.

Contents

Artificial Intelligence: A trillion dollar question

Often dubbed the modern industrial revolution, the emergence of artificial intelligence (AI) has supercharged the global data centre sector from a robust investment opportunity into a capital-agnostic arms race, where competitive dynamics are driven by speed, scale, and ambition, not quarterly budgets.

THE DEMAND SHOCK

Knight Frank forecasts global data centre capacity to expand from 62 gigawatts (GW) in 2025 to over 110GW by the end of 2028. AI-related volumes, both through self-build development and colocation lease agreements, could expand from c.8GW to c.27GW by 2028, representing a rise from 12.9% of workloads to 24.5% over the forecast period.

Over the course of 2026, it is expected that capital expenditure from US 'Big Tech' – Microsoft, Google, Amazon Web Services (AWS), and Meta – could surpass $650 billion, up from a combined $376 billion in 2025. This growth is being driven primarily by AWS and Google, which have upped annual capital expenditure forecasts to $200 billion and $185 billion, respectively. Microsoft, on the other hand, has not provided any full-year forecasts, only stating that its quarterly capital expenditure is expected to 'decrease sequentially' following a company record period during the quarter ended December 2025, where capital expenditure jumped 54.1% quarter-on-quarter. Despite this, however, Microsoft is still expected to deploy between $110-$130 billion during the year.

The current demand cycle has helped net colocation take-up reach 15.8GW signed in 2025, 37% of which is driven by AI-related demand. A little under 6GW of new AI-related colocation lease agreements were signed during the year, three times the volume signed in 2024.

In many markets, vacancy is hitting record lows. Global wholesale colocation vacancy has dipped to 8.1%, down from 9.1% at the end of 2024, with primary-market vacancy tightening further. Across Europe, the FLAP-D markets hit an overall vacancy of less-than 4%, with London at 3.1% and Frankfurt at <1%, whilst in North America and Asia-Pacific, Ashburn vacancy has dropped to <1% and Singapore hits 2.2%.

POWER-FIRST, NOT LOCATION-FIRST

The International Energy Agency (IEA) "Energy & AI" analysis emphasises that access to reliable, scalable power is the real bottleneck to "intelligence at scale", not land or GPUs. As power

> "Over the course of 2026, it is expected that capital expenditure from US 'Big Tech' – Microsoft, Google, Amazon Web Services, and Meta – could surpass $650 billion."

AI supply growth
Live GW

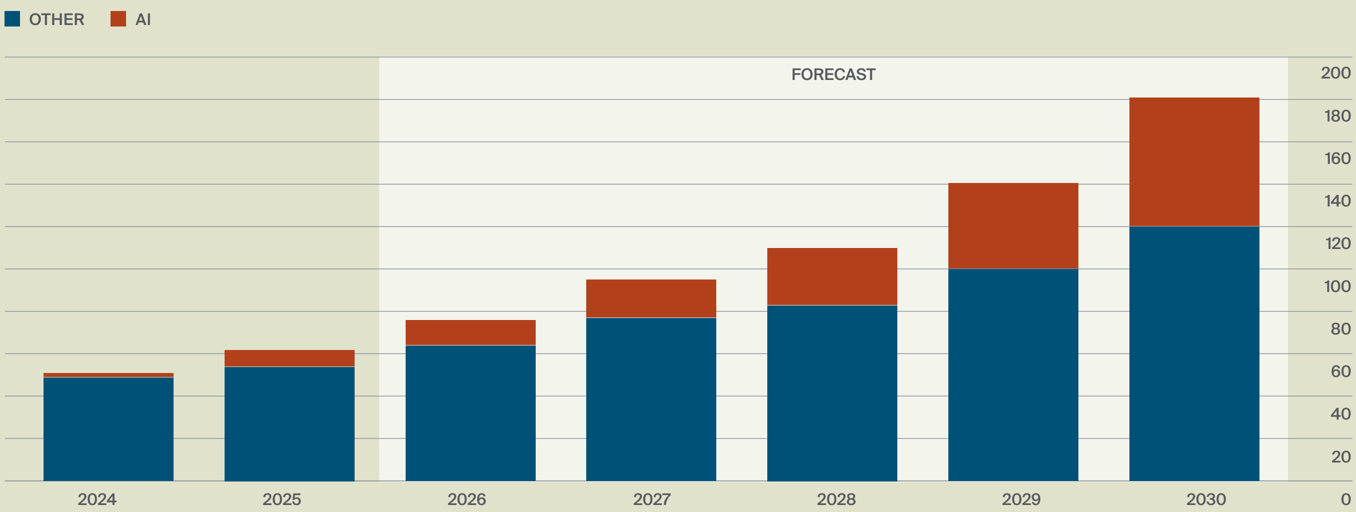

Source: Knight Frank Insight

generation capabilities ramp up across global markets, the widely acknowledged constraint is the transmission grid. Decades of underinvestment, have created a structural deficit causing multi-year connection queues, with key areas, such as London, approaching close to a 10-year connection wait time. Grid capacity across core data centre metros is predominantly reserved through to 2030, prompting long-term pre-leasing tendencies and a transition to on-site generation.

To correct this deficit, significant investment is required across most global markets. BloombergNEF estimates requirements in the United States may reach $150 billion by 2030, whilst the UK faces an estimated $107 billion cost to move renewable power generated in the north to core demand regions in the south. China, which operates the worlds largest grid, could require as much as $3.8 trillion by 2050, while for Australia the figure is roughly $300 billion through to 2050.

What this means for site selection is that deliverable power has become the prime variable. Investors are prioritising powered land, transmission line-proximate parcels, and campus-scale plots capable of phased power delivery, in the hundreds of megawatts, to host the c.20-30GW compute capacity to be delivered in support of AI over the coming 36 months.

A TRILLION-DOLLAR QUESTION

By 2030, global AI-related data centre capacity deployment volumes could hit 60GW, up from 8GW in 2025. This increase of 52GW will come at a development cost of between $676-$780 billion, whilst tenants are likely to spend the same or more to fit out space with necessary IT. The infrastructure to support the AI build-out over the coming 5 years will require an infrastructure investment of $1.4-$1.6 trillion.

Global data centre spending is projected to reach $2.8-$3.2 trillion over the next five years, with much of the new capacity already pre-leased to hyperscalers, reducing vacancy risk but heightening single-tenant exposure. The partnership between Oracle and OpenAI, announced in September 2025, quickly became the most high-profile AI infrastructure program. The agreement, a $300 billion, five-year cloud computing contract beginning in 2027 that would make Oracle the primary provider of compute for OpenAI and position it alongside Microsoft, AWS, and Google as a hyperscaler capable of delivering next-gen AI compute infrastructure.

> "The infrastructure necessary to support the AI build-out over the coming 5 years will require an investment of $1.4-$1.6 trillion."

'Big Tech' annual CapEx

$ million

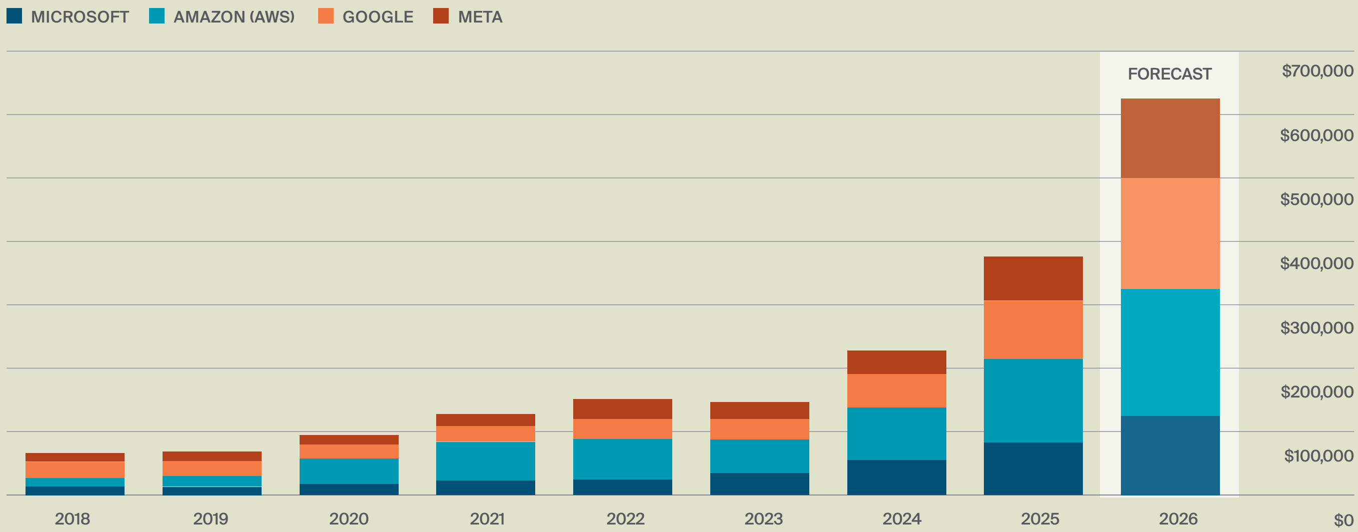

Source: Knight Frank Insight

Yet, OpenAI's annualised revenue of $10 billion (as of Q2 2025) was overshadowed by the projected $60 billion yearly cloud bill, potentially forcing it to consider major cost cuts or fresh capital raises. Credit agencies warned the deal could strain Oracle's balance sheet, with Moody's highlighting leverage and counterparty risk. By March 2026, the companies cancelled an initial 600 megawatt (MW) Abilene expansion, citing funding complications and volatile demand forecasts.

Going forward, OpenAI has $600 billion in projected cloud compute spend by 2030, which was recently cut back from an earlier pledge to spend $1.4 trillion by 2033. Oracle is set to raise $50 billion in debt and equity to fund the OpenAI build-out, with large-scale layoffs expected to help cover costs. Similarly, fellow Stargate backer SoftBank is turning to debt to fund OpenAI's expansion, seeking $40 billion in loans.

In other news, Anthropic announced a $50 billion investment in US AI infrastructure to build data centres in Texas and New York with Fluidstack. OpenAI completed a $40 billion funding round, the largest ever by a private tech company, whilst also completing a $38 billion deal with AWS, providing it with access to AWS' infrastructure to scale its core AI workloads. Meta & Blue Owl Capital announced a $27 billion joint-venture for the development of the Hyperion AI campus. Nebius announced GPU contracts with Microsoft and Meta for $19.4 billion and $3 billion, respectively, whilst similarly receiving a $2 billion investment from Nvidia to help fund its AI data centre expansion. Similarly, CoreWeave announced a

Key AI infrastructure agreements

Year	Quarter	Deal	Parties	Value ($ million)	Notes
2025	Q1	Stargate Project	OpenAI, SoftBank, Oracle & MGX	$500,000	Four-year investment plan to develop new AI infrastructure
2025	Q3	AI Compute Contract	Oracle & OpenAI	$300,000	Five-year AI compute contract beginning in 2027
2026	Q1	Funding Round	OpenAI	$122,000	FUNDING ROUND VALUED OPENAI AT $852 BILLION
2025	Q4	AI Infrastructure Announcement	Anthropic, Fluidstack	$50,000	$50 billion investment in US AI Infrastructure via Fluidstack
2025	Q4	Aligned Data Centers	AIP, MGX & GIP	$40,000	The largest data centre acquisition deal ever announced
2025	Q4	AI Infrastructure Contract	OpenAI & AWS	$38,000	Provides OpenAI immediate & increasing access to AWS infrastructure
2025	Q4	Hyperion Project	Meta & Blue Owl Capital	$27,000	JV between Meta (20%) and Blue Owl Capital (80%)
2025	Q4	GPU Contract	Nebius & Microsoft	$19,400	AI infrastructure contract through to 2031
2025	Q2	Strategic Investment	Meta & Scale AI	$14,800	Meta acquired a 49% stake in Scale AI
2025	Q3	GPU Contract	Coreweave & Meta	$14,200	AI infrastructure contract through to 2031
2025	Q2	Lancaster County	CoreWeave	$6,000	AI data centre development in Lancaster, Pennsylvania
2025	Q4	GPU Contract	Nebius & Meta	$3,000	Five-year, $3 billion AI infrastructure deal
2026	Q1	Strategic Investment	Nebius & Nvidia	$2,000	Funding to accelerate AI data centre expansion

Key AI infrastructure projects

Operator	Project	Type	Capacity	Location	RFS Timeline	Notes
Meta	Richland Parish	Self-build	2GW (scalable to 5GW)	United States	2027 - 2030	$27 billion JV with Blue Owl Capital (80%)
Core42 & Oracle	Stargate Abu Dhabi	Build-to-suit	1GW	United Arab Emirates	2026 - 2029	200MW due in September 2026, 12 months after construction started
Digital Connexion	Andhra Pradesh	Wholesale	1GW	India	2027 - 2030	$11 billion, 400 acre AI-ready data centre campus
AWS	Razor5	Self-build	928MW	United States	2025 - 2028	$11 billion AI training campus (Project Rainer) for Anthropic
Crusoe	Stargate Abilene I & II	Build-to-suit	806MW	United States	2025 - 2026	Build-to-suit for Oracle, with capacity consumed by OpenAI
NEXTDC	S7	Wholesale	800MW	Australia	2027 - 2030	OpenAI partnered with NEXTDC to build a AUD$7 billion computing cluster
SoftBank	Lordstown	Self-build	767MW	United States	2027 - 2031	Site acquired for $375 million and will be a part of OpenAI-Oracle initiative
Google	Project Mica	Self-build	538MW	United States	2027 - 2028	Construction began in Q1 2026 with 18-24 month development timeline
Galaxy Digital	Helios Campus	Wholesale	393MW	United States	2027 - 2028	CoreWeave signed a 15-year lease for 460MW of critical capacity
Microsoft	Fairwater I	Self-build	338MW	United States	2025 - 2026	$3.3 billion investment, plus a further $4 billion over the next three year
HUMAIN	SPARK	Build-to-suit	300MW (scalable to 1.2GW)	Saudi Arabia	2026 - 2034	HUMAIN aims to reach 1.9GW by 2030, and 6.6GW by 2034
xAI	Colossus 1	Self-build	212MW	United States	2024 - 2025	Initial 106MW built in 122 days, with phase 2 constructed in 92 days
xAI	Colossus 2	Self-build	1.2GW	United States	2025 - 2026	Phase one constructed in 168 days, with full fit out expected during 2026

Source: Knight Frank Insight

$6 billion AI data centre project in Lancaster, Pennsylvania.

Regarding the cancelled Abilene expansion, and another revealing signal of the instability of the Oracle & OpenAI partnership, came from Nvidia. The company paid a deposit of $150 million to Crusoe to secure the site, ensuring that its hardware is deployed at the site instead of a competitor. Both Meta and Microsoft are considering occupying the unused space.

SOWING THE SEEDS FOR AN OVERBUILD?

With global data centre capacity set to triple by 2030, and AI-based capacity to grow by over ten times, the question remains over whether the current build-out roadmap is steering towards an over-build. At present, there exists an undersupply of data centre capacity, driving record levels of pre-leasing activity from hyperscalers, with multi-year roadmaps, public balance sheets, and a known desire for geographic redundancy underpinning landlord visibility.

Over the next 24-36 months, power availability, not capital, will cap supply, keeping core fundamentals strong across Tier I & Tier II markets and pushing additional development to new power-rich regions.

> "Over the next 24-36 months, power availability, not capital, will cap supply, keeping core fundamentals strong across Tier I & Tier II markets and pushing additional development to new power-rich regions."

However, beyond 2028, two potential swing factors will dominate the direction of travel: Inference monetisation, and whether AI can move from promise to profit at scale, and grid build-out, and whether new generation capacity and transmission networks can catch up. What this argues for is a disciplined, phased build-out with flexible designs that can switch between high-density AI and lower-density cloud, and contract structures that share delivery risk.

WHAT DOES THIS MEAN FOR INTERESTED PARTIES?

For developers and investors, pursue 'power-first' land strategies, prioritising parcels with realistic pre-2028 energisation paths and proximity to high-voltage transmission and gas networks. Upon which, consider AI-ready shells, with allowances for liquid cooling fitouts and assuming rack densities of 50-100+kW.

For landlords and REITs, lean into deliverable megawatts. The pricing premium for AI-capable land, with near-term power, is set to remain, with rental growth and multi-phase pre-leasing to remain the norm. Balance tenant concentration risk, stagger expiry schedules, and consider heightened underwriting scrutiny given the sector's lack of investment-grade credit tenants.

For occupiers, secure power early. With new-lease rental rates expected to rise by 8-12% annually, take-on phased pre-lease options, particularly in power-scarce regions, rather than betting on a late-cycle market.

FINAL TAKE

The data centre sector is no longer a peripheral enabler of technology, but the industrial backbone of the next economic cycle. Real estate investors, developers, and operators now sit at the centre of this transformation, where the ability to secure and deliver power, navigate government policy, and build adaptable infrastructure will determine who thrives. The winners will be those who pair disciplined capital deployment with speed, scale, and sustainability.

Highest operable server density

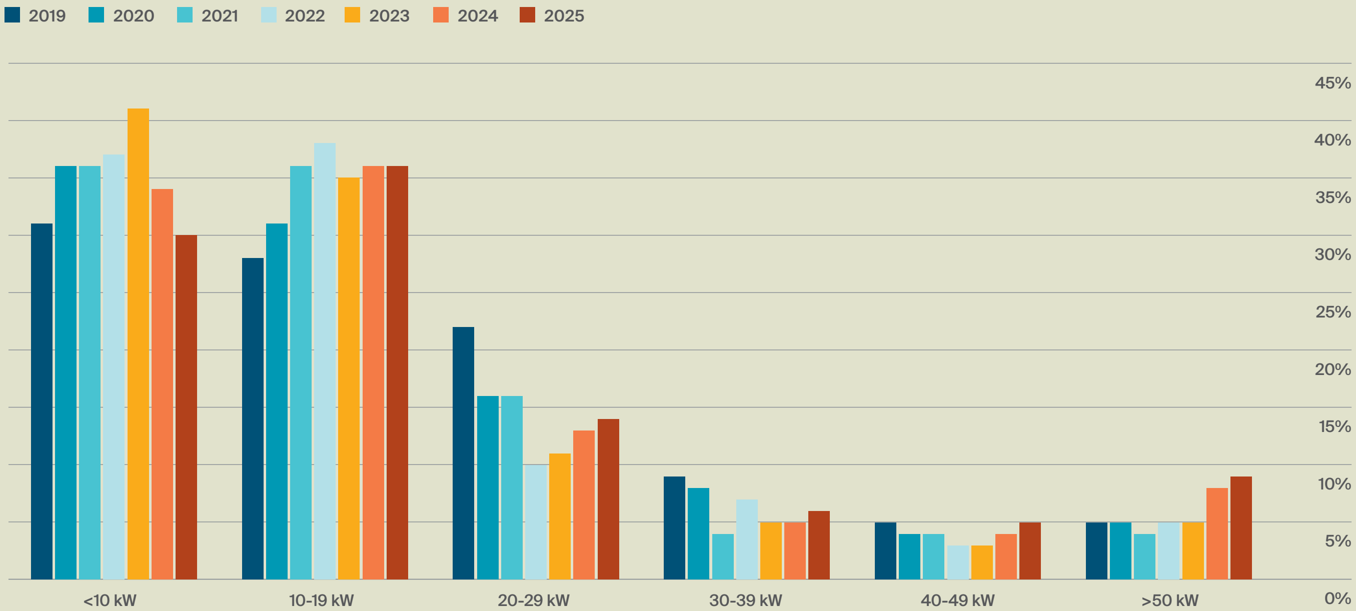

Source: Knight Frank Insight, Uptime Institute (2019-2024)

Neoclouds: What they are, why they exist, and why they matter

THE DEFINITION?
Neocloud is not just another buzzword in an industry already overflowing with them. It is structural shift in how organisations think about high performance computing (HPC) and AI era infrastructure. Neoclouds can be defined as specialised, AI-focused cloud providers built to deliver high-performance GPU infrastructure optimised for workloads such as AI training, inference and HPC. They differ from traditional hyperscalers by offering faster provisioning, more transparent per-GPU pricing, and environments tailor-made for compute-intensive demands. That is a useful starting point, but to really understand why neoclouds have emerged, it is necessary to look at the evolution of cloud itself.

For more than a decade, cloud adoption has been defined by hyperscalers, predominantly by the big five, AWS, Microsoft, Google, Meta and Oracle. These hyperscalers have shaped the modern expectations of cloud computing, which are elasticity, global reach, managed services and the ability to scale without owning or maintaining physical infrastructure. This model is still hugely valuable. Most organisations need reliable compute, storage, networking, identity and access management and platform services that can support common workloads and hyperscalers deliver this by providing broad, general-purpose cloud. But the rise of AI, in particular deep learning and generative AI, identified a gap in the market. Training large models requires many GPUs working together efficiently and inference at scale needs predictable performance and cost efficiency. This gap is exactly where neoclouds have positioned.

> "Crusoe, Nebius, Lambda Labs and CoreWeave were among the top companies to recognise the opportunity created by constrained GPU supply."

THE PLAYERS?
Crusoe, Nebius, Lambda and CoreWeave were among the companies to recognise the opportunity created by constrained GPU supply, positioning business services to fill that gap well before the recent acceleration in AI demand. Today, the ecosystem has broadened significantly. We estimate there are now close to 200 neocloud operators globally, with around $10 billion in investment allocated to them last year, according to S&P. That figure is likely to grow as the AI infrastructure market expands. Neocloud revenue is projected to increase from $23.9 billion in 2025 to $179.1 billion by the end of the decade, based on data from Synergy Research Group.

> "Training large models requires many GPUs working together efficiently, while inference at scale demands predictable performance and cost efficiency. This gap is exactly where neoclouds have positioned themselves."

Market size forecast
$ billion

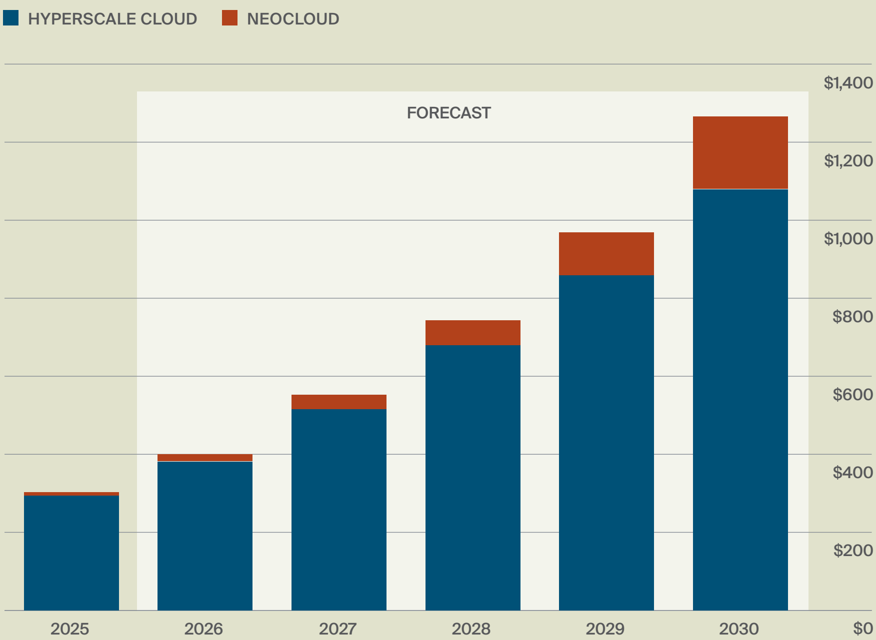

Source: Knight Frank Insight, Structure Research

“We estimate there are now close to 200 neocloud operators globally, with around $10 billion in investment allocated to them last year.”

Yet not all neocloud providers operate in the same way and their strategies create different growth profiles. Some have adopted capital intensive models, building and operating their own data centre infrastructure. CoreWeave, for example, reported an estimated 850MW of connected power at the end of 2025, with more than 3.1GW in its pipeline. Its capital expenditure (CapEx) trajectory reflects this, growing 18.5% year on year, with 2025 CapEx reaching $10.3 billion and forecast to more than double in 2026. Nscale has followed a similar path. With facilities across the UK and Nordics, it has scaled capacity to nearly 700MW and has a pipeline of 1.3GW. Others rely more heavily on leasing colocation capacity. Leasing agreements for AI/HPC-focused colocation reached nearly 4.1GW in 2025, up from 2.2GW in 2024. North America and APAC were the main hotspots, accounting for 45% and 39% of 2025 leasing activity, respectively.

How the GPU supply chains are structured is another differentiator. CoreWeave, Nebius, Lambda, Crusoe and Nscale, to name a few, all have partnerships with Nvidia, giving them preferential access to next-generation GPUs, early hardware roadmaps and, in some cases, capital support. The most recent of these agreements was announced in March this year, when Nvidia committed $2 billion to a strategic partnership with Nebius. Beyond Nvidia, alternative GPU partnerships remain more limited and tactical. Lambda and CoreWeave have both introduced AMD Instinct MI300-based instances, primarily to diversify supply and serve specific training or inference workloads, but these sit alongside, rather than replace, their Nvidia deployments.

THE RISKS?

Neoclouds are increasingly funding expansion with debt. At CoreWeave, total liabilities increased to $45.97 billion at the end of 2025, up from $16.53 billion a year earlier. This sits against $49.3 billion of total assets, leaving a comparatively thin equity buffer for a business still operating in hyper-growth mode. While this is not “bad debt” per se, it reflects a deliberate capital strategy. Neoclouds are deploying significant capital upfront, with returns monetised over multi-year customer contracts, making their financial profiles closer to infrastructure finance than to traditional cloud economics.

Early advantage, in particular the speed to access GPUs, may not be long lived. S&P Global Market Intelligence has noted that improving chip availability could erode one of the sector’s key differentiators. As supply normalises, hyperscalers may increasingly internalise AI workloads or compete more aggressively on price and performance, leveraging balance sheets, integrated platforms and global footprints.

These pressures are likely to drive consolidation as the market matures. Larger neoclouds as well as hyperscalers seeking to expand AI infrastructure capacity, may look to acquire smaller GPU cloud providers that are unable to sustain the capital intensity of growth on their own.

Power: Capacity, constraint, and cost

DOES MORE DATA MEAN MORE POWER?

As Microsoft's chief executive Satya Nadella put it in a podcast interview last year, "the biggest issue we are now having is not a compute glut, but it's power." AI and cloud computing are pushing energy, and particularly electricity, demand to record highs. Forecasts from the IEA suggest data centres could almost double their electricity consumption between 2025 and 2030, reaching over 945TWh, which is just under 3% of total global electricity demand. The headline risk is not simply whether the world can generate enough electricity. It is whether power systems, networks, interconnections and market structures can deliver it to the right places, at the right time, at the speed the digital economy now demands. That mismatch is already influencing policy, investment priorities and the economics of building new digital infrastructure.

WHAT ROLE CAN NUCLEAR POWER PLAY?

Nuclear energy is well positioned to support a more electrified, digital economy because it provides reliable power 24 hours a day, 365 days a year, with relatively low operational emissions. Today, almost 380GW of nuclear capacity is operational globally, concentrated in a relatively small group of countries. The United States is the largest producer with 97GW. France follows with 63GW and generates almost 70% of its electricity from nuclear, one of the highest shares globally. Russia operates 28GW, while in the Asia-Pacific region, South Korea (25.6GW) and Japan (12.6GW) remain key generators.

Nuclear hotspots

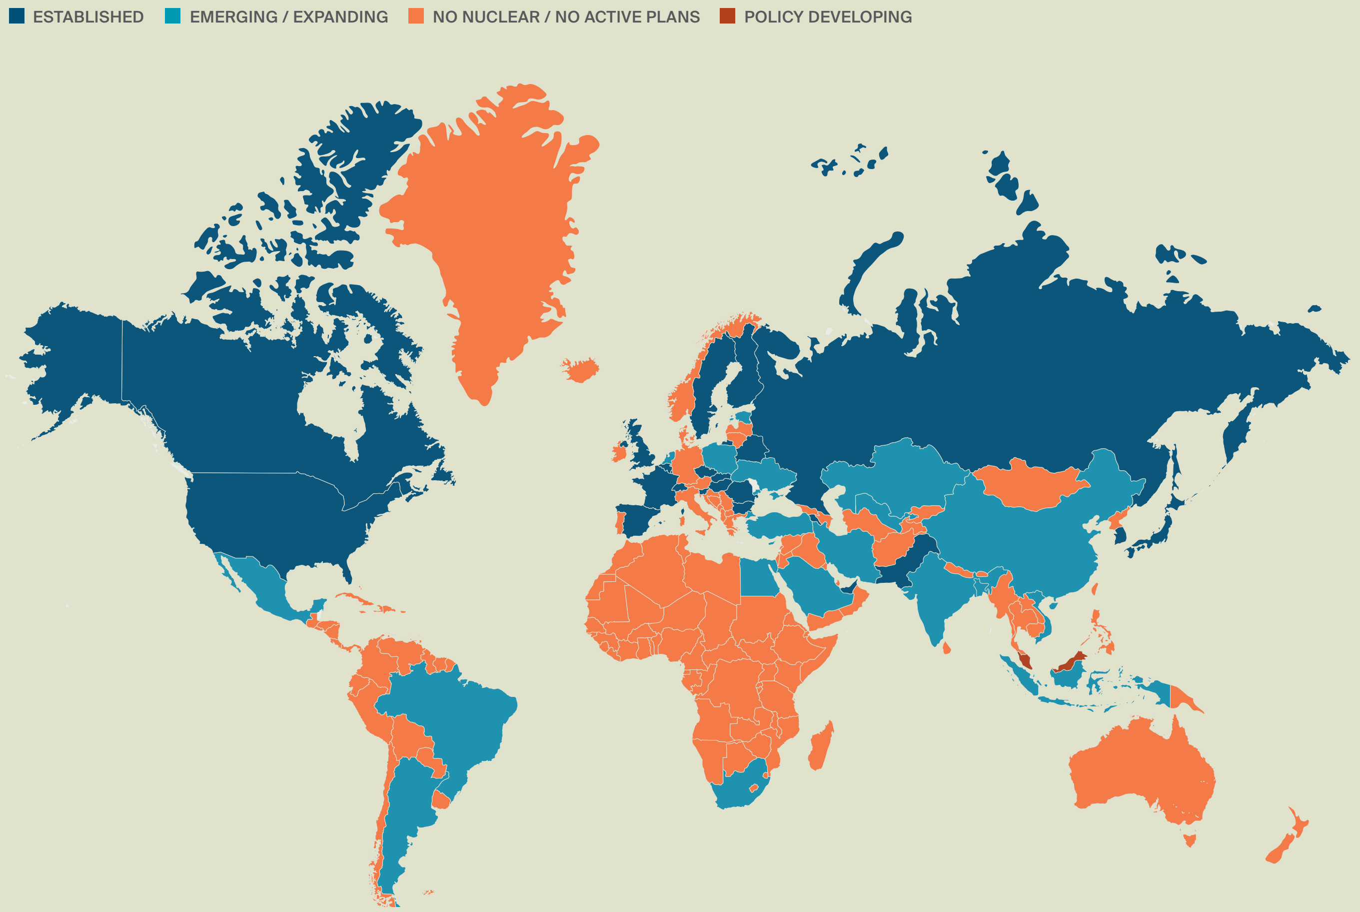

Source: Knight Frank Insight

However, meeting future demand will require significant expansion. Goldman Sachs estimates that AI-driven electricity needs alone could require up to 90GW of additional nuclear capacity. Governments are also signalling broader support for nuclear power, with the strongest near-term expansion expected in China, which has 257.1GW planned or under construction. India follows with 45.2GW, Russia with 26.2GW and a growing group of emerging European markets including Poland (13.8GW) and Ukraine (13.3GW). Across the wider APAC region, installed nuclear capacity is projected to triple from 112.2GW to 462.0GW in the coming years.

Countries are also reassessing earlier limitations on nuclear energy. Switzerland is reconsidering its long-standing position, with the government presenting draft legislation in August last year aimed at ending its ban on constructing new nuclear power plants. In the Middle East, Saudi Arabia has made progress on the technical specifications for its first nuclear power plant, an important milestone since the Energy Minister first signalled the country's interest in nuclear development in 2023.

WHY ARE HYPERSCALERS MOVING UPSTREAM?

Large technology companies are no longer passive buyers of electricity. Hyperscalers including Microsoft, Amazon, Google and Meta are now directly investing in nuclear energy through long-term power purchase agreements (PPAs).

Microsoft has signed a 20-year, 835MW nuclear PPA to restart and offtake 100% of the output from the Three Mile Island facility. Amazon has finalised an $18 billion PPA to secure up to 1.9GW with Talen Energy. Google has contracted up to 500MW of Small Modular Reactors (SMRs) capacity through its agreement with Kairos Power, while Meta has issued nuclear procurement requests and signed long-term nuclear deals exceeding 6.6GW, including a 20-year PPA with Constellation's Clinton Clean Energy Center in Illinois. Collectively, these agreements form a major share of the more than 55.9 GW of new PPAs signed in the last year, albeit largely concentrated in the United States, where there is an established market.

> "Across major markets, investment in renewables, nuclear and storage is occurring far faster than expansion of transmission and distribution networks."

IS THE CONSTRAINT THE GRID RATHER THAN GENERATION?

While generation is growing, the constraint is increasingly the grid, a bottleneck that is now widely acknowledged. Across major markets, investment in renewables, nuclear and storage is occurring far faster than expansion of transmission and distribution networks. The result is congestion and multi-year interconnection queues. According to Ember, developers now face average grid-connection waits of roughly nine years in the UK, seven years in Germany, five in Spain, three in Italy, and two in Norway.

> "Hyperscalers including Microsoft, Amazon, Google and Meta are now directly investing in nuclear energy through long-term power purchase agreements (PPAs)."

This imbalance is visible in global investment patterns. While annual spending on electricity generation is estimated at $1 trillion, grids receive only about $400 billion, according to the IEA. Although this gap is beginning to narrow, with grid investment expected to double to $800 billion by 2030, decades of underinvestment have created a structural deficit that cannot be closed quickly. The scale of reinforcement required varies significantly by country.

> “Developers now believe that around a third of their data centres will rely on 100% onsite generation by 2030, with almost half (44%) by 2035.”

BloombergNEF estimates that Germany may need around $89 billion in transmission and $50 billion in distribution upgrades by 2032. China, which operates the world’s largest grid, could require as much as $3.8 trillion by 2050. Italy may need $12 billion to strengthen its north-south transmission corridor. In the US, grid investment requirements may reach $150 billion by 2030, while for Australia the figure is roughly $300 billion through to 2050. The UK faces an estimated $107 billion requirement to move renewable power generated in the north to demand centres in the south.

WHO PAYS FOR THE GRID?

In last year’s report, we highlighted how operators in Dublin were already exploring alternatives such as onsite gas generation, private transmission builds and connections to their own renewable assets to overcome long connection delays. Since then, policy frameworks have tightened. The four-year de facto moratorium on new data centre grid connections in Dublin was lifted in December 2025, but only under strict new conditions. Developers may only move forward if they install on-site generation capable of covering 100% of their requested grid-import capacity and demonstrate that 80% of their annual energy consumption will come from newly-built renewable projects located in Ireland. Elsewhere, in the US, a survey by Bloom Energy of hyperscalers and colocation providers shows a clear shift in developers’ expectations for onsite power deployment. Developers now believe that around a third of their data centres will rely on 100% onsite generation by 2030, with almost half (44%) by 2035.

Governments and utilities globally are adopting a similar stance. Increasingly, they are encouraging (or requiring) large power users to coinvest in the grid as part of broader system reinforcement plans. In the United States, the federal government has signalled that data centre operators must directly pay for the grid infrastructure they require, rather than shifting costs to households. This was formalised through the Ratepayer Protection Pledge, announced by the White House in March and signed by major hyperscalers including Amazon, Google, Meta, Microsoft, OpenAI, Oracle and xAI. Under the pledge, companies must fund the infrastructure needed to support their facilities. At the state level, aligned expectations are taking shape. In New York, for example, the Energize NY Development initiative will require large power users, especially data centres that do not deliver sufficient local economic value, to either self-generate electricity or pay higher rates to fund grid investment.

HOW ARE DIGITAL AND ENERGY INFRASTRUCTURE BECOMING INTERTWINED?

Access to power, rather than demand for compute, is increasingly determining where and how data centres can be built. As grid connections become slower, more constrained and more costly, developers are rethinking location, scale and design, with greater reliance on onsite generation, private networks and dedicated renewable or nuclear supply. Investors are responding by deploying capital across digital and energy infrastructure in tandem, recognising that secure, timely power has become a critical driver of value.

Global grid investment forecast
$ billion

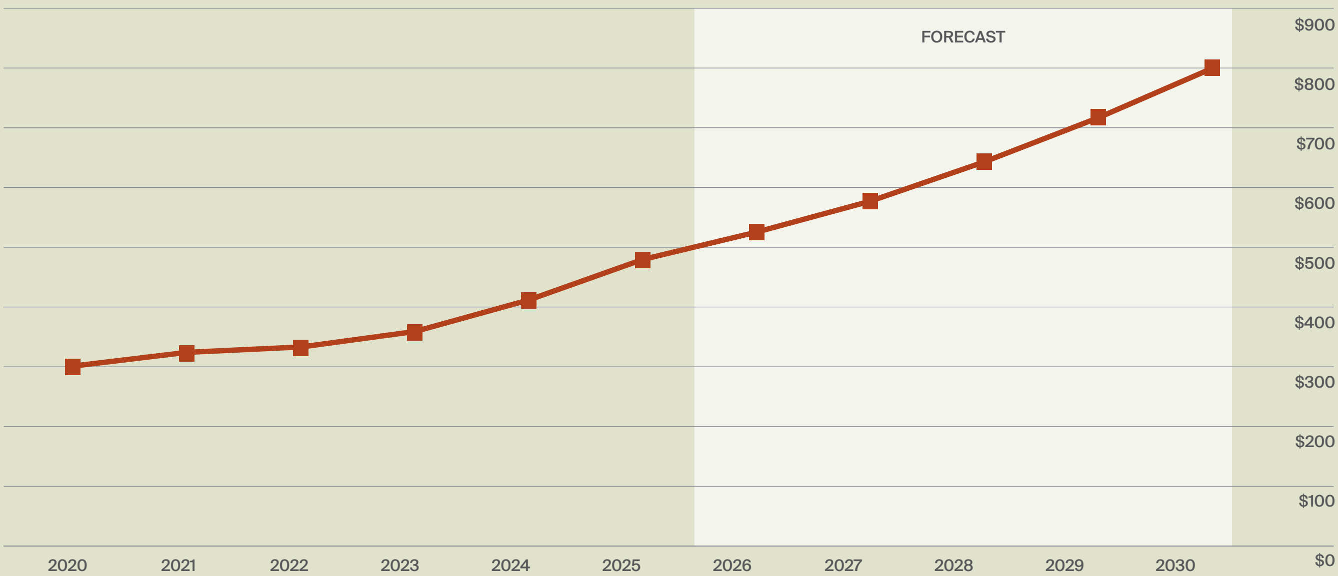

Source: Knight Frank Insight, BloombergNEF

Power, Policy and Perception: What makes a site viable for data centre development?

The global race to build digital infrastructure is accelerating at a pace few anticipated. What was once a relatively niche segment of real estate has become a cornerstone of national economic policy, security and technological competitiveness.

Yet despite, or maybe because of the scale of investment flowing into the sector, a misconception is growing. The assumption that any well-located parcel of land can support a data centre development. Viability, however, is defined by a far more complex equation. Development requires the intersection of power systems, planning frameworks, environmental constraints and cost pressures. Land availability alone is not the principal limiting factor.

POWER IS KING?

Access to power is now the critical determinant of whether a project can proceed.

Across mature markets, grid capacity is tightening under the combined pressure of electrification, de-carbonisation and digital demand. In West London, for example, development in boroughs such as Hillingdon, Ealing and Hounslow is challenged because the network cannot meet the scale of demand. Grid reinforcement is planned, but delivery remains years away.

This is not an isolated case. In Ireland, the relationship between power and planning has effectively been formalised. Following constraints in the Dublin region, new data centre connections are now conditional on-site generation or storage, alongside alignment with new renewable energy supply. In practice, this means that energy strategy must be resolved before detailed design begins.

In the US, same challenge, but at a different scale. Northern Virginia, the world's most established data centre market, is facing transmission constraints and increasing scrutiny over new substations. As a result, data centre developers are shifting attention towards power advantaged states such as Texas and New Mexico, where large scale energy availability is considered easier to secure.

> "Access to electricity has always been important, but it is now the critical determinant of whether a project can proceed."

These examples demonstrate that power is no longer just as part of the process, but is the first hurdle in the viability equation.

Without a credible energisation pathway, supported by realistic timelines, a site cannot support development regardless of its other attributes.

GOVERNMENTS: FRIEND OR FOE?

As power becomes more constrained, national and local governments are playing an increasingly decisive role in where and how data centres are built. This intervention reflects a mix of economic ambition, strategic necessity and environmental responsibility.

In the US, state level incentives are common from, tax abatements, sales tax exemptions and other financial incentives. Each are often tied to capital expenditure thresholds or employment targets, even though data centres are relatively low employment once operational. The broader objective is economic, aimed to stimulate construction, anchor digital ecosystems and attract ancillary industries.

> "As power becomes more constrained, national and local governments are playing an increasingly decisive role in where and how data centres are built."

Europe, currently, presents a more nuanced picture. Governments remain broadly supportive of data centre growth but also demonstrate stricter environmental standards and constrained energy systems. In markets such as London, Frankfurt and Amsterdam, development challenged by securing power and meeting regulatory requirements. This has meant more coordinated planning, including the identification of development zones and accelerated investment in energy infrastructure. In the UK, a 'Power First' policy is being reflected. Data centres are designated as Critical National Infrastructure, with reforms considered to accelerate viable projects through grid connection queues while reducing speculative applications.

In Asia Pacific, governments are also taking a more selective approach. Singapore's decision to pause development before adding capacity through a competitive allocation process reflects an aim to balance growth with sustainability and resource constraints. Meanwhile, markets such as India, South Korea and Malaysia are using regulatory reform and targeted incentives to expand national digital infrastructure and support economic strategies.

> "Government support is forthcoming but dependent on alignment with energy policy, environmental objectives and national priorities."

Across all regions, one theme is clear. Government support is forthcoming but dependent on alignment with energy policy, environmental objectives and national priorities.

ENVIRONMENT VS GROWTH

Planning authorities increasingly assess data centres as infrastructure systems with measurable impacts.

In many jurisdictions, projects must demonstrate access to low carbon electricity, often with requirements

Data centre build cost range
$m per MW

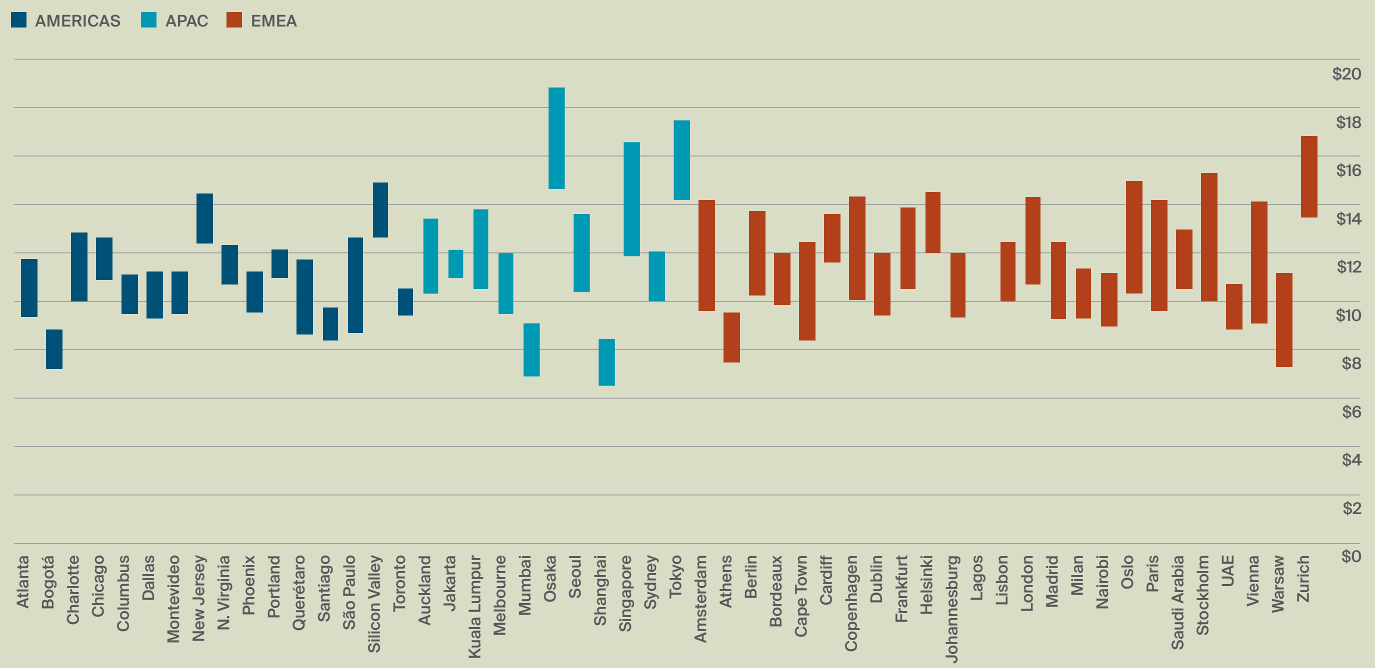

Source: Knight Frank Insight

"Water use is also coming under greater scrutiny, particularly as data centre density increases with the rise of AI workloads."

for new renewable generation, with a preference for 'behind the meter' microgrid solutions rather than total reliance on existing supply. This reflects both regulatory expectations and corporate sustainability commitments.

Water use is also coming under greater scrutiny, particularly as data centre density increases with the rise of AI workloads.

In water stressed regions, this has become a critical issue. Projections in cities such as Sydney suggest that without changes in technology or sourcing, data centres could account for a significant share of potable water demand within the next decade. This has led to closer examination of water strategies, including recycled water use and alternative cooling approaches.

At the same time, some global regions are demonstrating how environmental challenges can be turned into opportunities. In the Nordics, data centres are being integrated into district heating networks, exporting waste heat to local communities. This approach improves overall efficiency while strengthening the social licence to operate.

WHAT ABOUT CONNECTIVITY?

While power has overtaken connectivity as the primary constraint, network infrastructure remains a critical factor in determining both site viability and long-term competitiveness.

Data centres derive value not only from the ability to process data, but from the focal position within global and regional networks. Connectivity underpins latency, resilience and access to users, cloud platforms and enterprise ecosystems. For certain workloads, particularly financial trading, content delivery and interconnection heavy services, proximity to network hubs is critical.

Established locations illustrate the strength of connectivity driven markets. London Docklands, anchored by the London Internet Exchange (LINX), has evolved into one of the most densely interconnected ecosystems in the world. Its proximity to a high concentration of carriers creates a level of network gravity that is difficult to replicate.

Frankfurt offers a comparable model in continental Europe. The presence of DE CIX has helped establish a highly interconnected environment that continues to attract both hyperscale and colocation operators.

In the US, Ashburn in Northern Virginia demonstrates how connectivity and scale reinforce one another. The region benefits from a dense concentration of fibre routes and interconnection facilities.

"In the Nordics, data centres are being integrated into district heating networks, exporting waste heat to local communities."

However, balance is shifting. Power constraints are beginning to challenge further expansion, highlighting that connectivity alone is no longer sufficient to sustain growth.

In Asia Pacific, Singapore remains a leading connectivity hub due to its strategic position on major cable routes linking Asia, Europe and the United States. Emerging hubs such as Johor in Malaysia are additionally benefiting from proximity to Singapore, offering access to similar network routes while providing greater availability of land and power. This reflects a broader trend where secondary markets develop alongside established hubs, capturing spillover demand.

IS THE COST OF BUILDING BECOMING PROHIBITIVE?

Construction costs have increased globally, with benchmarks indicating rises of around 5 to 6% year on year in 2025. For AI ready, liquid cooled facilities, an additional premium of 7 to 10% is becoming typical. In mature markets, this can push total build costs beyond $10million per MW.

Supply chain constraints are a major contributor. Critical components such as high voltage transformers and switchgear are subject to lead times of 100 to 120 weeks or more. Procurement strategy has therefore become central to delivery.

Operational costs are also evolving. Electricity prices remain volatile and structurally higher than pre pandemic levels. In the UK, rising network charges are expected to further increase delivered energy costs from 2026 onwards.

There are, however, countervailing trends. Renewable power purchase agreement prices have softened in parts of Europe, particularly for solar, while battery storage costs have declined significantly. Co located storage is increasingly being used to manage peak demand and improve cost certainty.

As a result, developers are adopting more integrated energy strategies, combining renewable procurement with storage and grid optimisation. While this adds complexity, it can strengthen long term project viability.

> “Network infrastructure remains a critical factor in determining both site viability and long-term competitiveness.”

WHY DO VIABLE SITES STILL FAIL?

Even where sites appear to meet the core criteria of land availability, connectivity and initial power access, projects can and frequently do fail to progress. The reasons are rarely singular. Instead, they tend to reflect a combination of technical, regulatory and social risks that emerge over time.

Grid access and queue uncertainty remains the most immediate and

Data centre regional build cost forecasts

$m per MW

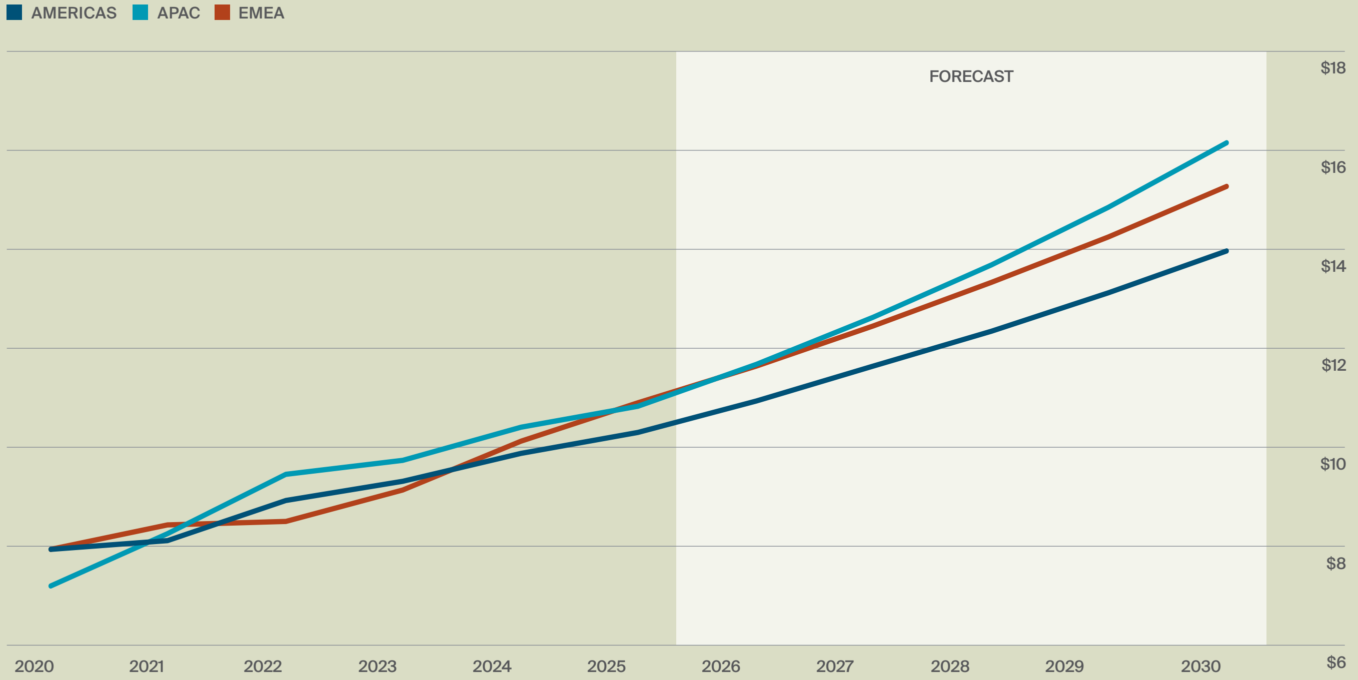

Source: Knight Frank Insight

> "Construction costs have increased globally, with benchmarks indicating rises of around 5 to 6% year on year in 2025."

often decisive issue. In many mature markets, connection queues are congested, with projects competing for limited capacity across multiyear delivery horizons.

Closely linked to this is infrastructure dependency. Data centres rely on substations, transmission upgrades and wider network reinforcement, much of which sits outside the developer's control. Delays in these enabling works can stall projects regardless of progress on site.

Political and regulatory risk adds complexity. Projects often span multiple layers of governance, and approvals at one level do not guarantee delivery. Policy shifts, legal challenges or public scrutiny can alter outcomes even late in the process.

Community is increasingly a factor. Concerns around energy use, water consumption and land allocation are influencing decision making, particularly in constrained regions where trade-offs between competing uses are more visible.

This extends into environmental compliance with the latest challenge to the Woodlands park development in the UK a good example.

ARE WE AT A TIPPING POINT FOR DATA CENTRE DEVELOPMENT?

Data centre development has moved into a strategic phase. The rapid growth of digital infrastructure is now meeting the practical limits of energy systems, environmental capacity and regulation.

Securing a plot is only the beginning. The ability to source reliable power, understand policy constraints, mitigate environmental impacts and manage rising construction and energy costs has become central to project viability. The commercial opportunity remains strong, but the route to delivery is increasingly complex.

For governments, the challenge is equally significant. Data centres support economic growth, national security and technological competitiveness, yet also place pressure on power networks, land supply and local communities. Policymakers must balance these competing priorities across energy planning, environmental regulation and industrial strategy. Supportive policy can unlock investment, but it must be matched with considerations that address public expectations and long-term sustainability.

Ultimately, successful sites will be those that combine the ambitions across a wide spectrum of differing priorities. As digital capability becomes more critical to modern economies, the question is shifting. It is no longer simply about where data centres can be built, but where they can be developed responsibly, timely and at scale.

> "Data centres rely on a combination of assets and attributes, much of which sits outside the developer's control."

Investment: Assessing value, risk, and returns

The global data centre market has entered its most aggressive investment and build-out phase in its history. Demand is being reshaped by the AI super cycle, traditional hyperscale cloud providers, and escalating power constraints that increasingly dictate market behaviour. Investors are navigating an environment of unprecedented capital influx, rising tenant concentration risk, and a widening divergence between power-rich and power-constrained regions.

ARE HYPERSCALE ASSETS STILL A SAFE BET?

Hyperscale data centres have traditionally been considered one of the most resilient real estate investment classes, characterised by long leases, often extending to 15+ years, and strong tenant covenant, with demand growth tied to structural changes from Internet of Things (IoT) and global cloud adoption. Hyperscalers expanded globally, while enterprises outsourced infrastructure needs to colocation providers. Emerging technologies, 5G adoption, and increasing data consumption volumes fuelled growth, particularly in high-density urban markets.

Data centres require high upfront CapEx for land acquisition, facility construction, and power & cooling infrastructure. Modern facilities built today can expect to incur build costs of £15 million per MW, plus an additional 15% for land and power acquisitions. Modern Tier-IV builds incorporate high-redundancy systems, advanced security, and energy-efficient designs, driving costs further. Ongoing OpEx includes energy, cooling, maintenance, and staffing, usually at an annual expectation of 20%, but developers can mitigate these costs through modular expansions and economies of scale.

Despite high initial costs, data centres offer attractive return on Investment (ROI) due to heightened demand, stable cash flows, and significant appreciation potential. Capitalisation rates tend to exceed traditional real estate sectors but have not been immune to global macroeconomic pressures. There has been a period of re-pricing following global interest rate shocks, however prime markets have now stabilised at 4.5%-5.5%. For development projects, with power and planning, discount rates hover between 12-13%, eventually decreasing to 6%-7% once stabilised and fully let, usually over a 4-5 year period. Investors can achieve enhanced returns through lease escalations, expansion projects, and energy cost optimisation. Moreover, the growing demand for sustainable solutions, edge computing, and AI inference facilities provides new avenues for value creation.

> "Despite high initial costs, data centres offer attractive ROI due to heightened demand, stable cash flows, and significant appreciation potential."

Going into 2026, markets are evolving from the initial AI "gold rush" and into

Data centre yield profile

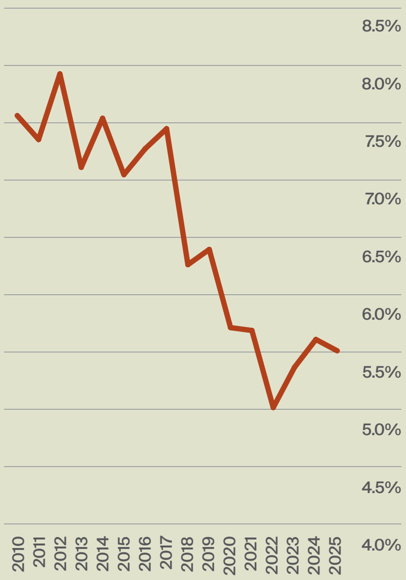

Source: Knight Frank Insight

Global data centre quarterly real estate investment volumes
£ million

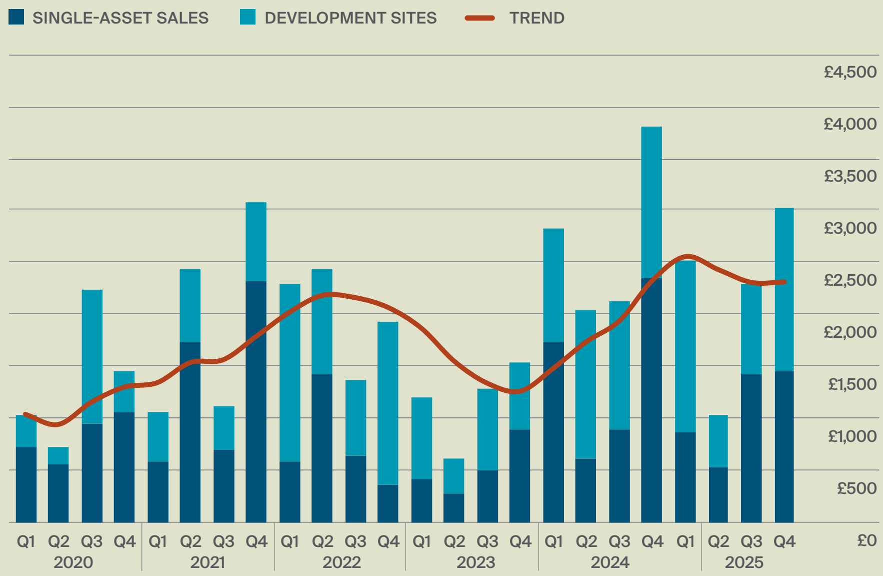

Source: Knight Frank Insight

large-scale industrial development, driven by AI, specialised neocloud providers and high-density campuses in non-traditional locations. Financing models have shifted away from traditional real estate structures towards sophisticated infrastructure vehicles, such as yield companies, joint ventures, continuation funds, and specialised REITs, to address liquidity challenges associated with large asset sizes. Single-asset transactions remain illiquid, so operators favour platform-level value retention. However, an increase in single disposals is expected, enabling faster capital recycling and compressing cap rates later in 2026.

> "Investors can achieve enhanced returns through lease escalations, expansion projects, and energy cost optimisation."

AN AI SUPER-CYCLE?

AI has become the leading driver of investment into the data centre sector, with demand for the sector dominating pipelines and long-term investment focus. During 2020 and 2021, quarterly data centre real estate investment volumes, through single-asset-sales and development site acquisitions, averaged between £1-£1.5 billion, whilst in 2025, this volume has doubled to between £2.5-£3 billion. Investment into stabilised assets has remained relatively static, with volumes averaging £1 billion per quarter, whilst the growth has been driven by an increasing focus on development sites. Development site quarterly investment volumes have grown from £500 million in 2020-21, up to £1.3 billion across 2024-25.

Data centre investment markets had to navigate some challenges during 2025, particularly in relation to the launch of DeepSeek in Q1 and Microsoft lease cancellations, resulting in a "lost" second quarter in Q2. The launch of DeepSeek created concerns over current US investment models into AI, whilst

> "During 2020 and 2021, quarterly data centre real estate investment volumes, through single-asset-sales and development site acquisitions, averaged between £1-£1.5 billion, whilst in 2025, this volume has doubled to between £2.5-£3 billion."

Microsoft's cancellation of 2GW of data centre leases creating suggestions of market over-supply. This saw investment volumes in Q2 drop 50%, based on Knight Frank's rolling four-quarter average, to £1 billion.

Markets have since recovered from these shocks, owing to questions over the credibility of true DeepSeek costs, and Microsoft lease obligations surging in the quarters since from $92.7 billion in Q2 up to $155.1 billion in Q4. Real estate transactions recorded in 2025,

Future data centre lease obligations
$ million

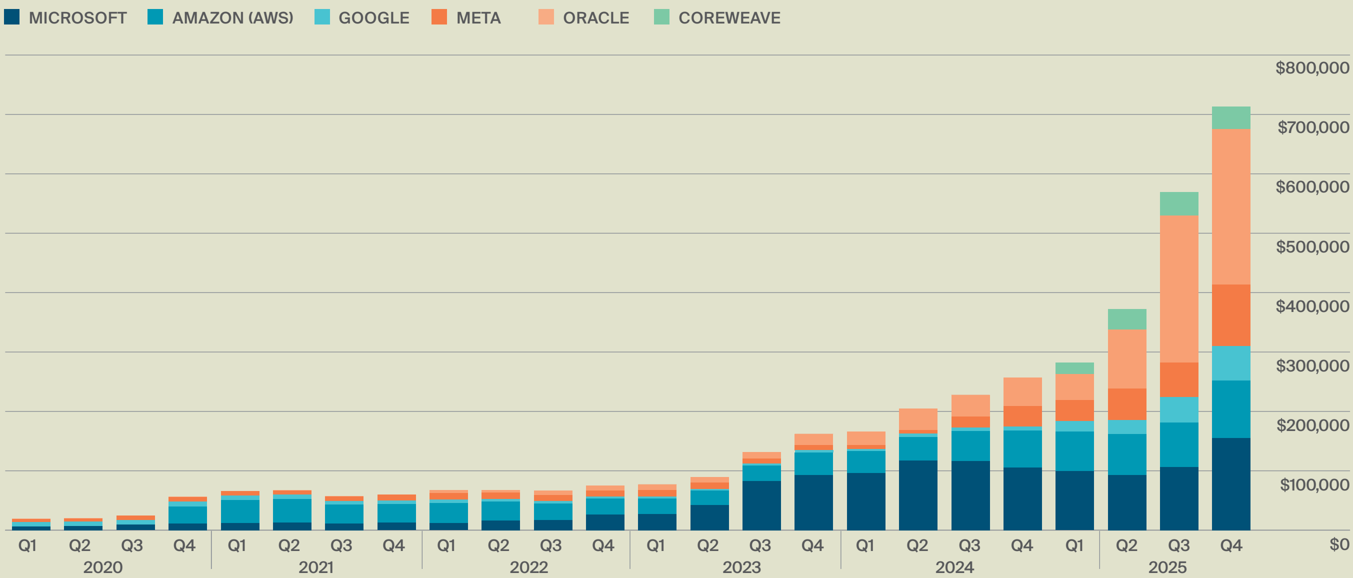

Source: Knight Frank Insight

> "2025 saw a record-breaking deal: the $40 billion acquisition of Aligned Data Centers by a consortium comprising AIP, GIP, and MGX."

not including portfolio acquisitions, saw a 17.5% drop compared with 2024, although this is the result of the temporary Q2 slowdown.

M&A and JV investment volumes were also not immune to the Q2 slowdown, but annual volumes increased 40% on 2024 volumes, with $82 billion recorded, and entered 2026 with $137 billion worth of agreed but yet-to-complete deals. 2025 saw a record-breaking deal: the $40 billion acquisition of Aligned Data Centers by a consortium comprising AIP, GIP, and MGX. The transaction was agreed in Q4 2025 and is scheduled to close in H1 2026.

TENANT CONCENTRATION RISK

One key concern from investors, and one cause of the Q2 slowdown, is the identification of growing tenant concentration and whether this is becoming too high in certain hyperscale portfolios. Global hyperscalers (such as Microsoft, AWS, Google, and more) are the largest consumers of data centre capacity across the globe. Globally, 11GW of new capacity was delivered in 2025, 66% of which was absorbed by these hyperscalers, with the largest consumers being AWS and Microsoft, absorbing 2.3GW and 1.7GW, respectively, 40% of this new capacity.

Hyperscaler dominance of colocation leasing events and new self-build capacity has seen market representation for these participants rise from 35% in 2017 to 69% in 2025. This has been the result of an aggressive build-out and leasing strategy by these operators, having been responsible for building and leasing 79% of new global data centre capacity since 2017. Resultingly, global

Global data centre M&A investment
$ million

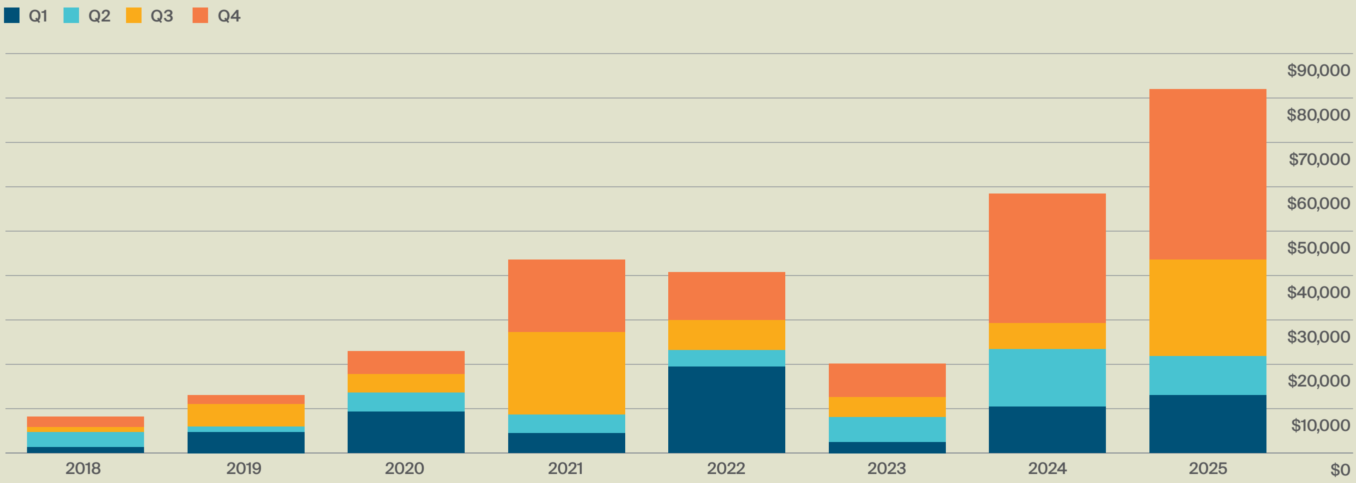

Source: Knight Frank Insight

Global cloud service provider revenue
$ million

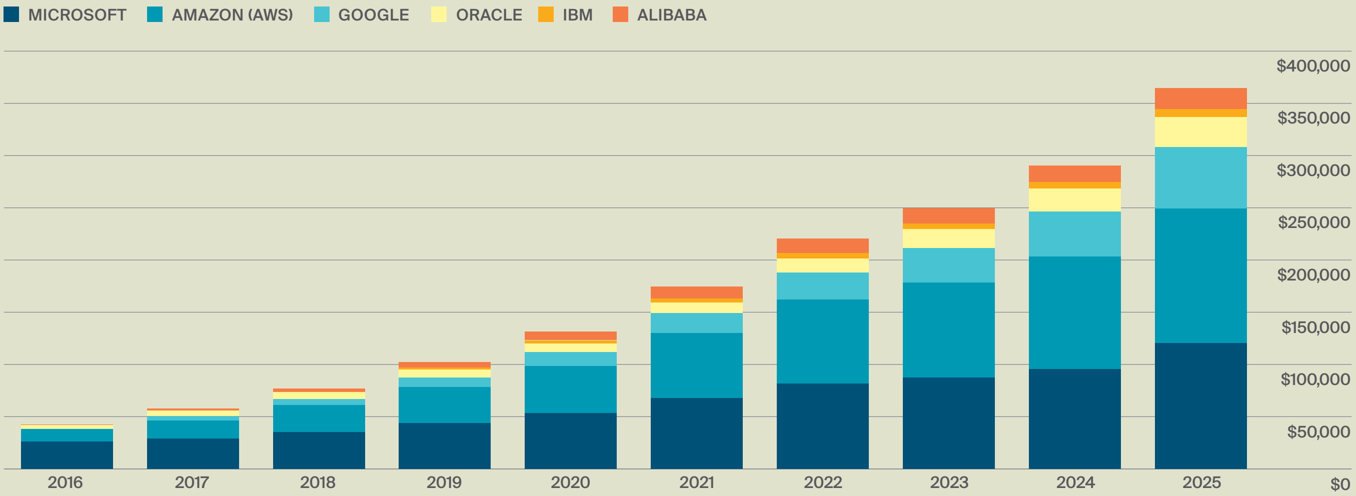

Source: Knight Frank Insight

hyperscalers have amassed a combined portfolio of 40.2GW, via leased colocation and self-build capacities.

The consequence for the industry has been a rise in tenant concentration risk, with a growing dependence on the financial conditions and strategic decisions of a limited number of major market participants. In response, developers are increasingly being advised to mitigate exposure by securing leases from multiple tenants and delivering capacity in phases, instead of relying on a single hyperscaler to pre lease an entire campus.

THE POWER PREMIUM

As global demand for data centres has accelerated over the past decade, coupled with mounting pressure on grid networks around key metropolitan areas, the value of land with access to near-term power has risen sharply. In core data centre markets, sites capable of delivering 50MW+ grid connections within a three-year time frame now command significant premiums. Knight Frank analysis indicates this premium has pushed power land in tier-I markets to as much as three-to-four times traditional industrial land values.

> "In core data centre markets, sites capable of delivering 50MW+ grid connections within a three-year time frame now command significant premiums."

For active real estate investors without substantial financing capacity, this trend underscores a strategic opportunity in land banking within high-demand regions. By acquiring land early and focussing on securing grid connections, investors can subsequently dispose of these power-ready sites to data centre developers and realise the associated premiums.

HYPERSCALER / REIT PERFORMANCE

Public cloud providers are responsible for the majority of global data centre leasing activity, being the largest occupiers of both self-build and colocation-based data centre capacity. Cloud occupiers were responsible for 49% of leasing transactions in 2025 and have been responsible for 69% of aggregate leasing activity to date. The growth and stability of cloud-based revenue is integral to the health of the data centre industry and is an important metric for evaluating growth prospects for the market. Combined cloud revenues for the four major US public cloud service providers – Microsoft, Amazon, Google, Oracle

> "Combined cloud revenues for the four major US public cloud service providers – Microsoft, Amazon Web Services, Google, Oracle – grew by 25.5% in 2025 ."

– grew by 25.5% in 2025 and have expanded at a compound-annual-growth-rate (CAGR) of 23% over the last five years.

Since 2017, the S&P 500, S&P 500 IT, and S&P 500 Real Estate indexes have grown by 206, 604, and 36 bps, respectively. At the same time, US global cloud service providers Microsoft, Amazon, Google, and Oracle, have expanded by 678, 516, 690, and 407 bps, respectively, each of which exceeding index growth volumes for the S&P 500 index, with both Microsoft and Google exceeding the S&P 500 IT index. Specialised real estate investment trusts (REITs) Equinix, Digital Realty, Iron Mountain, and GDS Holdings have each grown by 159, 66, 212, and 203 bps, respectively, all of which exceeding indexed growth rates of the S&P 500 Real Estate index, alongside three of which expanding faster than the S&P 500 index as well.

Indexed price returns

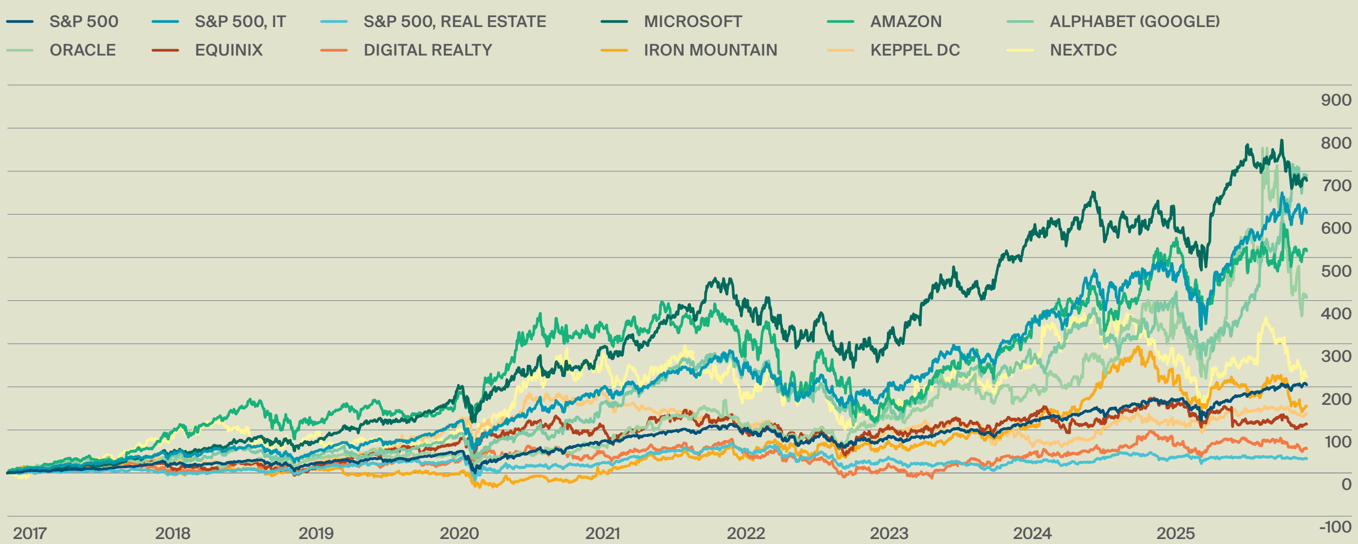

Source: Knight Frank Insight

We like questions, if you've got one about our research, or would like some property advice, we would love to hear from you.

Global

Stephen Beard
Global Head of Data Centres
stephen.beard@me.knightfrank.com

UK and Europe

Oscar Matthews
Partner
oscar.matthews@knightfrank.com

Celeste McGinley
Senior Surveyor
celeste.mcginley@knightfrank.com

Olamide Emiloju
Surveyor
olamide.emiloju@knightfrank.com

Archie MacColl
Surveyor
archie.maccoll@knightfrank.com

Valuation

Alex Burgoyne
Global Head of Data Centres Valuations
alex.burgoyne@knightfrank.com

Agne Berzinskaite
Associate
agne.berzinskaite@knightfrank.com

Jonathan Oliver
Associate
jonathan.oliver@knightfrank.com

Isobel Green
Senior Surveyor
isobel.green@knightfrank.com

Jamie Hall
Surveyor
jamie.hall@knightfrank.com

Americas

Michael Morris
President, Data Center Capital Markets & Advisory
mmorris@cresa.com

MENA

Si Lau
Surveyor
si.lau@me.knightfrank.com

Reed Santos
Surveyor
reed.santos@knightfrank.com

APAC

Jiya Agrawal
Analyst, APAC
jiya.agrawal@asia.knightfrank.com

Power & MEP Consultancy

Christopher Jones
Head of Power Procurement & MEP Consultancy
chris.jones@knightfrank.com

Insight

Darren Mansfield
Head of Data Centre Insight
darren.mansfield@knightfrank.com

Harry Hannam
Associate
harry.hannam@knightfrank.com

Nicola Ryan
Senior Analyst
nicola.ryan@knightfrank.com

Knight Frank