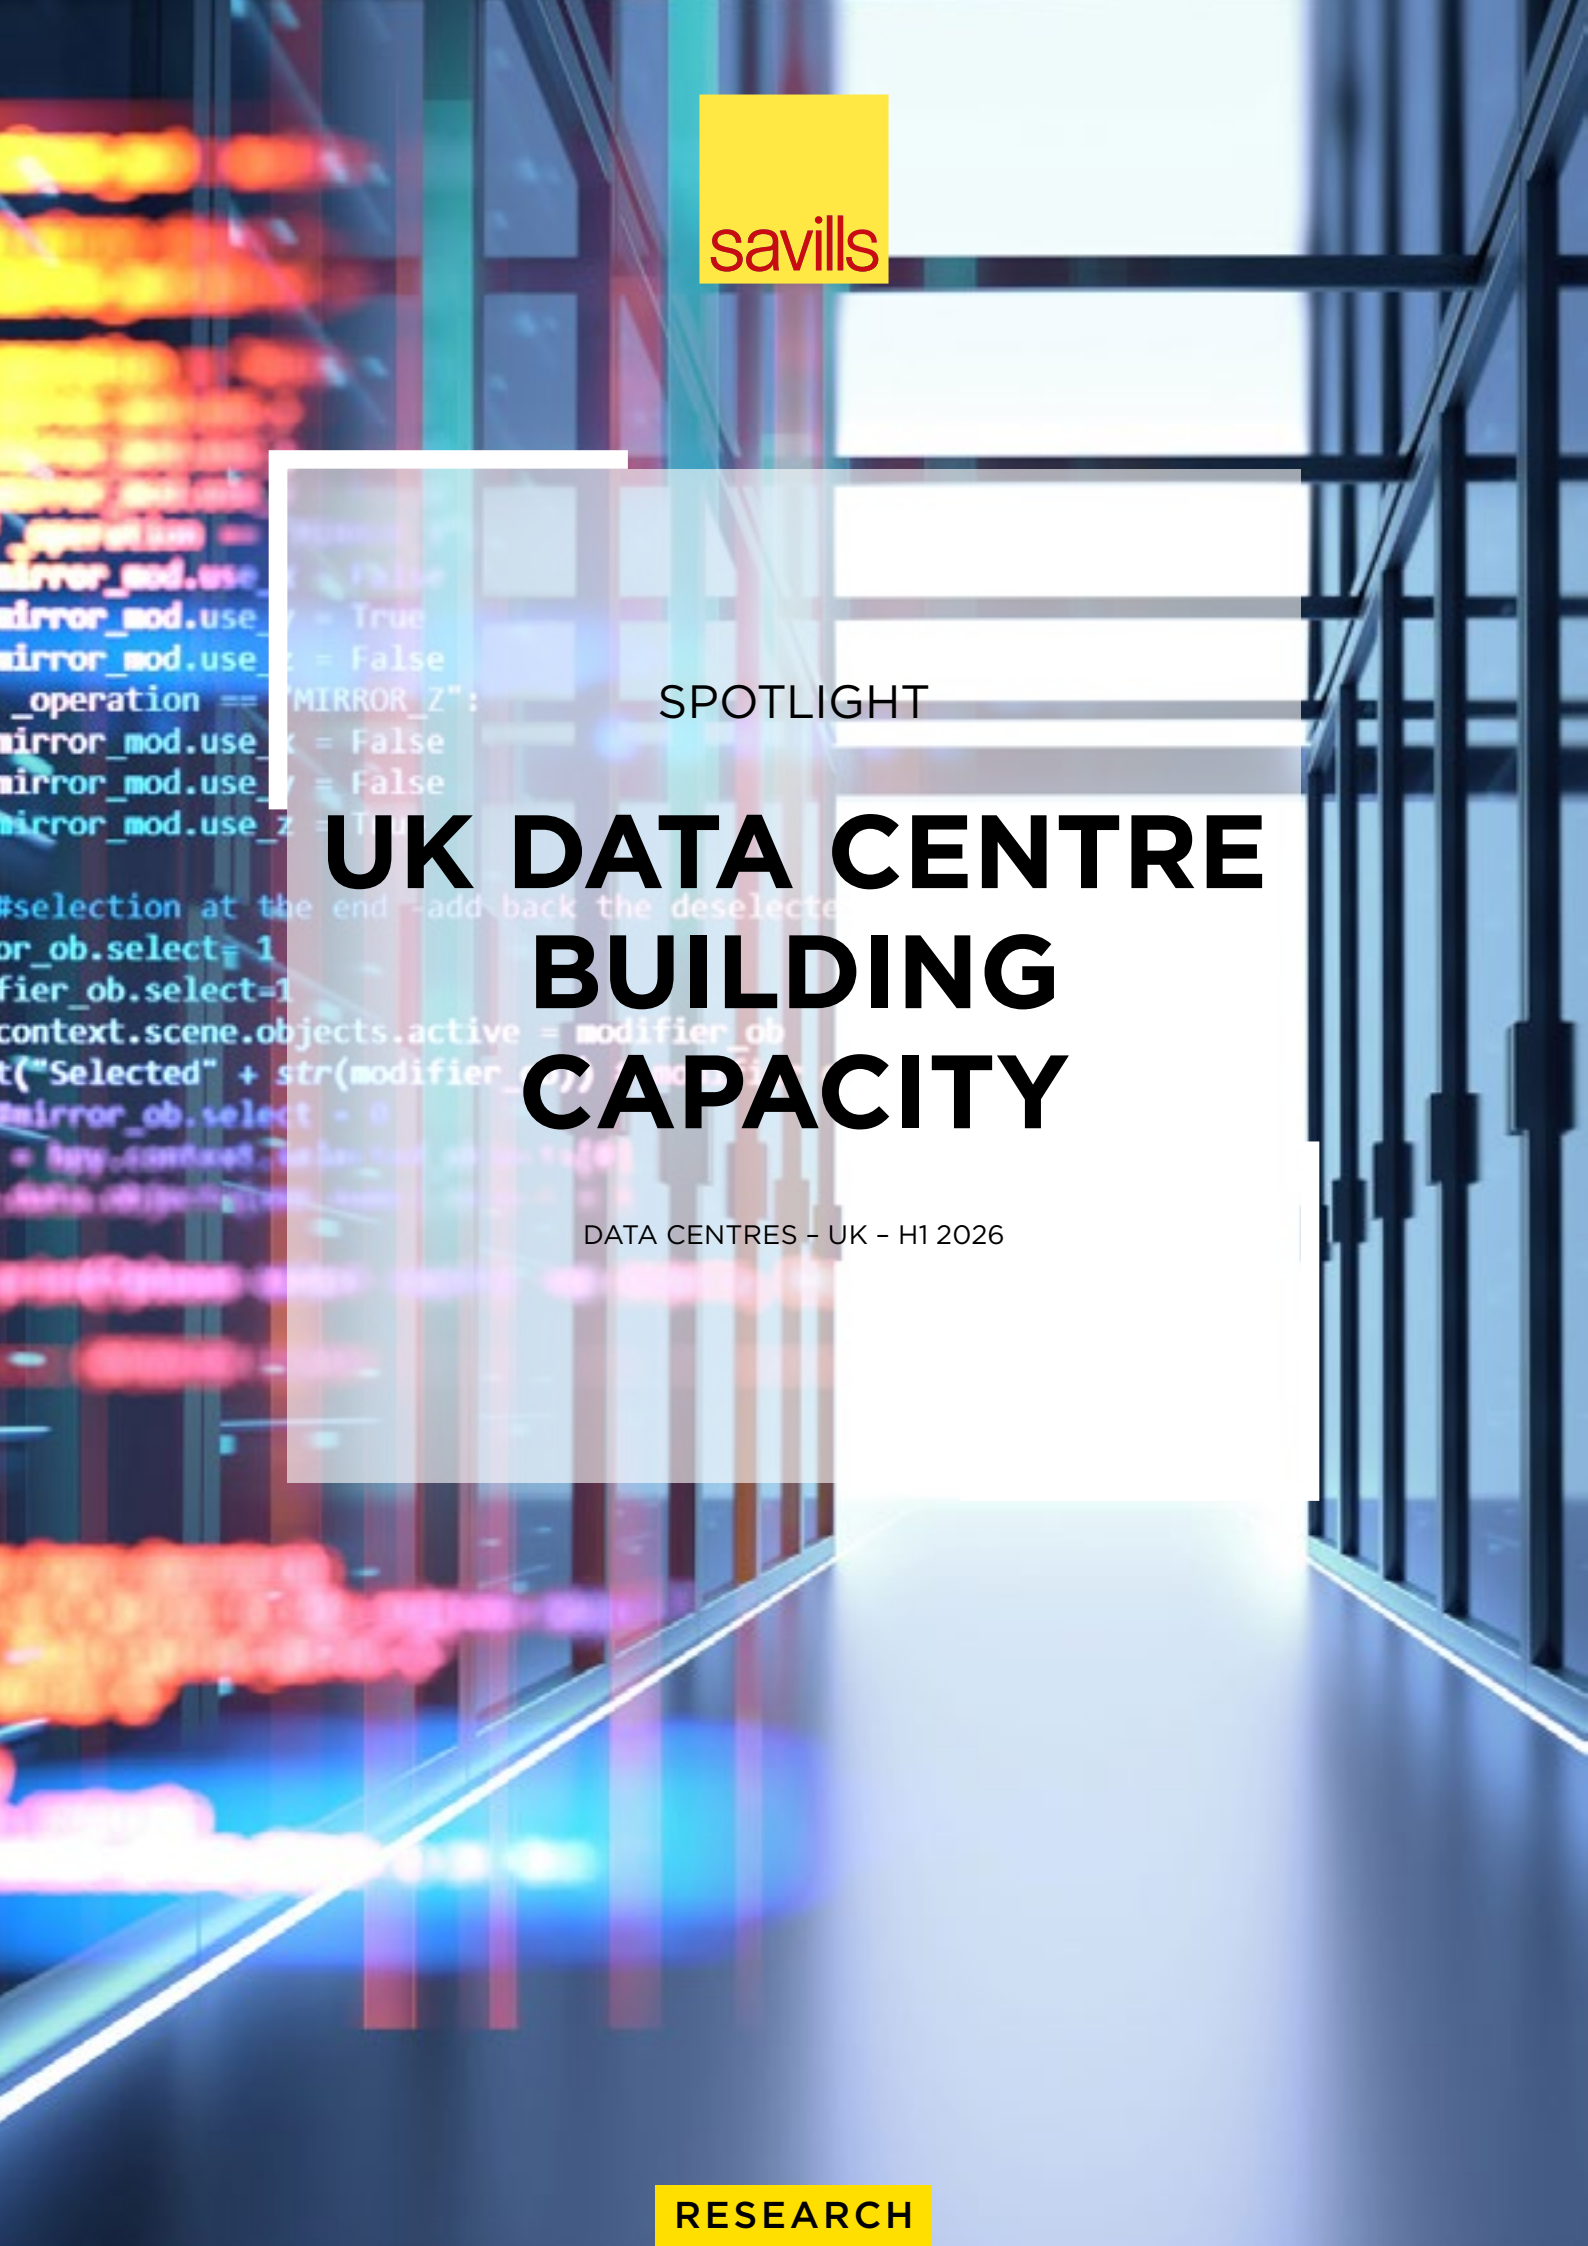

savills

SPOTLIGHT

UK DATA CENTRE BUILDING CAPACITY

DATA CENTRES – UK – H1 2026

RESEARCH

KEY TAKEAWAYS

01

The UK data centre market remains structurally attractive, supported by hyperscale cloud, AI and resilient enterprise demand. The constraint is not appetite, but the availability of genuinely deliverable capacity.

02

London continues to dominate because its fibre density, cloud ecosystems and availability zones are difficult to replicate. However, that same concentration intensifies pressure on power, land and delivery timelines.

03

Supply and demand remain exceptionally tight. New deliveries are being rapidly absorbed, leaving the market increasingly reliant on pre-commitment, while the vacancy rate has fallen from 27% in 2016 to 8% in Q1 2026.

04

Supply is expanding on paper, but much of the pipeline remains early-stage and uncertain. Pre-letting has become a core feature of the market, with future capacity increasingly committed well before completion.

05

Power is now the primary gatekeeper of growth. Grid reform should improve queue discipline, but it will not eliminate the underlying capacity scarcity or the long lead times required for reinforcement.

06

Energy costs also matter more than before. In a European context, the UK remains a relatively expensive power market, which weakens its competitiveness for the most energy-intensive deployments.

07

Land strategy is no longer just about the site itself, but about securing the wider corridor of rights needed to make development possible. Third-party land consents for cable routes, pipework and access are increasingly part of the critical path.

08

Planning conditions have become more supportive, helped by stronger policy recognition of data centres and broader consenting routes. Even so, consent is increasingly conditional on schemes demonstrating credible mitigation and locally visible benefits.

09

Construction has become a major execution filter rather than a straightforward delivery step. Rising costs, specialist labour shortages and long lead times for key equipment are widening the gap between consented projects and operational stock.

10

Social impact is now a material part of viability. Developers that can translate national strategic value into local benefits, such as skills, partnerships or heat reuse, are likely to be better placed to secure support and reduce planning risk.

11

Investment appetite remains strong, but capital is more selective and infrastructure-led. Value is concentrating in power-secured land, operational platforms, and projects with a clear route through planning, energisation and delivery.

12

Looking ahead, the UK is unlikely to lack demand, capital or strategic relevance. The decisive question is which projects can convert those advantages into deliverable capacity in a more constrained, regulated and execution-sensitive environment.

UK DATA CENTRE - BUILDING CAPACITY

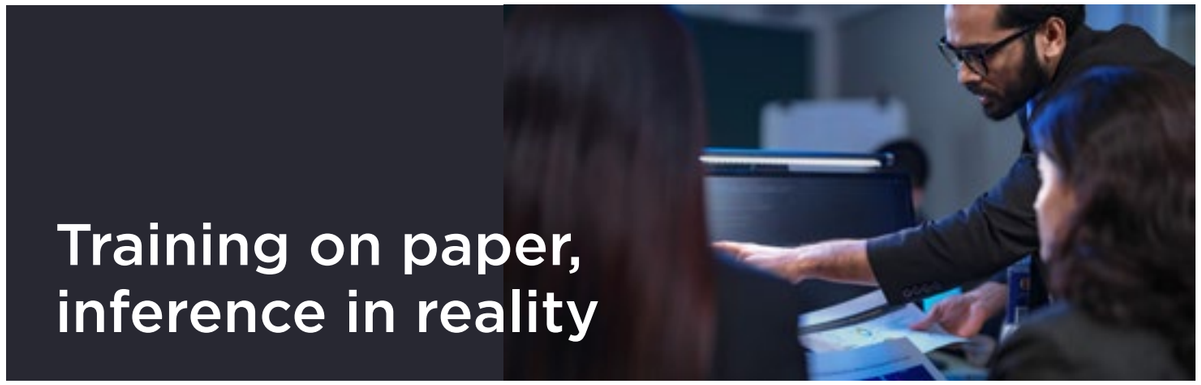

Training on paper, inference in reality

London leads the market, yet grid constraints will shape what comes next.

Over the past two years, the UK data centre market has entered a new phase, driven less by traditional enterprise cloud adoption and increasingly by artificial intelligence workloads. Hyperscale cloud, AI native “neocloud” providers, and sovereign compute requirements are converging, creating unprecedented demand for high-density, power-intensive infrastructure. Nowhere is this more visible than in and around London, which remains Europe’s dominant data centre hub, accounting for more than 80% of UK supply and acting as the primary landing zone for cloud availability zones, AI platforms and interconnection ecosystems.

AI is fundamentally reshaping demand. Large-scale training workloads require concentrated campuses with thousands of GPUs, extreme rack densities and access to tens or even hundreds of megawatts of power. This has given rise to a new category of infrastructure providers, often referred to as neoclouds, that offer GPU-as-a-service platforms optimised for AI training and inference. These players are now competing directly with hyperscalers for capacity, accelerating take-up and pushing facilities to technical limits that were unthinkable only a few years ago.

However, while announced investment and planning pipelines suggest explosive growth, structural constraints are becoming decisive. Power availability, grid connection delays, planning friction and sustainability pressures are no longer secondary risks; they are now the primary filters determining which projects will actually be delivered. Ofgem data shows that proposed UK data centre projects are collectively requesting around 50GW of grid capacity, more than the country’s current peak electricity demand, making it mathematically impossible for all schemes to proceed as planned.

This tension raises a critical question for the UK market: are we overbuilding AI training capacity on paper while underestimating the longer-term, more distributed growth of inference-driven data centres?

Evidence increasingly suggests that while a limited number of very large AI training campuses will be built, large training clusters are also more likely to reach saturation once a handful of frontier-model builders and hyperscalers have secured the scarce combination of land, cooling and 50–200MW+ grid connections they require. Notably, many of the largest training campuses are being developed in the United States, where the frontier AI ecosystem is most concentrated and where hyperscalers and specialist operators are deploying multi-hundred-megawatt sites at unprecedented scale. This concentration can further narrow the addressable market for additional UK-based mega-campuses beyond a few flagship builds.

Training demand is inherently “lumpy”, and cyclical capacity is added in step-changes around new model generations. Then utilisation can soften as workloads shift to optimisation, fine-tuning, and the reuse of trained models, meaning additional campuses quickly compete for the same relatively narrow customer base.

> “AI training drives headlines; inference drives the long-term demand reality. Over time, delivery will be shaped by power availability, grid access and planning constraints.
>
> **CAMERON BELL**

By contrast, inference capacity, closer to users, businesses, and population centres, is likely to grow steadily for much longer. It scales with the breadth of AI adoption across sectors, needs lower latency and regional resilience, and can be deployed in smaller increments across existing colocation and edge footprints, tracking economic activity rather than model-training cycles.

In the UK, this points to sustained demand not only within the London metro, where financial services, media, enterprise SaaS and dense interconnection ecosystems create persistent low latency requirements, but also across emerging regional clusters such as the Midlands, the North West and Scotland, as organisations embed AI into day-to-day workflows. Crucially, inference build-out can often proceed in modular tranches, such as incremental 5–20MW deployments, and through retrofits or expansions of existing campuses, making it better aligned with the UK’s constrained grid-connection environment and planning realities than the next wave of ultra-large training sites.

London's gravity, hyperscale volatility & the next constraint

The UK market has entered a structurally tighter phase: demand has broadened and deepened, but deliverable supply remains constrained by power, land and time.

Occupier market: three engines, one choke-point

The UK data centre market is no longer best explained by a single technology cycle. Demand is being propelled by three overlapping engines: cloud-led scale, AI-driven power intensity, and a persistent enterprise and colocation undercurrent. All are converging on a limited pool of capacity that can be delivered in the right places and within occupiers' timeframes. That tension is most acute in London, where cloud ecosystems, network density and established availability zones continue to anchor site selection and, in doing so, make infrastructure constraints the practical limit on choice.

Recent metrics illustrate the consequences. Annual take-up has moved onto a higher runway since 2020, reaching 217MW in 2025, a record high. The subdued 19MW recorded in Q1 2026 is better read as a symptom of limited near-term options than a reversal in requirement depth. Vacancy has compressed sharply, from 29% in 2006 to 8% in Q1 2026, with London tighter at 7%. Commitment timing has therefore shifted forward: in 2024, pre-sold capacity from future projects totalled 257MW, alongside 212MW of take-up and against deliveries of 190MW, reinforcing that a significant share of demand is being secured well before assets become operational.

Cloud-led demand remains the principal volume driver, but procurement has become more fluid. Hyperscalers still set the marginal price for well-located, power-secure campuses, yet capacity is increasingly accessed through wholesale platforms and self-build. Wholesale colocation accounts for 1,163MW of live UK capacity, and rises to 2,729MW when live, under-construction and committed schemes are combined.

> The UK market is delivering record capacity, but demand is absorbing it even faster. Hence, vacancy has dropped over the past decade, from 27% in 2016 to just 8% in Q1 2026.
>
> **LYDIA BRISSY**

AI-related demand is changing the character of requirements as much as their scale. Compared with traditional cloud deployments, AI places disproportionate weight on power density, cooling specification and energy economics, and it tends to arrive in larger, more time-sensitive blocks. Market feedback suggests early European momentum was interrupted while design standards and investment priorities were recalibrated, but confidence returned later in 2025 as major occupiers re-entered the market. Reflecting that shift, AI now appears as a discrete category in the UK development landscape, with 50MW recorded within self-build AI capacity in the combined live-plus-pipeline view.

Enterprise and broader colocation demand provide the market's constant, and it is increasingly expressed through hybrid strategies that blend public cloud with third-party facilities for control, resilience and compliance. Retail colocation remains material, with 347MW of live capacity, but tighter supply conditions are encouraging larger suites and longer commitments that are beginning to resemble wholesale procurement. For investors, this convergence blurs traditional segmentation and increases the premium placed on platforms that can flex across tenant types without relying on a single anchor to underwrite the business case.

FIG 1: UK SUPPLY & DEMAND DYNAMICS

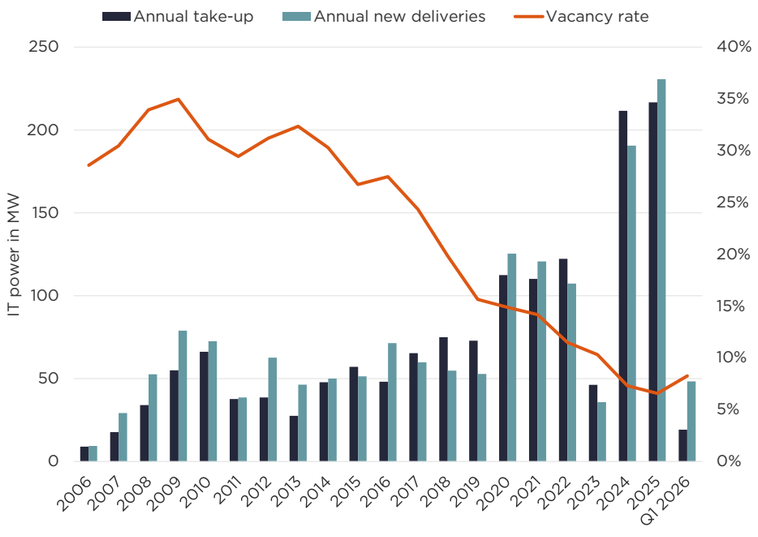

Source: Savills Research based on DC Byte

UK DATA CENTRE - BUILDING CAPACITY

London's gravity, supply reality

London's dominance is not a legacy artefact; it has strengthened as the market has scaled. UK live capacity has reached 1,803MW in 2026, and London accounts for 1,637MW, around 91% of the national stock, up from 79% in 2006. This concentration reflects the continued pull of the capital's connectivity and cloud ecosystems, but it also means that UK-wide capacity expansion does not automatically translate into choice for occupiers. Outside the London orbit, the cloud market remains comparatively shallow and less able to absorb requirements at pace.

As a result, London operates as both the market's demand engine and its constraint amplifier. Where power and land are scarce, timelines elongate, and occupiers face a tighter trade-off between location fidelity and delivery certainty. It is this dynamic, rather than any sudden weakening of London's role, that is prompting a broader search for alternative geographies.

Market discussions reinforce that decentralisation is, for now, incremental. Activity remains anchored to tier-one locations around London, with only selective regional depth emerging. Manchester is the most cited secondary market, but follow-through beyond London and Manchester has been limited, reflecting how difficult it remains to replicate the capital's combination of network density, occupier ecosystems and timely access to power.

> Regional interest is growing due to market constraints; however, the majority of requirements still gravitate to London's availability zones, as the ecosystem is hard to replicate.
>
> **NIAMH DOHERTY**

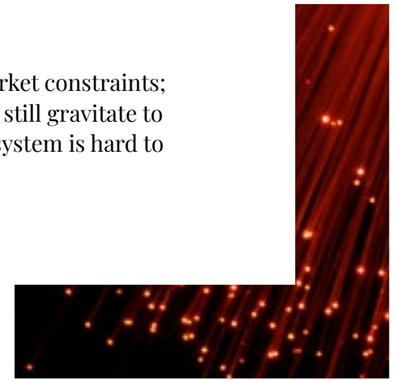

Where regional opportunities do look credible, they are typically those that can offer what London cannot: scale, land and a clearer pathway to power at scale. This is not a redirection of cloud demand, which continues to cluster around established availability zones, but a structural addition driven by AI's fundamentally different site-selection criteria: geographically agnostic, less latency-sensitive, and able to follow power and land rather than population centres. The North East is increasingly positioned in that bracket, supported by renewable infrastructure and grid connectivity, and by a planning context that is more receptive to large-scale investment.

Even so, this is best understood as a medium-term diversification story rather than an immediate release valve for the broader market.

91%

SHARE OF LONDON LIVE IT POWER IN THE UK

FIG 2: THE LONDON DOMINANCE STORY

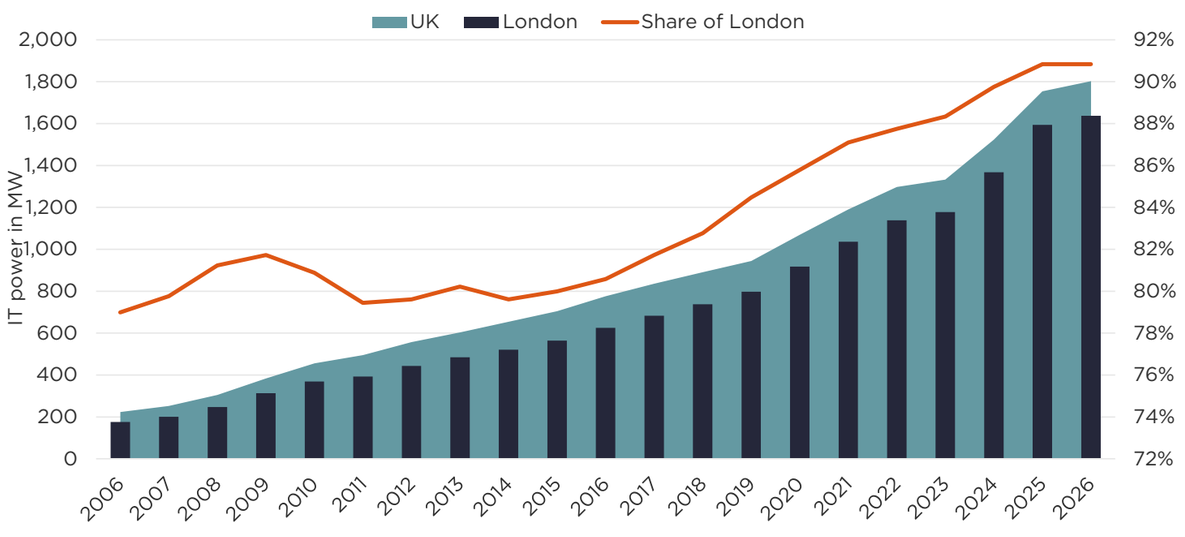

Source: Savills Research based on DC Byte

Bigger facilities, fewer marginal locations
Alongside geographic concentration, the market's development profile is shifting, with growth increasingly delivered through fewer, larger facilities. Average UK data centre size has risen from 2.1MW in 2006 to 6.3MW in 2026, and while sub 5MW sites still dominate (171 of 258 existing facilities), incremental capacity is being delivered through fewer, much larger projects. The direction of travel is unambiguous: there is only one operational facility above 100MW today, Vantage Data Centres' Cardiff (CWL 1) in Newport, offering 130MW of IT power, while 28 projects of that scale sit in the wider pipeline.

This re-scaling concentrates both opportunity and risk. Larger campuses can relieve scarcity through step-changes in capacity, but they also increase exposure to programme and energisation risk, and they deepen reliance on a smaller pool of anchor occupiers. Operator concentration remains high today, with the top ten operators controlling 62% of UK live capacity. These structural dynamics explain why the pipeline appears extensive, and why deliverability, rather than headline megawatts, is the critical lens.

A large pipeline on paper, a narrow window in practice
Last year, 231MW of IT capacity was delivered in the UK, following 191MW in 2024, both record highs. That step-change has not yet repeated in 2026, with 48MW delivered in Q1. Of the 242MW currently under construction across the UK, around 66MW is scheduled to complete within the remainder of 2026 (Q2–Q4), underscoring how limited the immediate delivery window is relative to the headline pipeline.

Headline pipeline numbers can be misleading without a clear view of what is actually deliverable. The UK pipeline totals around 13,500MW, approximately 7.5 times current live stock, yet the executable component is concentrated in a relatively small number of schemes. In addition to the 242MW under construction, a further 2,084MW is categorised as committed, while the balance (9,371MW) remains early stage, where grid access, planning outcomes and financing are still uncertain. London's dominance persists in the most advanced stages, accounting for 240MW of capacity under construction and 1,652MW of committed capacity.

In a FLAPD context, this concentration underpins London's continued leadership over the medium term: by 2030, London's stock is expected to reach around 2,679MW, still well ahead of other FLAPD markets. That implies London grows by roughly 1.1 GW between 2025 and 2030, materially outpacing all other FLAPD cities over the same period.

7.5 X
UK PIPELINE COMPARED TO LIVE STOCK

Even within the under-construction cohort, availability is thin. Of the c.66MW scheduled for completion in 2026, 96% is already pre-sold, leaving less than 3MW available to transact. The forward picture is only marginally looser, across the 37MW currently programmed for 2027 delivery, just over half is already taken up. The market, therefore, does not move from tight to balanced simply because cranes are visible. Instead, pre-letting has become part of the supply mechanism, pulling future capacity forward into today's decision set and reducing the stock of walk-up options for both hyperscale and enterprise occupiers.

FIG 3: UK DATA CENTRES GROWING BIGGER

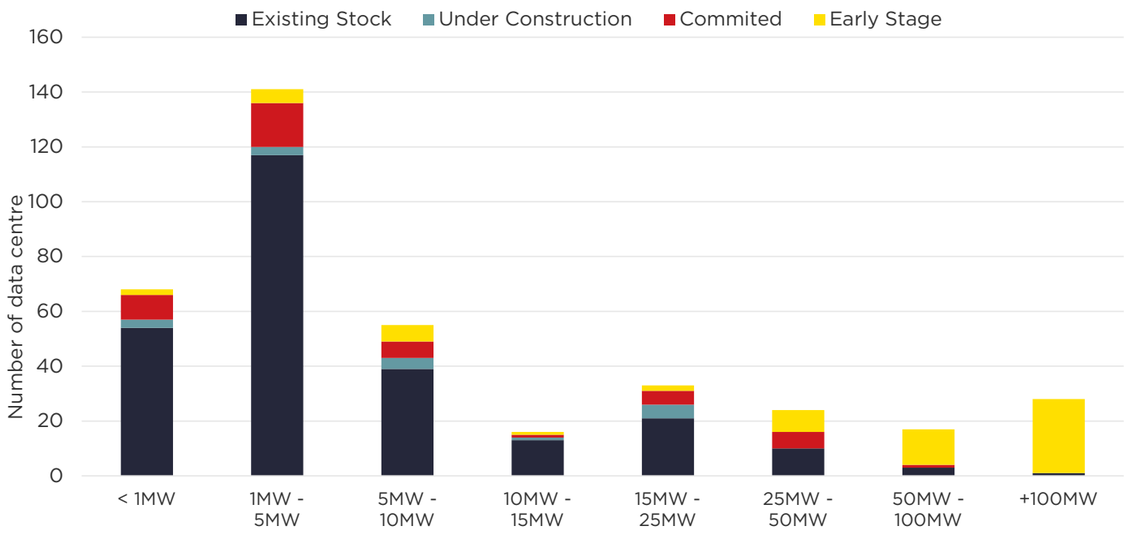

Source: Savills Research based on DC Byte

Constraints and consequences: power, land and the economics of time
The UK's supply–demand imbalance is ultimately a constraint story, and power sits at the apex. London's growing share of national capacity reflects rational demand clustering around established cloud zones, but it also heightens exposure to grid scarcity and long connection lead times. Market conversations suggest that, in some of the most constrained West London locations, developers are now planning for multi-year waits for firm capacity, long enough to undermine traditional development sequencing and force a rethink of how projects are energised. The direction of travel is clear: delivery certainty is increasingly shaped by energy strategy as much as by real estate fundamentals.

Land constraints reinforce the power challenge. Where power and connectivity coincide, land values and competitive tension rise sharply, and the market can generate speculative behaviour, applications submitted in the hope of capturing an outsized uplift if power is secured. However, the volume of proposals should not be mistaken for deliverable supply. Practitioner feedback indicates that a significant proportion of speculative applications are unlikely to progress to development, whether due to grid allocation outcomes, planning friction, or the capital intensity required to meet modern design requirements. For developers, this places a premium on origination capability and stakeholder management; for occupiers, it increases the value of partnering early with platforms that can demonstrate credible pathways through planning, procurement and energisation rather than relying on announcements.

Construction timelines and delivery risk further shape market outcomes. The UK's average facility size has increased, reflecting the shift towards larger, more complex facilities where mechanical and electrical fit-out drives both cost and programme. While retrofitting existing industrial shells can be technically feasible in some cases, the economics are often less compelling than occupiers expect because the critical cost sits in M&E rather than the envelope. For investors, these dynamics sharpen the distinction between 'pipeline narrative' and investable reality: consenting and delivering powered capacity becomes the scarce capability, while underwriting must account for the possibility that timing slippage turns today's contracted demand into tomorrow's renegotiation.

> Alternative power solutions are moving from contingency to strategy, operators are using them to bridge the gap to grid capacity and secure position in the right locations.
>
> **RUPERT DUCKWORTH**

Taken together, these forces sustain a landlord-leaning market even after periods of heavy building. Tight vacancy, high pre-letting and volatile delivery mean that pricing and lease structures are increasingly anchored in scarcity and certainty rather than marginal cost alone. At the same time, constraints are beginning to create a selective rationale for activity beyond the traditional London core, not because the market has become location-agnostic, but because the economics of time and power increasingly reward places where those constraints are less binding.

Queue discipline: Separating signal from noise in UK power access

In the UK, access to power is increasingly the defining constraint for data centre growth, shaping where schemes locate, how quickly they progress, and how investable they are.

Energy as the gatekeeper

Energy is now the primary delivery risk for UK data centre development. While the sector's growth is anchored in cloud adoption and AI-led compute demand, a scheme's viability is increasingly determined by whether power can be secured at the right scale, in the right location, and within an investable timeframe. This has brought grid capacity, connection timescales and regulatory change to the forefront of site selection and planning strategy, particularly in London and the South East, where demand is concentrated and existing infrastructure is most constrained.

At a high level, the UK electricity system is split between transmission and distribution networks, coordinated by the National Energy System Operator (NESO). Transmission networks move power at high voltage and are operated by three Transmission Owners: National Grid Electricity Transmission (England and Wales), Scottish Power Transmission and SSE Transmission (Scotland). Power then feeds into lower-voltage distribution networks managed by regional Distribution Network Operators (DNOs). NESO, re-nationalised in September 2024 from the former ESO, sits above these operators, coordinating system planning and, critically for large loads, administering the transmission connection process and queue management.

The tightening of capacity is the product of both supply side and demand side shifts. On the demand side, the electrification of transport and heating is increasing baseline electricity demand, while data centres represent a step change in large, continuous loads. On the supply side, increasing volumes of renewable generation are connecting at lower voltages, changing historic one-directional power flows and increasing the need for reinforcement on local networks. The result is a 'perfect storm', where grid upgrades are required not only to accommodate rapidly rising demand, but also to adapt the network itself to a fundamentally different pattern of electricity generation and power flows.

This pressure is most visible in the connection queue. By the time the application window closed in March 2025, NESO had identified around 125 GW of connection requests, around 1.7 times the UK's total grid capacity of approximately 75 GW. Around 140 data centre projects alone were seeking transmission connections, representing approximately 50 GW of requested capacity. While not all projects will proceed, the scale of oversubscription has reinforced the need to distinguish credible demand from speculative applications. Reflecting this, NESO's connection reforms have allocated around 11 GW of demand connections due to be connected before 2030, providing a more realistic indication of the capacity likely to be delivered over the remainder of the decade.

Cost of power and competitiveness gap

Cost is the other side of the power access equation. For data centres, electricity pricing is driven not only by wholesale power markets but also by a growing layer of regulated and policy-related charges, including transmission and distribution network charges and system-balancing costs. This matters because the UK's non domestic electricity prices remain elevated by recent historical standards, even as wholesale markets have eased since the 2022 energy crisis.

In a European context, the UK's cost position weakens its relative location appeal for power intensive data centres. Wholesale electricity prices across Europe have diverged sharply since 2021, with Nordic and Iberian markets benefiting from stronger renewable penetration and interconnection, translating into materially lower average prices (fig 4). While wholesale prices are only one component of the delivered tariff, the UK consistently sits towards the higher end of the European range. For developers and occupiers running high-utilisation compute, these differentials feed straight into total cost of ownership, making power cost, alongside latency and connectivity, a decisive factor in location selection.

As connection reform favours 'ready and needed' projects, power-cost credibility is becoming a competitive differentiator. Actual electricity outcomes vary widely by site, driven by regional network charges, voltage levels, and contract structures, including the balance between fixed and pass-through costs. In practice, this means that securing capacity alone is no longer sufficient. As the bar rises on deliverability, investors and operators increasingly expect schemes to demonstrate not just a viable connection pathway, but a clear, financeable strategy to manage long-term power costs. Those able to evidence both are materially better positioned to progress in a constrained and more selective UK grid environment.

125 GW

DC PROJECTS IN THE TRANSMISSION QUEUE

UK DATA CENTRE - BUILDING CAPACITY

From queue to strategy: reforming access and reshaping geography

In response, the UK is moving away from a pure 'first come, first served' model and towards prioritising projects that are demonstrably deliverable. Connection reform began with consultations in 2023 and was implemented from January 2025, alongside a temporary pause on new transmission applications (with demand-side windows closing in March 2025). Under the reformed approach, connection rights are more tightly linked to real-world deliverability, most notably through strengthened evidence requirements around land control (including defined site boundaries and, at a minimum, multi-year option agreements). Queue ordering is also increasingly influenced by planning status and demonstrable need, with a two-gate framework separating projects that meet the new criteria (Gate 2) from those placed in a holding position without a firm connection date or price (Gate 1). The practical effect is to raise the bar for data centre developers: power strategy must now be built around secured land, a credible consenting pathway and robust, defensible load assumptions.

Alongside NESO's queue reordering, Ofgem, DESNZ and NESO are implementing a broader Demand Connections Reform programme built around three pillars: Curate, Plan and Connect. Together, these reforms aim to remove speculative projects from the queue, prioritise strategically important demand, including data centres, and accelerate the delivery of viable connections. This is expected to translate into stronger queue entry and progression requirements, including clearer evidence of planning progress, financial backing and land rights, alongside mechanisms that discourage non credible applications. A further proposal under consideration is requiring applicants to demonstrate committed end user demand before securing capacity, helping ensure that grid access is allocated to genuine projects rather than land and power speculation. With application windows unlikely to reopen before late 2026, the current pause provides developers with an opportunity to complete technical due diligence, refine power requirements and strengthen commercial arrangements ahead of the next allocation round.

These reforms also sharpen the geographic dimension of UK data centre delivery. In England and Wales, distribution connections can, in some cases, support medium-scale loads, but practical limits mean the largest campuses often require transmission-level access. In Scotland, where transmission responsibilities extend further down the voltage hierarchy, large data centres are more directly dependent on transmission applications, making the current closure of the transmission window particularly constraining. Looking ahead, further prioritisation is expected to focus on areas where hyperscalers already cluster (often referred to as 'availability zones') and on government-designated growth locations, including emerging AI growth zones. This suggests that 'where to build' is increasingly intertwined with national infrastructure planning: developers may need to align site choices not only with fibre and latency, but with where grid reinforcement is planned, where policy support is clearest, and where delivery can be evidenced earliest.

To support more proactive decision-making, NESO is also developing a multi-tier spatial energy planning framework spanning local-area energy plans, regional energy spatial plans, and a national strategic plan. The intent is to model the system to 2050 and coordinate reinforcement ahead of need, rather than reacting once connection requests arrive. For data centres, this matters because it can improve transparency around where capacity will emerge and where constraints are likely to persist, information that can be used earlier in the development cycle to reduce aborted site costs, inform master planning and engage planning authorities with a clearer evidence base on infrastructure feasibility.

> By prioritising credible, well-prepared projects, the UK can unlock investment faster and give developers clearer routes to power.
>
> **PHIL PEARSON**

FIG 4: AVERAGE EUROPEAN WHOLESALE ELECTRICITY PRICE

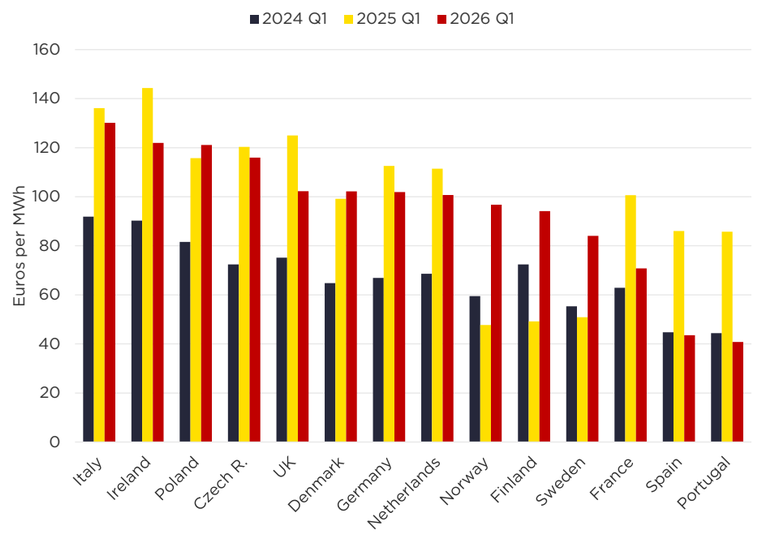

Source: Savills Research based on Ember

Beyond the red line: The land consents that make or break data centre delivery

Land is no longer priced simply on location or scale, but on deliverability; the speed at which planning certainty and critical off site land rights can be secured to make the project a reality.

Policy tailwinds and new consenting routes
A major shift since late 2024 is that planning policy in England has moved from treating data centres implicitly to explicitly referencing them. The updated National Planning Policy Framework (NPPF) requires local planning authorities to consider data centres when preparing local plans and determining applications, giving greater weight to demonstrated need and to the sector's specific locational requirements.

This policy direction strengthens the strategic case for development, but it does not eliminate scrutiny. The land strategy still needs to show how a proposal fits its setting and how key local impacts are mitigated, particularly in constrained markets, where competing uses, limited industrial land, and community concerns (See Social Impact section of the report) can intensify the planning debate.

At the same time, consenting routes are broadening. Larger schemes can opt into the Nationally Significant Infrastructure Project (NSIP) regime and seek approval via a nationally determined Development Consent Order (DCO), while AI Growth Zones are being positioned as another acceleration mechanism for major AI-related infrastructure. In combination, these changes mean consent strategy is increasingly about selecting the right route for scale and context, and assembling evidence and landing early enough to withstand potential challenges.

The hidden critical path: land rights beyond the red line
Even with supportive policy, projects can stall on land consents that sit outside the site boundary. In practice, the most time-sensitive agreements often relate to third party land needed for cable and pipe routes, surveys, temporary working areas and construction access. When this work starts late, sponsors can face programme risk and inflated compensation expectations, especially where a small number of owners control unavoidable corridors.

Best practice is to run land rights as an early, parallel workstream: map ownership, test constraints through targeted surveys, and secure voluntary rights, such as easements or option agreements, before construction timetables harden. Where voluntary agreements cannot be reached on reasonable terms, a Compulsory Purchase Order (CPO) can provide a statutory back stop, but it is not a quick fix. It requires early preparation, a clear record of reasonable engagement and offers, and acceptance that determinations can be uncertain. Used well, CPO preparation reduces exposure to last minute 'ransom' situations while keeping negotiated solutions as the preferred route.

Because these consents touch multiple parties, developers, contractors, utilities and landowners, risk allocation matters. Misalignment in approach, documentation, or the timetable can create avoidable exposure. Treating land consents as a governed package, with clear decision points and realistic lead times, is increasingly central to protecting delivery and value.

> "Critical land is rarely the Data Centre site itself, but the corridor that unlocks it. If the land rights within the corridor are not secured early, the programme to operate is not real.
>
> **JAMES PERNA**

Investor and developer checklist: what to underwrite

- Start beyond the boundary early: treat third party land rights, surveys, access and reinstatement as a critical path workstream, not a post design task.
- Price negotiation and delay risk: assume corridor control can translate into schedule pressure and higher compensation expectations if engagement starts late.
- Use CPO strategically (not reactively): if a Compulsory Purchase Order may be needed as a back stop, begin evidence-building early and plan for timetable uncertainty.
- Choose the consenting route deliberately: align project scale with the most appropriate pathway (local planning vs NSIP/DCO vs acceleration mechanisms) and tailor the evidence accordingly.
- Govern interfaces and accountability: ensure delivery partners share the same approach to land rights, documentation and decision-making so the programme is not undermined by misalignment.

From pipeline growth to execution risk

Bigger pipelines, tighter capacity: execution is the new risk.

Cost escalation across the delivery chain
As the UK data centre pipeline scales rapidly, the sector is entering a new phase where the main challenge is no longer simply securing planning consent or identifying demand, but delivering increasingly complex schemes within an environment defined by rising costs, constrained supply chains and tightening construction capacity. Delivery risk has therefore become a defining feature of the market and, increasingly, a key differentiator between projects that remain theoretical and those capable of reaching operation within investable timeframes.

Construction costs across the UK continue to rise, broadly in line with wider EMEA trends. According to Turner & Townsend data, London remains one of the most expensive data centre markets in Europe, with the average costs at US$12.02/W in 2025, up from US$11.20/W in 2024, representing an annual growth of approximately 7.3%. Over the same period, the EMEA average increased from US$10.16/W to US$10.92/W, suggesting that while cost inflation is now relatively broad-based across the region, London still commands a premium of roughly 10% above the wider European average.

Several structural factors continue to underpin this upward pressure. Labour shortages remain particularly acute across specialist electrical and mechanical disciplines, while inflation in core materials such as steel, copper and concrete continues to feed through into project costs.

At the same time, the technical complexity of new facilities is increasing materially. Higher rack densities, more sophisticated cooling requirements, greater redundancy expectations and tighter sustainability standards are all contributing to more capital-intensive developments. In parallel, lead times for critical components, including transformers, generators and cooling equipment, remain elevated despite the easing of some of the global supply chain disruption seen during 2021 and 2022.

Execution risk is increasingly becoming a grid-industrialisation issue rather than a pure construction challenge. Multi-year lead times for transformers, switchgear and cable packages mean delivery programmes must be structured around equipment availability and utility sequencing, not idealised build schedules. In practice, this shifts critical-path control upstream into electrical infrastructure procurement and grid interface planning.

FIG 5: CONSTRUCTION COSTS IN EUROPE

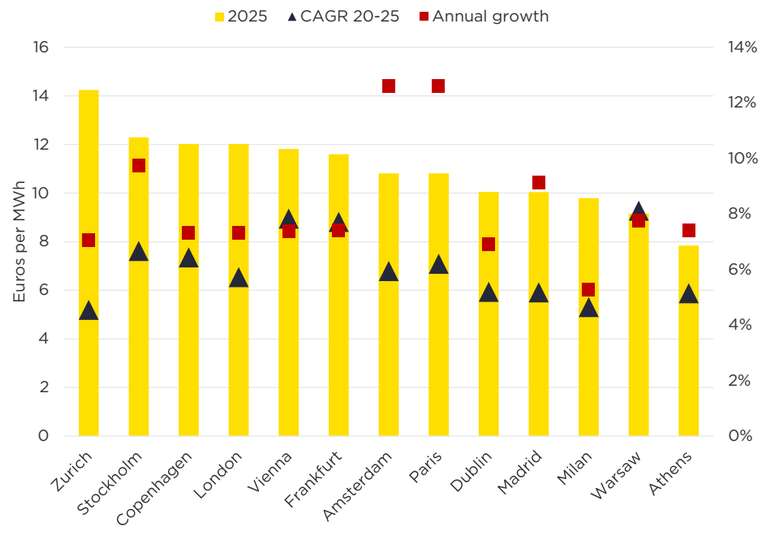

Source: Savills Research based on Turner & Townsend

> Construction is no longer the rate limiting step in isolation; it is the interaction between cost inflation, supply chain friction and delivery certainty that now defines risk in the UK data centre pipeline
>
> **MARC EDMONDSON**

A widening gap between scale and deliverability

Labour availability has also become a central constraint shaping construction risk. Although the UK benefits from a relatively mature contractor ecosystem and deeper technical expertise than many secondary European markets, demand for specialist skills continues to exceed supply. Competition for experienced commissioning engineers, project managers and electrical specialists has intensified significantly, contributing to wage inflation and increasing reliance on a limited pool of Tier 1 contractors.

This dynamic is reinforcing an increasingly polarised construction landscape. Larger operators and hyperscalers are generally better positioned to secure labour and equipment capacity through long-term framework agreements, forward procurement and established supplier relationships. By contrast, smaller developers and new market entrants face greater exposure to programme uncertainty and contractor availability. As a result, delivery certainty is becoming increasingly correlated with balance sheet strength, procurement capability and operational scale, rather than planning status alone.

Programme risk has consequently risen materially over the past two years. Although headline construction durations have not shifted dramatically, variability around delivery schedules has widened, particularly during fit-out and commissioning phases. Delays are most commonly associated with late delivery of electrical infrastructure, revisions linked to evolving compute density requirements, coordination issues between contractors and utility providers, and labour shortages during peak delivery periods.

Against this backdrop, the most effective risk-mitigation tool is often not contingency but disciplined change control. Late shifts in density, cooling architecture or technical specifications can quickly erode the benefits of modular design, forward procurement and contractor alignment. Governance from project initiation is therefore becoming more important than programme float in preserving cost certainty and delivery credibility.

Execution capability as the new competitive advantage

In response, developers are placing greater emphasis on resilience and execution certainty, even where this extends overall programme timelines. This is driving a decisive shift towards earlier capital commitment, procurement-led delivery strategies, modular construction approaches and deeper supplier integration, alongside greater design flexibility capable of accommodating future technological evolution. In a market where demand fundamentals remain exceptionally strong but delivery conditions are becoming increasingly complex, competitive advantage is progressively shifting towards operators capable not only of securing power and land, but also of demonstrating that projects can realistically be delivered at scale, on time and within budget.

For investors, construction risk is now a core underwriting consideration rather than a secondary development variable. Rising build costs, longer procurement cycles and increasing programme variability directly influence income timing, yield assumptions and exit liquidity. Consequently, the ability to demonstrate credible execution capability is becoming as important as location or power availability in assessing scheme viability.

Underwriting also needs to recognise commissioning as a revenue-critical phase rather than a purely technical close-out exercise. Defects resolution, security compliance, owner-furnished equipment integration and energisation windows can all delay handover and defer income, even where core construction is substantially complete. As a result, commissioning risk should be assessed as a commercial exposure with direct implications for lease commencement, stabilisation timing and exit assumptions.

FIG 6: FACTORS DELAYING OR PREVENTING AI INFRASTRUCTURE DEPLOYMENT

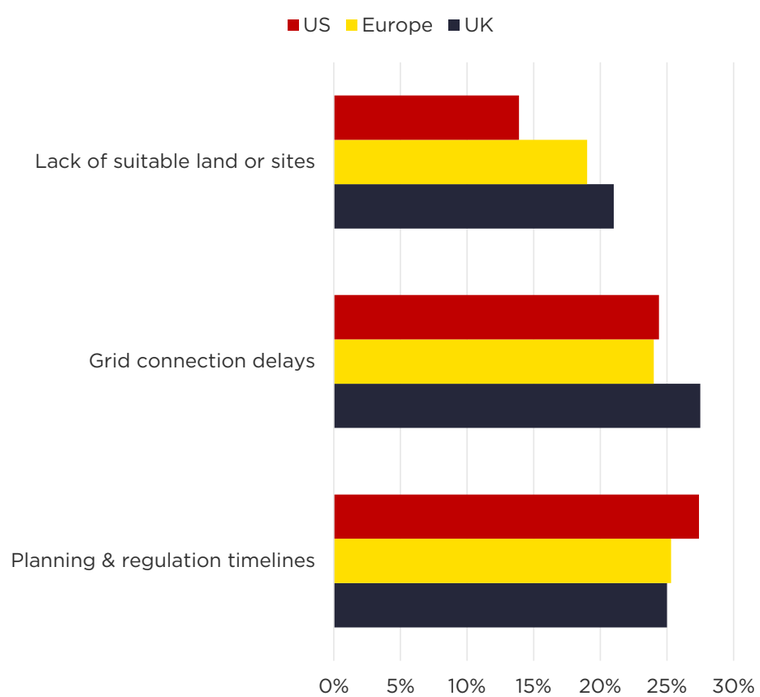

Source: Savills Research based on CUDO

The gap between strategic value and local acceptance

How social impact is shaping the consent and viability of data centre development.

> Failure to address community concerns and demonstrate local benefits at an early stage can materially increase planning risk, delay consent and, in some cases, undermine scheme viability.
>
> **GEORGINA NIMMO**

Data centres are now a core component of the UK's digital and economic infrastructure, underpinning cloud computing, artificial intelligence and the delivery of essential public and private services. Their designation as Critical National Infrastructure reflects this strategic role. At the same time, the rapid expansion of data-centre capacity has elevated social-impact considerations in planning and investment decision-making. In the UK context, these impacts are best assessed through two complementary lenses: community acceptance and associated planning risk, often described as the 'social licence to operate', and the socio-economic contribution generated by data-centre development at national and local levels.

From a community perspective, data centres present a challenging proposition. While nationally strategic, they are often perceived locally as large, opaque and resource-intensive assets, with limited visible employment relative to their physical footprint and power demand. Common concerns raised during the planning process include land take, visual impact, construction disruption, noise and competition for scarce grid capacity. These issues are particularly acute in semi-urban and edge-of-town locations, where data centres may sit alongside residential development or sensitive land uses. As a result, even policy-compliant schemes can face elevated planning risk when local engagement is weak or when communities and local planning authorities perceive an imbalance between local costs and benefits.

This tension sits in contrast to the sector's demonstrable economic contribution at a national level. According to techUK (2024), data centres currently contribute around £4.7 billion in gross value added (GVA) per year to the UK economy, support approximately 43,500 jobs, and generate around £640 million in annual tax revenues.

These impacts reflect the enabling role of digital infrastructure across finance, professional services, life sciences, retail and public administration rather than direct employment alone.

At a local scale, however, these macro-economic benefits are often less visible to host communities. Operational employment at individual sites is typically modest once facilities are live, reinforcing the perception that data centres deliver limited direct benefits relative to their land and energy requirements. As a result, greater emphasis is being placed on how national-level benefits translate into tangible local outcomes, including skills development,

apprenticeships, local procurement, and long-term partnerships with community groups, charities, and education providers, to ensure that investment in the sector delivers wide-ranging benefits.

Skills pathways are among the most effective ways to translate these enabling benefits into outcomes visible at the local level. In practice, this includes structured apprenticeships, work experience placements, and education partnerships linked to live developments. Ark's Union Park scheme provides a clear illustration, with the project projected to deliver 4,347 weeks of employment and training, including 2,273 weeks of apprenticeship and pre apprenticeship opportunities targeted at local residents, alongside sustained engagement with schools and careers programmes. Similar approaches are multiplying elsewhere. Ada Infrastructure's Docklands Data Centre Campus, for example, includes the delivery of a multi functional building intended to support technology training and align with Newham Council's wider Newham Data growth ambitions. At the occupier level, initiatives such as Google's Google for Education programme demonstrate how data centre operators can contribute to wider digital skills development by supporting schools with training and access to technology, thereby strengthening longer-term local skills pipelines.

In parallel, planning authorities are increasingly exploring mechanisms through which data centres can contribute more directly to local sustainability and decarbonisation objectives, thereby helping strengthen community acceptance. One such mechanism is the potential reuse of waste heat, which has attracted growing policy and local interest. In West London, the Old Oak and Park Royal Energy Network illustrates how heat reuse is being positioned within planning and regeneration policy, with the scheme designed to supply more than 9,000 homes and businesses once operational. However, heat reuse remains selective and highly context specific and should be understood as one of several possible tools to support community acceptance. A more detailed assessment is provided in the heat reuse section of this report.

Taken together, these dynamics underline that social impact is now a core component of data-centre viability in the UK. Community acceptance cannot be assumed, even in a supportive national policy environment, and planning outcomes are increasingly shaped by perceptions of fairness, visibility and local benefit. While data centres deliver substantial socio-economic value at a national level, the challenge for developers, investors and policymakers lies in bridging the gap between national strategic narratives and locally meaningful outcomes, ensuring that future data-centre growth is both economically productive and socially sustainable.

FIG 7: UK DATA CENTRE JOBS

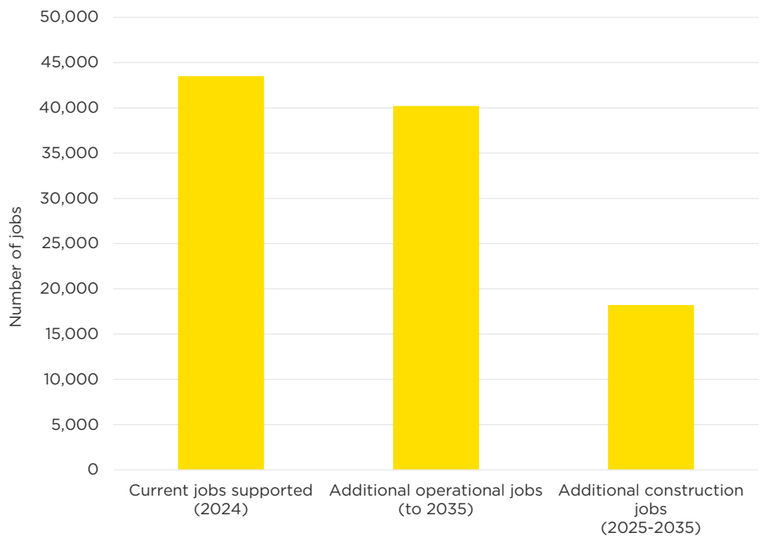

Source: techUK 'Foundations for the Future: How Data Centres Can Supercharge UK Economic Growth' (November 2024)

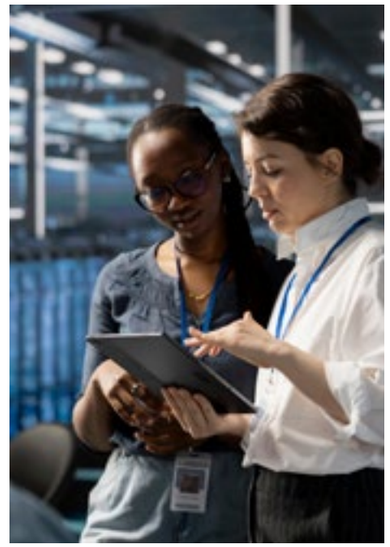

FIG 8: CASE STUDIES

Illustrative case studies	Case study 1	Case study 2
DC project & location	ARK Data Centres – Former Mercure Site & land to the south east, Tylers Way, WD25 8JH (Elstree/Hertsmere)	Land at Manor Farm, Poyle (Slough Availability Zone), Slough Borough Council
Client & role	ARK Estates 5 Ltd / Ark Data Centres – promoter/ developer; needs, economic and social value evidence prepared by Savills	Manor Farm Propco Limited – appellant; Proof of Evidence (Needs Assessment) prepared by Savills
Social impact integration	Social value positioned alongside need and economic case (jobs, skills, wellbeing and inclusion) Commitments framed around construction and long term operational benefits Local baseline and challenges used to target interventions (unemployment, youth retention, localised deprivation) Delivery model evidenced through Ark's track record (e.g., Union Park skills/apprenticeships; Spring Park community investment)	Needs case developed for a hyperscale data centre within the Slough Availability Zone (SAZ) Evidence addresses alternative sites arguments and locational requirements (power, fibre, latency, scale) Demand–supply analysis used to demonstrate a significant capacity shortfall to 2030 Case positioned as critical national infrastructure supporting the digital economy and resilience
End result	Strengthened narrative for consent by linking critical infrastructure delivery to local priorities Clear articulation of local economic benefits (construction and operational jobs; high quality roles) Defined pathway to longer term community value through partnerships, training and additionality initiatives	Evidence supports an urgent needs case within SAZ, including quantified demand and pipeline supply Positions the scheme as helping address a material capacity shortfall and the immediacy of live requirements Strengthened planning case on strategic importance and consequences of unmet need
Key takeaway	A combined needs, economic and social value narrative can materially improve deliverability by translating strategic infrastructure into tangible, locally relevant benefits	A robust, quantified needs case—grounded in availability-zone requirements and supply constraints- can help de-risk planning decisions for hyperscale proposals in sensitive locations

Source: Savills Research

Heat reuse: turning data centres into urban heat sources

Data centres run 24/7, so does their waste heat. With the right infrastructure, that by-product can become a dependable input to low carbon heating.

Heat reuse is becoming an increasingly relevant theme in UK data centre planning because it links a fast growing digital land use to two urgent public priorities: decarbonising heat and delivering locally visible community benefit. Heat networks are now a central part of the UK's transition strategy; the government expects their share of UK heat to rise from around 2% today to 18% by 2050 and has committed over £350m to support low carbon heat network deployment. In parallel, London, where around 80% of UK data centre capacity is concentrated, has both high heat-density and growing planning scrutiny of energy and carbon. This combination is pushing heat reuse from an "interesting idea" towards an increasingly practical planning consideration, particularly in locations where new development, heat network zones and data centre growth overlap and reducing overall electricity grid demand.

How it works: from server heat to street heat

Almost all electricity entering IT equipment ultimately leaves as heat. In conventional operation, cooling systems capture and reject it to ambient air or water. Heat reuse changes the destination: a portion of that heat is intercepted and exported as a usable thermal source. In practice, heat is captured from the cooling loop, passed through a plate heat exchanger into a separate offtake circuit, and lifted by water-source heat pumps before supplying a heat consumer or network. Air-cooled systems typically need this uplift, while liquid cooling can provide higher-grade heat and improve recovery economics. Typical output heat is around 30–35°C, which heat pumps can lift to around 70–80°C for space heating and domestic hot water. Advanced liquid-cooled systems can deliver higher recoverable temperatures, reducing the required lift and making direct or near-direct use more realistic in some applications.

In the UK, heat reuse most commonly involves supplying district or communal heat networks, as serving multiple buildings makes the required infrastructure viable. Typical end uses include residential space heating and domestic hot water, as well as public buildings such as hospitals, universities, schools and leisure centres. Data centre heat can also be used on-site or campus-scale, for instance, to heat offices, support domestic hot water, or reduce reliance on gas boilers when simultaneous heating and cooling loads exist. Beyond heat networks, newer concepts are emerging: "edge" and micro data centres positioned close to heat demands (including new developments) to shorten pipe runs and improve viability; and diversification into non building heat demands (e.g., leisure facilities or controlled-environment agriculture), which can help balance seasonal demand profiles.

The hard realities of pipes, risk and payback

The core constraint is that most data centre waste heat remains low grade, meaning extraction at scale typically relies on heat pumps, additional electrical input and significant pipework, particularly where temperature differentials are limited. Viability therefore depends on proximity to dense, stable heat demand, concentrating opportunities in urban locations where data centres and large residential or institutional uses overlap.

However, new generations of ambient loop heat networks offer a more flexible model. By circulating low temperature water, these systems can harness low grade data centre heat directly, with distributed end user heat pumps uplifting temperatures for heating or reversing to provide cooling, maximising utilisation of this resource.

The primary barrier is not the technology, but the cost, routing, type and ownership of the infrastructure required to connect supply with a viable offtake.

> "Capturing this surplus heat turns a waste stream of the digital economy into a strategic asset, decarbonising existing heat networks and underpinning the next generation of low carbon district heating.
>
> **NICK GREEN**

Operationally, data centres have very low risk tolerance: they must retain heat-rejection capacity, full cooling resilience and avoid disruption that could threaten uptime or customer SLAs. Commercially, heat network developers seek long-term certainty, while data centre heat availability can be difficult to guarantee because colocation occupancy may ramp up over years, IT utilisation can change, and customers may limit disclosure of future loads. Seasonality also matters: data centres produce heat year-round, while demand peaks in winter, so viable schemes usually need diversified heat demand, storage or complementary loads. The UK's relatively immature heat network market also means fewer ready-made networks and more coordination between developers, operators and local authorities.

London-led projects and policy tailwinds

Heat reuse is moving from pilots to early delivery in the UK, with London offering the strongest cluster of opportunities because data centres and heat demand are co-located. The Greater London Authority commissioned a study estimating theoretical potential to recover up to 1.6 TWh of heat annually across 27 modelled networks, enough to supply around

UK DATA CENTRE - BUILDING CAPACITY

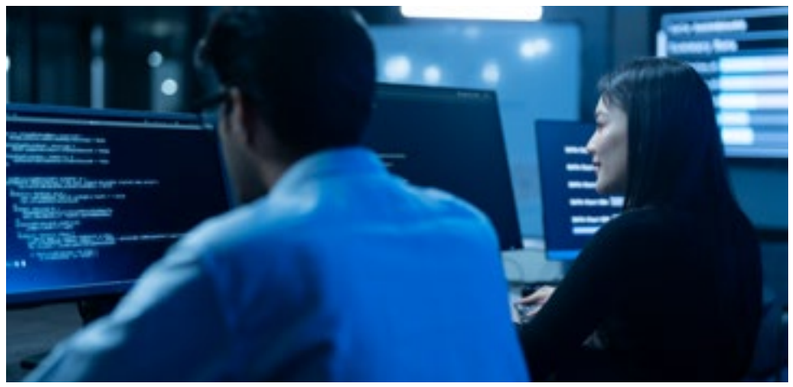

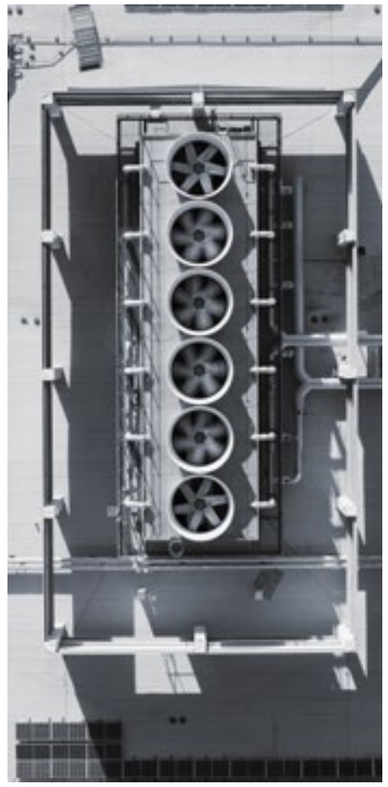

350,000 homes. Nationally, government support is helping first-of-a-kind schemes: in November 2023, a £36m award backed a West London district heating system intended to share data centre waste heat with up to 10,000 new homes.

Policy is also beginning to materially shift the market. Imminent Energy Act zoning regulations are expected to accelerate investment, expansion and delivery of district heat networks by identifying priority areas and strengthening the case for connection. The City of London has already established a fast tracked zoning programme and is actively progressing network development, including consideration of data centre waste heat integration.

However, these figures are potential rather than committed supply: delivery depends on network build-out, offtake agreements, timing and facility configuration. Planning is starting to respond, with London borough approaches and area guidance increasingly expecting proposals to be "heat export ready", safeguarding space and routes for offtake equipment and pipework.

Momentum is also being reinforced by changing technology and measurement. Industry stakeholders report average annual IT utilisation of around 75% and up to 90% for AI workloads, according to the Greater London Authority report, and a broader shift towards liquid cooling as rack densities increase, both of which tend to increase the quantity and temperature of recoverable heat. Typical recoverable temperatures span roughly 15–25°C for direct/indirect air-cooled systems, 20–30°C for liquid-to-air systems, and 30–60°C for liquid-to-liquid systems, with higher-temperature liquid cooling offering "very high" suitability for heat recovery. This matters because higher source temperatures improve heat pump performance or allow direct heat supplied to networks further reducing operating cost and emissions.

At the same time, reporting frameworks are making heat reuse more visible: alongside PUE, the Energy Reuse Factor (ERF) is emerging as a key indicator, with ISO/IEC 30134-6:2021 providing a standard method for its calculation; one reference estimate suggests ERF could reach up to 69% in winter months and 86% in summer months, subject to cooling system type and ambient conditions[12]. Together, these trends are helping shift heat reuse from an optional add-on to a credible component of "net zero ready" data centre design and local energy planning.

For planners and communities, the key issue is not that every data centre must export heat from day one, but that developments are designed early on to allow for it, with the right cooling setup, space for heat connections, and links to nearby heat demand. Furthermore, district network developers should consider modern infrastructure types and plan for connection for low carbon Data Centres waste heat sources.

Capital chasing capacity

In an increasingly crowded market, control of power, pipeline and execution is becoming the key differentiator.

> What was once a niche alternative is now being priced, financed and competed for as strategic infrastructure. Larger platforms, deeper partnerships and more complex capital structures are redefining how investors compete in the sector.
>
> **LYDIA BRISSY**

The UK data centre investment market is moving further away from conventional real estate pricing logic and towards a model shaped by infrastructure-style capital, operational capability and strategic control over future capacity. Capital remains abundant, but access to product is constrained, and that imbalance is increasingly determining how investors enter the market, what they are prepared to pay and how risk is assessed. With investor appetite expanding faster than the supply of existing stabilised assets, buyers are increasingly forced to find alternative routes into the sector, including M&A, strategic minority stakes, development-led entry strategies, and powered land positions that offer a clearer path to deployment. In that sense, the market is no longer defined simply by asset ownership, but by the ability to secure deliverability, scale and long-term relevance within an increasingly competitive digital infrastructure landscape.

This depth of capital is evident globally and is feeding directly into the UK market. Private funds investing in data centres and telecom infrastructure raised record commitments in 2025, with specialist vehicles alone raising US$26 billion and broader generalist funds with digital infrastructure exposure raising a further US$131.6 billion, while total digital infrastructure assets under management with data centre and telecom exposure were estimated at US$251.7 billion, including US$65.9 billion of dry powder, according to Pitchbook. That fundraising momentum reflects a broad view among allocators that data centres sit at the intersection of structural demand growth, critical infrastructure and long-duration income. It also underlines an important shift in buyer composition. Capital targeting the sector increasingly comes from infrastructure funds, private equity, sovereign and pension capital, alongside selected real estate strategies, reinforcing the sector's transition from a niche alternative use class into a mainstream strategic allocation.

That capital is encountering a market where traditional routes to scale remain limited. Stabilised, institutional-grade assets are scarce, owners are often reluctant to sell, and single-asset transactions do not offer the critical mass sought by large-ticket investors. As a result, investors are becoming more strategic in how they establish exposure. M&A and recapitalisations are increasingly attractive because they provide not only operational estates, but also scale, customer footprint, pipelines and, crucially, a degree of control over future execution. The global rise in data centre M&A illustrates this clearly: deal value increased from US$26 billion in 2023 to US$77 billion in 2024 before reaching US$69 billion in 2025, following a much lower base of US$7.5 billion in 2015. One notable recent example is SoftBank's approximately

FIG 9: GLOBAL VALUE OF DATA CENTRE M&A

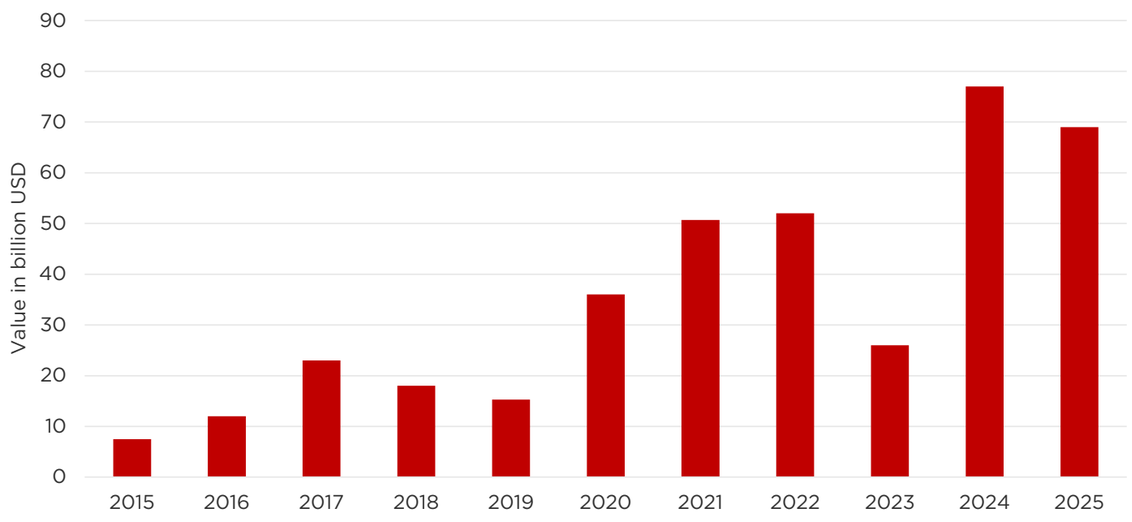

Source: Savills Research based on Synergy group, DGTIInfra

UK DATA CENTRE - BUILDING CAPACITY

> AI-led demand is reinforcing the sector's shift towards larger, more strategic and more capital-intensive investment structures. Data centres are increasingly being underwritten not simply as property, but as long-term digital and energy infrastructure.
>
> **CAMERON BELL**

$40BN

LARGEST DATA CENTRE DEAL OCT 2025

US$4 billion acquisition of DigitalBridge, announced in December 2025 and expected to be completed in the second half of 2026, as part of its broader strategy to build AI-related infrastructure at scale.

Strategic partnerships are also becoming a more prominent route into the market. For many investors, particularly those without operating expertise, joint ventures and minority positions offer a more effective way to gain exposure while mitigating execution risk. The rationale is straightforward. Underwriting now goes far beyond rent, covenant and location. It also includes power procurement, technical design, delivery capability and customer concentration. Partnerships allow investors not only to access specialist operating expertise but also to assemble the scale and financing capacity required to secure and deliver increasingly capital-intensive opportunities. A clear example of that shift is the October 2025 acquisition of Aligned Data Centers, at an implied enterprise value of around US$40 billion, the largest data centre deal signed so far. Led by a consortium including Global Infrastructure Partners, MGX and the AI Infrastructure Partnership, the transaction highlights the growing convergence between infrastructure capital, private equity-style platform strategies and hyperscaler-driven AI ecosystems, as investors increasingly seek direct exposure to long-term compute demand, power-secured development pipelines and operational digital infrastructure platforms. The deal also illustrates how data centres are progressively being viewed not simply as real estate assets, but as strategic AI and energy infrastructure.

That dynamic is particularly visible in the UK land market, where a meaningful share of recent activity has centred on development-led entry and campus. With stabilised assets scarce and hyperscale demand still ahead of new delivery, value is moving upstream towards control of land, power, and pipelines. This also makes entry more complex and capital-intensive, increasing the importance of partnerships that can combine funding, development capability and operational delivery. In this environment, competitive advantage depends less on capital access alone and more on the ability to secure and execute capacity at scale. Recent transactions also illustrate the growing dominance of development-led investment strategies in the UK market. In 2025, SEGRO and Pure Data Centres launched a £1 billion joint venture to develop a 56MW hyperscale campus at Premier Park in West London, with around 70MVA of secured power capacity already in place and phased delivery expected from 2028. Similarly, Colt Data Centre Services secured approval in late 2025 for a £2.5 billion expansion of its Hayes campus, adding 97MW of IT load with grid delivery expected by 2027 and first operations targeted for 2029. Meanwhile, the May 2026 acquisition by Stellanor Datacenters of eight UK facilities from Redcentric for £123 million was accompanied by plans to upgrade the assets into higher-density AI-ready facilities, illustrating how investors are increasingly targeting assets with secured power capacity and future repositioning potential

This has important implications for underwriting. Investors are increasingly differentiating between nominal site availability and genuinely bankable development opportunities. Where future delivery is supported by credible power arrangements, site control, or a realistic route through planning, pricing can remain aggressive even in a higher-cost-of-capital backdrop. Where those elements are uncertain, capital is becoming more selective. This helps explain why power, planning, and construction pressures matter chiefly to investors as pricing variables rather than as abstract market constraints. They affect the probability of delivering capacity, the timing of capital deployment and the degree of execution risk embedded in a deal. In parts of the London market in particular, access to power has become central to value formation, and sites with a stronger prospect of energisation are commanding a strategic premium relative to land with weaker or more speculative deliverability.

Financing structures are becoming more sophisticated. Across the sector, developers and operators are increasingly recycling capital through stabilised asset monetisation, portfolio refinancings and structured debt solutions, allowing capital to move more quickly through the development cycle. This is attracting a broader lender and investor base, from private credit and debt funds to insurers and infrastructure investors seeking senior exposure to long-duration cash flow. For the UK market, this matters because it broadens the pool of capital able to participate, even where direct acquisition opportunities remain limited. It also supports a wider range of risk appetites, from development and lease-up risk through to lower-risk income-backed strategies. As more assets mature and

require refinancing, debt capital is likely to become a more visible component of the competitive landscape, not simply a source of leverage but an alternative route to market participation.

The small pool of standing assets with long leases, high occupancy and strong counterparties remains in strong demand from core income-focused capital. Recent UK transactions underline that point. In October 2025, Stellanor agreed to acquire Redcentric Data Centres for up to £127 million, adding eight operating facilities to its platform, with completion following in May 2026. In September 2025, BlackRock and Digital Gravity Partners seeded their Gravity Edge venture with the acquisition of an enterprise data centre in West London, supported by an initial investment of more than £100 million. Although the market remains thin, the appetite for secure income and operational platforms is clearly intact. Where stabilised assets do trade, pricing continues to reflect both their scarcity and their increasingly infrastructure-like income profile. Prime UK yields have remained broadly stable over the past year, following a 50 basis point outward shift between 2024 and 2025, and currently stand in the 5.5% to 6.5% range. That is around 50 basis points softer than the wider European range, reflecting the relative movement in UK gilt rates versus Eurozone government bond yields.

Overall, investor strategy is therefore diverging by mandate and capability. Core income capital remains focused on stabilised assets with long leases and strong counterparties, but the shortage of such products is forcing many institutions to look earlier in the lifecycle or to back experienced operators through joint ventures and recapitalisations. Infrastructure and private equity capital are generally more willing to underwrite development, platform complexity and corporate transactions where scale and operational upside can be captured. Real estate capital remains active, but the old distinction between real estate and infrastructure has become less useful in practice, as both increasingly use overlapping ownership and financing models. In effect, the key differentiator is no longer asset class label but the ability to assess technical risk, secure specialist partnerships, and commit capital at the pace required by a highly competitive market.

FIG 10: PRIME YIELDS

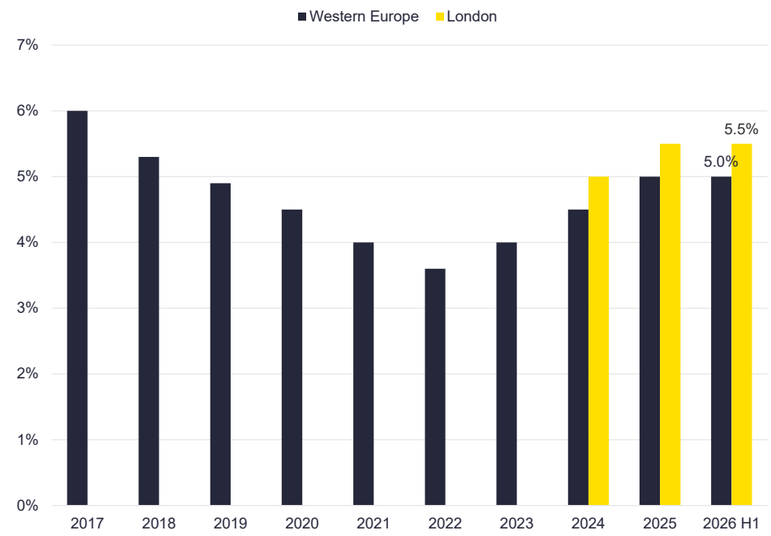

Source: techUK 'Foundations for the Future: How Data Centres Can Supercharge UK Economic Growth' (November 2024)

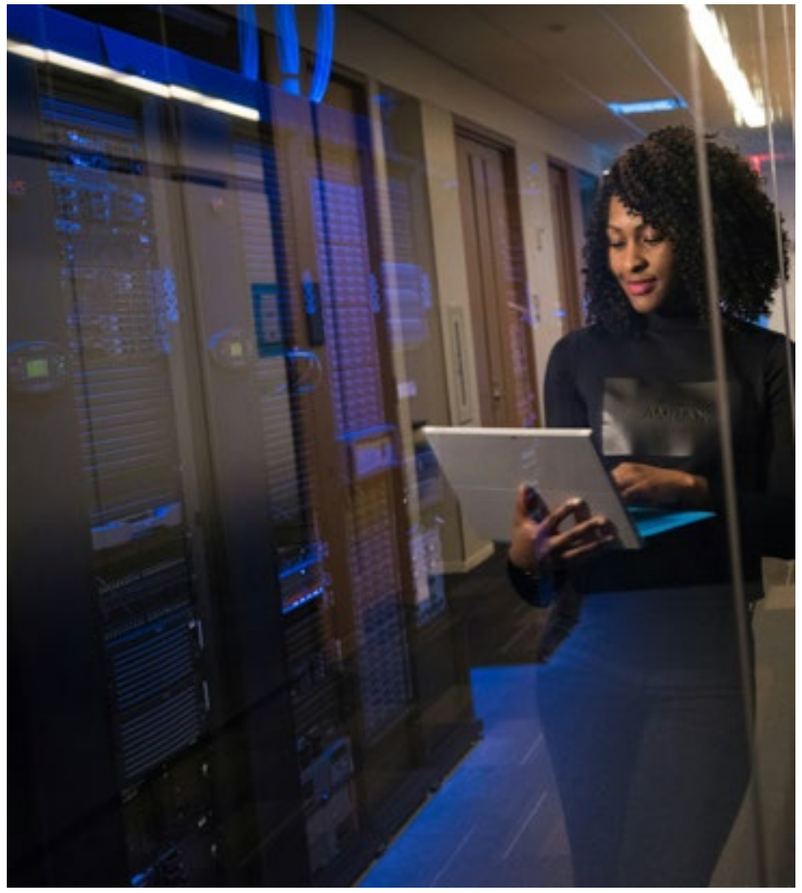

Delivering the next wave

The outlook remains positive, but growth will increasingly be filtered by power, planning and execution.

The UK data centre market is likely to remain structurally tight rather than cyclical. Demand should continue to be underpinned by three overlapping drivers: hyperscale cloud expansion, AI-related deployments and a resilient enterprise-colocation base. London will remain the dominant market because its availability zones, fibre density and cloud ecosystems are difficult to replicate, but that same concentration will continue to intensify pressure on land, power and delivery timelines. In practice, this means that take-up is likely to remain constrained less by appetite than by the availability of genuinely deliverable capacity. Even where headline pipeline numbers appear substantial, much of that pipeline is still early-stage and subject to grid access, planning progression, and funding discipline. The market is therefore entering a phase in which pre-letting, phased delivery, and forward commitments are likely to remain normal features of occupier strategy rather than temporary responses to scarcity. This also reinforces our view that the most durable demand story is likely to shift progressively from a narrow focus on very large AI training schemes towards a broader, more persistent need for inference, hybrid cloud, and latency-sensitive deployments closer to end users.

The central issue for the next phase of growth will be execution. Power is now the primary gatekeeper of UK data centre delivery, and recent reforms to the grid connections regime are pushing the market further away from speculative queueing and towards a readiness-based model in which projects must evidence land control, planning progression and strategic relevance to secure firmer connection positions. These policy changes are directionally supportive, but they will not eliminate the underlying capacity scarcity or the long lead times required for reinforcement. Planning conditions have also become more supportive than they were two years ago, with stronger recognition of data centres in national policy and the prospect of broader consenting routes for major schemes. However, support is increasingly conditional on schemes demonstrating local legitimacy through credible mitigation of visual, environmental and infrastructure impacts, as well as clearer community value. Construction adds a further filter. Rising build costs, persistent labour shortages and long lead times for critical electrical equipment mean that not every consented or powered project will translate smoothly into operational stock. As a result, the gap between nominal pipeline and executable delivery is likely to remain one of the defining features of the market into 2027 and beyond.

For investors, this points to a market where value will continue to migrate upstream towards power-secured land, operational platforms and delivery capability rather than conventional real estate exposure alone. In our view, prime yields are likely to remain relatively firm for scarce stabilised assets with strong counterparties, while development-led strategies should continue to command attention where there is a credible route through energisation, planning and procurement. The old distinction between real estate and infrastructure capital is likely to become even less useful as ownership structures, financing routes and underwriting assumptions continue to converge. Geographically, London should retain its leadership over the medium term, but the search for alternative locations is likely to intensify, as operators seek to secure scale, lower-cost power, and a more credible path through consent. That does not imply rapid decentralisation: rather, selective regional growth should emerge where infrastructure conditions align, while the London market remains the main demand anchor. Overall, the outlook is one of continued expansion, but more selective. The UK is unlikely to lack demand, capital or strategic relevance; the more important question is which projects can convert those advantages into deliverable capacity within a tighter, more regulated and more execution-sensitive environment.

> “The UK will remain one of Europe’s most strategically important data centre Markets, but future growth will become more selective, more infrastructure-led and more tightly controlled by access to power.
>
> **LYDIA BRISSY**

Savills Research

We're a dedicated team with an unrivalled reputation for producing well-informed and accurate analysis, research and commentary across all sectors of the UK property market. To view copies of our previous Spotlight publications, go to **www.savills.co.uk/insight-and-opinion/**

Cameron Bell
Director
EMEA Data Centre Advisory
+44 7870 555 982
CLBell@savills.com

Rupert Duckworth
Associate Director
EMEA Data Centre Advisory
+44 7976 185 216
rupert.duckworth@savills.com

Marc Edmondson
Director
BPC
+44 7812 085 405
marc.edmondson@savills.com

Niamh Doherty
Associate
Investment Advisory EMEA
+44 7815 032 102
niamh.doherty@savills.com

Phil Pearson
Director
Energy, Renewables and Infrastructure
+44 7974 139440
phil.pearson@savills.com

James Perna
Director
Utilities & Infrastructure
+44 7807 999 613
JPerna@savills.com

Georgina Nimmo
Associate
Social Value
+44 7855 986 932
georgina.nimmo@savills.com

Nick Green
Director
London Planning
+44 7807 999 235
NDGreen@savills.com

Lydia Brissy
Director
European Research
+44 624 623 644
LBrissy@savills.fr

Callum Hodgson
Associate
Planning Economics
+44 7929 774 692
callum.hodgson@savills.com

Stewart Clark
Director
Valuation
+44 7812 965 344
SCClark@savills.com

Luisa Mostarda
Associate Director
Property Management
+353 86 809 5279
luisa.mostarda@savills.ie

Savills plc is a global real estate services provider listed on the London Stock Exchange. We have an international network of more than 700 offices and associates throughout the Americas, the UK, continental Europe, Asia Pacific, Africa, India and the Middle East, offering a broad range of specialist advisory, management and transactional services to clients all over the world. This report is for general informative purposes only. It may not be published, reproduced or quoted in part or in whole, nor may it be used as a basis for any contract, prospectus, agreement or other document without prior consent. While every effort has been made to ensure its accuracy, Savills accepts no liability whatsoever for any direct or consequential loss arising from its use. The content is strictly copyright and reproduction of the whole or part of it in any form is prohibited without written permission from Savills Research.