DATA CENTRES AND INFRASTRUCTURE
INVESTMENT ON THE RISE

Agnieszka Klich, Mark Thompson, Lennard Ong, Luke Ong and Tosin Oredola of Covington & Burling LLP examine the surge in private equity investment in data centres and global infrastructure assets.

Private equity investment in global infrastructure projects has seen a significant resurgence, with 2025 marking a record year for fundraising: over $200 billion was raised globally, surpassing the previous high of $180 billion in 2022. Private infrastructure assets under management reached an all-time high of $1.6 trillion in 2025, representing approximately 10% of all alternative assets.

Despite this record year, it is estimated that the world still requires an additional $106 trillion in global infrastructure investment by 2040. Sectors that require such increased investment span the energy transition, digital infrastructure (particularly data centres), transport and social infrastructure sectors.

This article looks at:

- The market context of this surge in projects.
- The key characteristics of the AI boom and energy markets that have driven the need for substantial financing in data centres.
- Some of the key issues that arise in data centre project finance transactions.
- The implications of these trends for other infrastructure projects.

MARKET CONTEXT

Within the broader infrastructure investment growth, the battle to build data centres has become the centre of attention (*see feature article "Data centres: the key issues for businesses", www.practicallaw.com/w-045-7240*). With demand outpacing supply and high growth predicted, debt and equity opportunities are available across the full development and operational lifecycle of data centres. These facilities have transitioned from niche alternative assets to a staple of modern infrastructure portfolios, turning data centres into a core infrastructure asset class. High-net-worth investors and institutional funds are particularly drawn to features such as predictable cash flows, long-term leases and the large amounts of capital that are required to meet future capacity demands.

The scale of this opportunity is reflected in market projections. By 2030, the global data centre market is projected to be worth approximately $900 billion, with the sector experiencing an infrastructure investment boom requiring up to $3 trillion by the end of the decade. Digital infrastructure is no longer solely the domain of telecoms companies. 2025 saw the largest-ever deal in the sector, with Aligned, a US data centre firm, being bought for $40 billion

Global data centre electricity consumption

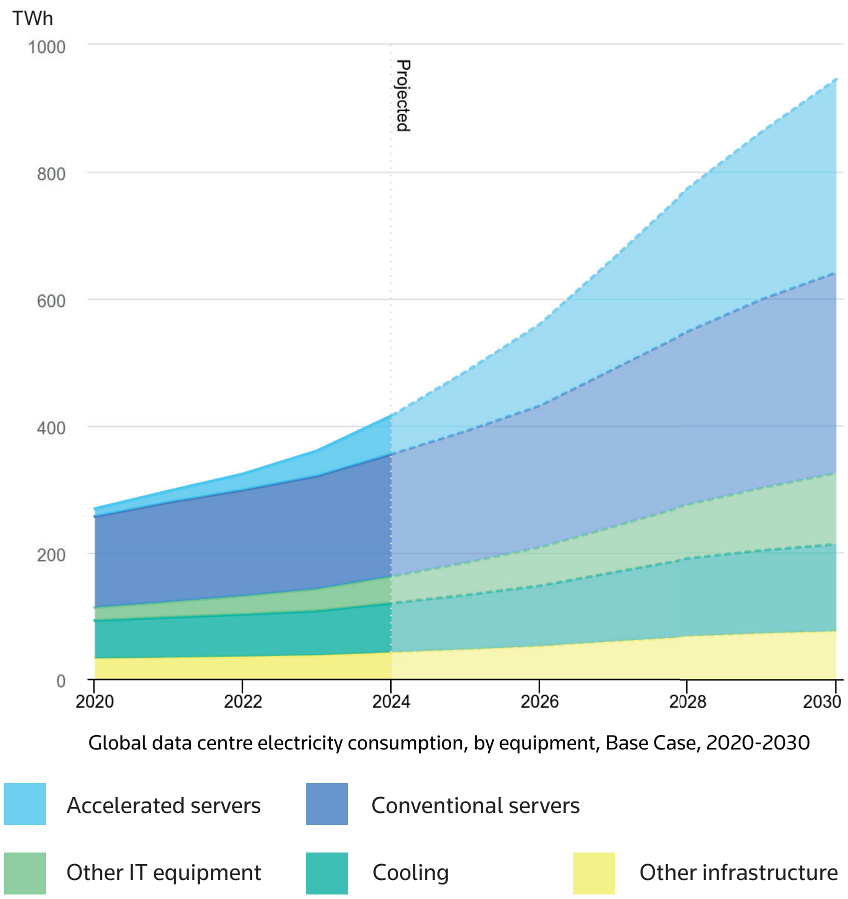

Global data centre electricity consumption, by equipment, Base Case, 2020-2030

IEA. Licence: CC BY 4.0

by a consortium that included BlackRock, Global Infrastructure Partners and Abu Dhabi's MGX. Data centre properties under construction globally have an estimated value of $550 billion.

The impact of the private equity industry, in particular, on the sector is evident, as is the appeal of the sector to private equity firms. Blackstone-backed QTS Realty Trust is now the largest independent data centre operator in the world and it is seeking to continue to grow in this area.

However, despite the surge in investment activity, entering this highly regulated asset class is complex and requires a clear understanding of transactional, policy and regulatory frameworks to ensure commercial success. The Financial Times reported that "an unprecedented 75 US data centre projects worth around $130 billion were blocked or delayed" in the first three months of 2026 (*www.ft.com/content/7800ba0f-1420-49fe-b260-9838632a19a4?syn-25a6b1a6=1*). Successfully navigating these risks requires sophisticated legal, regulatory and operational analysis throughout the investment lifecycle.

THE AI CATALYST

A key driver of this explosive growth in demand is the advent of generative AI. AI is driving up data demands, resulting in ever-larger data centre projects.

Exponential demand

Server racks for AI computing consume a lot more electricity than is generally used in cloud processes. It takes 5kW to 10kW to power a single rack of servers in a typical data centre, but a rack filled with advanced AI computing chips can demand over ten times that amount of power. This requires data centres that can support surges in workloads and more advanced cooling systems to handle the heat emitted. This in turn requires energy.

In 2025, data centres used approximately 485 terawatt hours (TWh) of electricity globally, with electricity consumption growing at around 15% per year, more than four times faster than the growth of total electricity consumption from all other sectors. The International Energy Agency projects that global electricity consumption for data centres will more than double to approximately 945 TWh by 2030, representing nearly 3% of total global electricity demand (*see box "Global data centre electricity consumption"*).

The hyperscaler race

The competitive dynamics among the hyperscalers (those companies that run enormous data centre networks, allowing them to offer computing resources and cloud services at a scale that smaller providers cannot match) reflect these energy pressures and are intensifying. Amazon has announced plans to spend almost $150 billion over the coming 15 years on data centres in a bid to maintain its position in the cloud services market, where it holds about twice the market share of its nearest competitor, Microsoft.

These demand pressures are also reflected in the pricing for the use of the data centres. Bloomberg reports that hyperscalers are now paying data centre operators monthly rates exceeding $100 per kilowatt in some regions, reaching approximately $150 per kilowatt in top-tier markets, compared to $70 to $80 per kilowatt just three years ago (*www.bloomberg.com/graphics/2025-ai-data-centers-electricity-prices/*). Because data centres provide access to power, rents are typically measured in kilowatts rather than square feet.

THE ENERGY NEXUS

The AI-driven explosion in data centre demand is having profound effects on global electricity markets. After years of stagnation, global electricity demand has recently started to surge, largely due to the exponential roll-out of data centres. The New York Times has reported that demand for electricity in the US, which stayed largely flat for two decades, has begun to surge as utilities confront an explosion in data centres, a resurgence in manufacturing and the connection of millions of electric vehicles (*www.nytimes.com/interactive/2024/03/13/climate/electric-power-climate-change.html*).

This surge in energy demand creates a significant tension with corporate sustainability commitments (*see "ESG and sustainability pressures" below*). Hyperscalers have pledged to power their operations with emissions-free power by 2030, but those commitments were made before the AI boom. Meta, Google, Amazon, and Microsoft have pledged to use only renewable energy and reach net-zero carbon emissions within the next two decades (*see box "Google and greenhouse gas emissions"*).

Hyperscalers now face an acute need to procure sustainable sources of energy to power their data centres at an unprecedented scale. In response, they are increasingly bypassing traditional grid connection processes in various ways. Technology companies and infrastructure fund portfolio companies are entering into corporate power purchase agreements (PPAs) to secure the supply of huge amounts of sustainable energy to power data centres, as well as trying to bring old nuclear plants back online.

Small modular reactors have also emerged as a promising solution for powering data centres with clean, reliable energy. Some fracking companies are switching to extract geothermal energy, rather than gas, to supply technology companies. On 20 September 2024, Constellation Energy announced the signing of a 20-year PPA with Microsoft that will pave the way for the restart of Three Mile Island Unit 1 (*www.constellationenergy.com/news/2024/Constellation-to-Launch-Crane-Clean-Energy-Center-Restoring-Jobs-and-Carbon-Free-Power-to-The-Grid.html*). As of 2026, analysts indicate that US data centre developers have announced around 100GW of off-grid gas generation.

These challenges are driving significant investment in green data centre solutions. The global green data centre market is projected to reach approximately $155.75 billion by 2030, up from $48.26 billion in 2025, representing a compound annual growth rate of 26.4%. On-site renewable energy production, such as wind or solar power, can decrease a data centre's carbon impact, electricity costs and burden on local power grids. Liquid cooling can reduce a data centre's overall power consumption and improve its power usage effectiveness. The global market for data centre liquid cooling was valued at approximately $5.5 billion in 2025 and is projected to reach $19 billion by 2030, according to industry estimates. However, these green solutions are not a silver bullet and bring their own legal and regulatory considerations. Liquid cooling systems, especially traditional evaporative cooling systems, can increase water consumption, requiring operators to secure water abstraction licences and comply with local water resource management regulations. On-site renewable generation may require planning permissions and grid connection agreements, while battery storage installations raise fire safety, hazardous materials and decommissioning obligations.

More broadly, the increased investment in data centres has yielded other issues. Supply chains have tightened for energy technologies such as gas turbines and transformers, as well as for advanced chips and IT components. The growing pipeline of data centre projects is also straining planning and regulatory systems, holding up grid connections and other necessary approvals. As of the end of 2025, over 2,060GW of generation and storage capacity were actively seeking connection to the US grid, according to the Lawrence Berkeley National Laboratory, with the median duration from interconnection request to commercial operation now exceeding five years (*https://emp.lbl.gov/queues*).

Google and greenhouse gas emissions

Google reported in its 2025 Environmental Report that its "ambition-based" greenhouse gas emissions increased by 11% in 2024 and were 51% higher than its 2019 baseline (*https://sustainability.google/reports/google-2025-environmental-report/*).

This increase occurred despite a 12% reduction in emissions associated with energy used by its data centres during 2024. The rise in overall emissions was attributed largely to the growing energy and supply-chain demands associated with the company's expansion, including the development of AI infrastructure.

Google has committed to halving its emissions by 2030, relative to its 2019 baseline, achieving net-zero emissions across its operations and value chain, and operating on carbon-free energy at all times. However, the company has acknowledged that rapid growth in AI-related electricity demand, the pace of clean-energy deployment and difficulties in decarbonising its supply chain are complicating its progress towards those targets.

FINANCING DATA CENTRES

The massive capital requirements of data centre development are also reshaping how these projects are financed.

Balance sheet or project financing

Data centres have historically been developed and financed on balance sheet by companies in real estate, facilities management and telecommunications sectors; that is, the amount the lenders have loaned was based on the balance sheet of the company building the asset. As outlined, the AI boom has resulted in ever-larger data centre projects. This in turn has forced market players to explore debt solutions.

Non-recourse, or limited-recourse, financing means that the lender's ability to recover its money if the borrower defaults is limited only to a specific asset or collateral tied to the loan; the lender has no recourse against the borrower's other assets. Lenders will finance a project based on its expected cash flows, in the form of the revenue from an offtake agreement where a buyer contracts to buy, or "offtake", the data centre's capacity, rather than the borrower's overall balance sheet. This type of financing will be key to meeting the exponential growth of data centre capacity.

Conflicting demands

Meeting this financing challenge will require a combination of specialisms that have historically operated independently from one another: real estate financing, energy and infrastructure project financing (particularly distributed solar and district cooling and heating), and, as applicable, the constraints and demands of large technology companies. Reconciling lender expectations with hyperscaler expectations and data centre market standards is a key challenge.

A core tension arises from differing priorities between lenders and the hyperscalers. Cash-rich hyperscalers have little regard for project or structured financing because the speed of deployment is the overriding priority for them. Data centre market standards for finance

Typical parties to a data centre investment

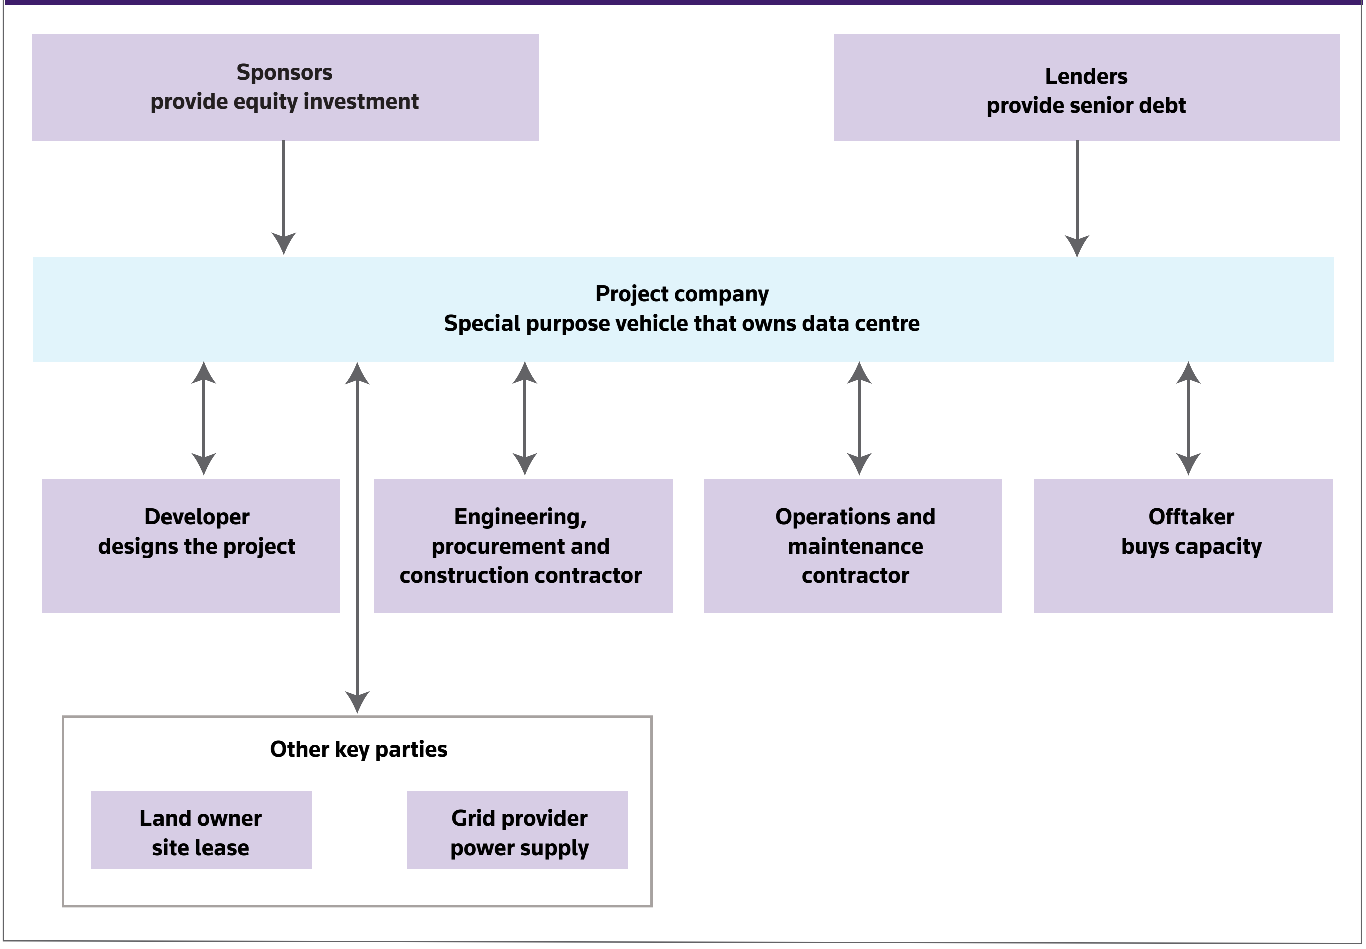

arrangements, such as short duration offtake agreements, are at odds with traditional project financing market standards that usually assume a long-term agreement. On a project finance transaction, lenders often want the offtake term to exceed the term of the debt (a "tail"), so as to cushion the servicing of the debt. Data centre offtake terms tend to be significantly shorter and can range from three to seven years. By contrast, terms of PPAs with US state utilities are typically 15 to 25 years.

The issue is that companies, including hyperscalers, sometimes have policies against entering into contracts of more than a certain number of years' duration. As such, the likelihood of offtake term renewal or non-renewal will be a core issue for the viability of the project finance. Lenders will look into the likelihood of a data centre offtaker (the buyer of the data centre's services) extending the term of their offtake agreement, which typically involves obtaining a market analysis to assess the risk of non-renewal.

Accordion finance

To address these structural challenges, one innovation that is gaining traction is accordion financing: this is a flexible debt financing structure that allows additional assets to be integrated incrementally into the same financing arrangement, as the development pipeline expands. Having financing conditions pre-agreed between a borrower and a lender accelerates the financing process and, ultimately, debt availability.

This is a financing structure previously used in connection with distributed solar projects to address a well-known financing conundrum where individual projects are too small to attract bank financing on a standalone basis, but are too big to be consistently funded on balance sheet.

The benefits of these portfolio structures are significant. The "portfolio effect", that is, mainly not needing to rely on a single revenue stream, should result in greater lender flexibility on project agreement terms, which is critical when dealing with hyperscalers that can be inflexible in their insistence on shorter offtake terms and reluctant to commit to contracts beyond their standard duration policies. In the context of a financing for a standalone data centre project, the termination of the sole offtake agreement would likely put the financing at risk. By contrast, under a portfolio or accordion financing structure, the termination of one offtake agreement should not risk the financing of the rest of the data centres, provided that the debt service ratios, which are applied on a portfolio basis, are maintained.

Other features

The market for data centre financing continues to evolve in other respects as well. Distribution lock-up ratios (covenants that restrict the borrower from paying out distributions unless certain performance thresholds are met) in data centre financings are an evolving area given the relative infancy of non-recourse and limited recourse financing for data centres. However, the market has seen a combination of interest cover ratios and leverage ratios, which are distinct from the typical debt service cover ratio that is used for distribution lock-ups in energy or infrastructure project finance or the loan-to-value ratio that is common in real estate.

Project sponsors conduct parallel corporate bridge loan and project finance debt negotiations in order to ensure that development work can progress unhindered by the structure of the finance arrangements.

KEY RISKS FOR INVESTORS

Private equity firms participate in data centre transactions in several capacities (*see box "Typical parties to a data centre investment"*). Most commonly, private equity funds act as equity sponsors, providing the risk capital that sits at the base of the capital structure beneath any senior or mezzanine debt. In this role, private equity firms are increasingly not only acting as a capital source but also as strategic owners, taking the lead in transactions and often bringing operational expertise across the asset lifecycle. Private equity sponsors also frequently structure transactions alongside other institutional investors. Sovereign wealth funds, pension funds and infrastructure-focused investors are common co-investors, particularly in larger platform transactions where capital requirements may exceed the capacity or risk appetite of a single fund.

In some cases, private equity firms have also expanded into credit strategies, providing debt financing through dedicated infrastructure debt funds or direct lending vehicles, allowing them to participate across the capital stack. This allows private equity firms to adjust their position in the capital structure depending on the risk-return characteristics of a given transaction, whether through senior secured lending or subordinated debt.

For the parties seeking to invest in infrastructure assets, there are several key areas of risk.

Foreign investment screening and national security

Data centres are increasingly classified as critical national infrastructure, making them a focal point for foreign direct investment (FDI) screening regimes worldwide (*see feature article "Foreign direct investment: national screening regimes proliferate", www.practicallaw.com/w-040-3591*).

In the US, the Committee on Foreign Investment in the United States (CFIUS) has expanded its scrutiny of transactions involving data centres, particularly those located near sensitive government or military installations, and can impose conditions on, or block, acquisitions by foreign-owned sponsors. The National Security and Investment Act 2021 in the UK designates data infrastructure as a mandatory notification sector, requiring private equity funds to file with the Investment Security Unit before completing acquisitions of entities that operate UK data centres, even for minority stakes above 25% (*see feature article "National Security and Investment Act 2021: taming the M&A dragon", www.practicallaw.com/w-032-2847*). EU member states have progressively broadened their screening mechanisms under the FDI Screening Regulation (*2019/452/EU*) framework, with several jurisdictions, including Germany, France and Italy, treating data centre operators as critical infrastructure targets that are subject to mandatory pre-closing approval.

Beyond ownership-level screening, investors must also navigate export control and supply chain security risks. US restrictions on advanced semiconductor and AI chip exports, particularly to China, can constrain a data centre's ability to source critical hardware, while the EU's Chips Act (*2023/1781/EU*) introduces its own supply chain resilience requirements.

For private equity sponsors, practical implications include extended deal timelines due to multijurisdictional filing obligations, the risk of remedies or undertakings being imposed as conditions to clearance, such as restrictions on access to data or physical site security requirements, and the potential for post-completion call-in powers in jurisdictions that operate retrospective review mechanisms.

Construction and development risk

Construction completion risk represents a particularly significant challenge for investors, given the unique characteristics of data centre projects. Lenders on data centre projects will be more prone to ask for enhanced support from a sponsor given the higher risks associated with a hyperscaler offtake agreement versus a traditional solar power project, for example.

A threshold question for data centre developers and their investors is the choice of procurement model for the construction phase. The two predominant delivery methods are:

- Engineering, procurement and construction (EPC) contracts.
- Design and build (D&B) contracts.

While both models consolidate design and construction responsibility under a single contractor, which provides the developer with a single point of accountability, they differ materially in the allocation of risk, the degree of owner involvement and the pricing structure.

Under an EPC contract, the contractor assumes full responsibility for the engineering, procurement of all materials and equipment, and construction of the facility, delivering the facility on a turnkey basis (that is, ready to operate) that meets the predetermined performance specifications. EPC contracts are typically structured on a lump-sum turnkey basis, which provides the developer with price certainty and transfers the vast majority of cost, schedule and performance risk to the contractor. The developer's involvement is minimal after the contract is awarded. It specifies the desired outcome, such as power capacity, uptime levels and cooling performance, and the EPC contractor bears responsibility for delivering a fully operational facility. This model is particularly attractive to financial investors and lenders as it provides a fixed-price, date-certain delivery commitment that is more straightforward to underwrite.

By contrast, a D&B contract delegates both design and construction to the contractor but typically preserves greater owner oversight over the design development process and may leave the procurement of certain specialist equipment, such as IT systems, cooling infrastructure or power distribution units, to the developer. D&B contracts tend to be more collaborative, with the developer retaining an active role in design review and approval, which can be advantageous where the developer has specific technical requirements or wishes to retain control over key technology choices. However, this greater involvement also means that the developer retains a larger share of the co-ordination and interface risk. D&B pricing structures are often more flexible, ranging from cost-plus to target-price with pain-sharing and gain-sharing mechanisms, which can accommodate design evolution but introduces greater cost variability for the developer.

The choice between D&B and EPC is driven by several factors that are specific to the data centre context. Where a developer is constructing a facility "shell" (that is, a powered and cooled building into which the hyperscaler or end user will install its own IT systems), a D&B contract may be the more natural fit as the contractor is not being asked to guarantee the performance of specialist equipment that it did not procure. Where, however, the developer requires a full turnkey solution, including specialist IT infrastructure that is capable of operating to defined performance specifications, an EPC contract provides a cleaner transfer of risk.

The speed at which the project can be delivered to market is a further critical consideration: while both models can achieve faster delivery than a traditional design-bid-build procurement structure, the more defined deliverables and reduced scope for design iteration under an EPC contract can give it a slight edge in accelerating project timelines.

The market continues to mature in this respect. While D&B arrangements have long been the dominant model for large-scale data centre construction, there is increasing momentum toward EPC structures as project sponsors and their financial backers look to transfer greater construction-phase exposure to the contractor. That said, the extraordinary capital requirements of AI-driven and hyperscale facilities, which often range in the billions of dollars, can exceed the risk capacity of even well-capitalised contractors that operate on a fixed-price basis.

In response, many such projects are now adopting target-cost frameworks with contractual mechanisms to share the upside and downside, occupying a middle ground between the traditional D&B and single contractor EPC models. In this target-cost model, the contractor is entitled to retain any savings but, by the same token, absorbs any cost overruns. A parallel shift is also underway across the wider construction industry, with parties moving away from rigid fixed-price models toward more collaborative, open-book approaches. This is driven in large part by persistent supply chain disruption and extended procurement timelines for critical components, such as transformers, gas turbines and advanced semiconductors.

Ultimately, no single template applies universally. The construction contract must be structured on a project-specific basis, and be tailored to the developer's risk tolerance, lender requirements, and the prevailing commercial conditions in the relevant market.

Counterparty and offtake risk

Hyperscaler relationships present their own set of challenges. Hyperscaler procurement often takes place through direct negotiation with various parties, such as contractors and technology vendors, and an established relationship is critical. Hyperscalers have a well-established contractual model and strong bargaining position. They are highly creditworthy and would be a sole or anchor offtaker in a data centre project financing. Each hyperscaler has its own house forms for offtake agreements and data centre service agreements and will typically accept nothing else.

Against this backdrop, lenders will generally seek the "back-to-back" principle to be reflected in the contractual documents. The "back-to-back" principle is a core project financing objective which dictates that risk ought to be passed through the entire suite of project development documents: essentially, the project company must contract on no worse terms with its contractors or subcontractors than it does with the offtaker. This ensures that the risk is allocated correctly from the project company to the relevant contractors and sub-contractors and is a key factor in determining whether a project is financially viable. This is more difficult to achieve when a data centre project is not procured by government.

Another issue is lender step-in rights. Governmental institutions are likely to be more amenable than corporate entities in providing direct agreements, which allow lenders to step into the project company's place under specified events of defaults, as they have experience of these arrangements from public-private partnership (PPP) projects. Corporates, especially hyperscalers, will struggle with the notion of contracting with a service provider's lenders, and especially the notion of a lender having step-in rights.

Further, corporate offtakers, especially hyperscalers, are less likely to agree to provide relief in the agreements for political force majeure events (including change in law), which raises the risk profile for sponsors.

Another key consideration is that hyperscalers will typically put the onus on the data centre operator to procure utilities, which is in contrast to typical PPP arrangements where the procurement authority more commonly retains responsibility for securing utility connections.

All of these issues require careful structuring and alignment with each sponsor's risk appetite. Engineering the suite of project development documents is a key focus.

ESG and sustainability pressures

Environmental, social and governance (ESG) compliance pressures are a significant risk factor. Hyperscaler sustainability strategies and sensitivities require careful attention. There are nuances in their approaches, such as whether they are seeking to achieve net-zero or negative carbon emissions, carbon-free operations, or carbon offsetting arrangements, including carbon additionality (that is, a project's emissions reductions or removals would not happen without revenue from selling carbon credits).

A June 2026 study by researchers from Harvard University, University of California, Los Angeles and other institutions estimated that 403 US hyperscale data centres were associated with approximately 37 million to 54 million metric tons of carbon dioxide emissions between May 2024 and April 2025 (*https://arxiv.org/pdf/2606.05420*). Under the study's central scenario, these facilities consumed around 1.8% of total US electricity, with approximately 54% of the electricity attributed to them generated from fossil-fuel sources. Their average carbon intensity was approximately 545 grams of carbon dioxide per kilowatt-hour, around 48% above the contemporaneous US national grid average. There are mounting pressures on large technology companies to make disclosures about their data centre carbon emissions, water and land use.

Financing considerations add another dimension to ESG compliance. If the data centre is intended to be off-grid while being exclusively powered by renewable energy, this will be a red flag from a financing perspective because of the potential risk of intermittent service threatening the offtake arrangements. Lenders will look at solutions to address renewable power intermittency, namely through back-up electricity solutions, whether battery storage or diesel generator system, and electricity transmission line redundancy.

Related information

This article is at practicallaw.com/w-050-8998

	Other links from uk.practicallaw.com/
Topics	
ESG and sustainability	topic/w-032-4726
Information technology	topic/5-103-2074
Private equity and venture capital	topic/0-103-1350
Project finance	topic/3-201-5282
Practice notes	
Anatomy of a ground-mounted solar power project	w-001-8826
Cloud services: overview	4-566-3285
Financing structures used in project finance transactions	w-013-5591
Managing cybersecurity risk and compliance	6-615-8326
NIS 2 Directive: overview	w-037-2098
Power projects: development phase	6-622-9625
Project finance: UK law overview	0-202-3293
Previous articles	
Data centres: the key issues for businesses (2025)	w-045-7240
Data-sharing frameworks: a state of play in the EU and the UK (2024)	w-043-3409
Preparing for large-scale disruption: risk and resilience (2024)	w-044-5697
Foreign direct investment: national screening regimes proliferate (2023)	w-040-3591
Cloud computing: trying to find the silver lining (2023)	w-039-4053
A sustainable future: will smart energy get us DER? (2022)	w-036-7558
National Security and Investment Act 2021: taming the M&A dragon (2021)	w-032-2847
AI and data protection: balancing tensions (2019)	w-020-9713

For subscription enquiries to Practical Law web materials please call +44 0345 600 9355

Regulatory and operational complexity

The regulatory landscape includes obtaining the relevant permits from governments and utilities in a timely fashion, FDI laws, anti-bribery and corruption issues, evolving cyber security laws (such as the Network and Information Security Directive 2 (*2022/2555/EU*) (NIS 2 Directive) and the Digital Operational Resilience Act (*2022/2554/EU*) (DORA) in the EU), evolving operational resilience laws and critical infrastructure requirements, changing sustainability laws, and requirements around data flow and data localisation that may require certain types of data to be kept within a jurisdiction. Geopolitical instability and sanctions risk add further complexity, especially, for example, the backdrop of US/Iran and US/China tensions.

WIDER INFRASTRUCTURE

While data centres have commanded an outsized share of recent investor attention, they represent one dimension of a far broader private equity infrastructure investment cycle. BCG's Infrastructure Strategy 2026 report confirmed that limited partners invested a total of $211 billion into the asset class in 2025, an increase of more than 60% over 2024 and 11% above the previous fundraising peak in 2022 (*www.bcg.com/publications/2026/year-of-increasing-scale-and-diversification*).

Capital is flowing across the full spectrum of infrastructure sectors. The energy transition remains a dominant theme, with global investment in energy transition infrastructure reaching $2.3 trillion in 2025, an 8% increase from 2024, and renewable energy alone attracted $690 billion (*https://about.bnef.com/insights/clean-energy/bloombergnef-finds-global-energy-transition-investment-reached-record-2-3-trillion-in-2025-up-8-from-2024/*). Battery storage transactions surged by over 60% in 2025, signalling the technology's maturation into an essential infrastructure component. Conventional thermal power M&A experienced a notable resurgence, headlined by Constellation Energy's acquisition of Calpine, and Vistra's acquisition of Cogentrix. Transport infrastructure, social infrastructure, and logistics assets are also attracting significant capital, with the convergence of private equity and traditional infrastructure funds accelerating around hybrid assets that straddle multiple sectors.

Private equity investors face very similar challenges to those covered above in the context of data centres, including the following principal issues.

Regulatory and permitting risk

Infrastructure projects are subject to multi-layered regulatory frameworks that span planning and environmental permitting, FDI screening regimes and sector-specific licensing requirements. Governments are increasingly tightening FDI screening, with critical infrastructure, including energy generation, transmission and transport, being subject to heightened scrutiny.

Investors must also navigate evolving cyber security and critical infrastructure regulation, including the NIS 2 Directive, DORA and analogous regimes in other jurisdictions.

Financing structures and capital deployment

Infrastructure financing is evolving to accommodate the scale and complexity of current investment requirements. The traditional project finance model, which is centred on non-recourse lending against long-term contracted cash flows, is being supplemented by innovative structures, such as portfolio and accordion financings, and hybrid instruments, that blend project finance and corporate lending features. Private credit has also emerged as an increasingly important capital source, with infrastructure continuing its multi-year expansion in this area.

ESG and sustainability obligations

The ESG regulatory landscape is in a state of significant flux. In the EU, the Environmental Omnibus proposals that are currently going through the legislative process seek to simplify and streamline sustainability reporting requirements (*see News brief "Streamlining ESG reporting: the EU's omnibus package", www.practicallaw.*

com/w-046-2734). In the US, ESG investing remains at the centre of a sustained legal and political debate, with executive orders, state-level anti-ESG legislation, and shifting proxy advisory standards creating a fragmented compliance environment.

Despite this regulatory uncertainty, institutional investors and lenders continue to require robust ESG frameworks, and infrastructure assets, particularly in the energy and transport sectors, are subject to increasing scrutiny on emissions, environmental impact and social licence to operate. Climate-related litigation remains a material risk and enforcement actions targeting greenwashing are intensifying in many jurisdictions (*see Briefing "Greenwashing: engaging with regulators on ESG concerns", www.practicallaw.com/w-036-3695*).

Contractual risk allocation

Infrastructure projects demand sophisticated contractual frameworks that appropriately allocate construction, operational and market risk among developers, contractors, offtakers and lenders. The choice of delivery model, whether EPC, D&B, EPCM (that is, an engineering, procurement and construction management model, where the contractor only manages and advises rather than taking responsibility), or emerging hybrid structures with pain-sharing and gain-sharing mechanisms, has significant implications for financial viability and investor returns.

Supply chain volatility, equipment lead times (particularly for transformers, turbines, and specialised components) and workforce scarcity are reducing contractors' willingness to accept traditional fixed-price risk. This is driving a market-wide shift towards more collaborative cost-sharing arrangements. Lenders continue to insist on robust step-in rights and back-to-back risk allocation, which creates structural tension with commercial counterparties.

Cross-border investment considerations

Infrastructure is inherently jurisdictional and cross-border investment introduces additional complexity. Investors must engage with divergent regulatory regimes, political risk (including expropriation and sanctions exposure), currency risk and the challenge of structuring investments to optimise tax efficiency across multiple jurisdictions. Geopolitical fragmentation, including trade tensions, sanctions regimes and the reorientation of global supply chains, is reshaping the cross-border investment landscape. The convergence of private equity and sovereign wealth fund capital in mega-infrastructure transactions adds a further dimension. Foreign ownership sensitivities are also becoming an increasingly prominent feature of the regulatory review of these projects.

SUBSTANTIAL OPPORTUNITY

The scale of the opportunity is substantial: private capital will be essential to meet the estimated $106 trillion that is needed for global infrastructure investment in the period up to 2040. This ensures that private equity will remain a central driver of infrastructure development. However, the legal and regulatory environment is becoming more complex, not less, and the ability to navigate these challenges will increasingly differentiate successful investors from the rest of the market.

Successfully navigating these risks requires sophisticated legal, regulatory and operational analysis throughout the investment lifecycle, from due diligence, acquisition and financing through operation and exit.

Agnieszka Klich and Mark Thompson are partners, Lennard Ong and Luke Ong are associates, and Tosin Oredola is a trainee solicitor, at Covington & Burling LLP.