


CYIENT

A GUIDE TO POWERING THE AI ERA

**Building Resilient, Scalable Energy
Infrastructure for Hyperscale Data Centres**



CONTENTS

Executive Summary	03
-------------------	----

The Scale of the Opportunity	03
------------------------------	----

Grid Capacity and Availability	04
--------------------------------	----

Long Lead Times on Critical Equipment	05
---------------------------------------	----

Meeting Tier IV Requirements Without Fossil Fuels	06
---	----

The Cost of Building at Scale	08
-------------------------------	----

The Workforce Challenge	09
-------------------------	----

Sustainability in the Age of Hyperscale	10
---	----

Facilitating the Transformation with Cyient	11
---	----

Conclusion	12
------------	----

About the Author	12
------------------	----

About Cyient	13
--------------	----

Executive Summary

The global data centre market is growing at a pace few industries have seen before. Capacity is expected to nearly double to 200 GW by 2030, driven primarily by the surge in demand for AI and generative AI workloads. Projections from JLL point to a 14% compound annual growth rate, with around \$1.2 trillion USD in real estate value creation and a further \$1–2 trillion USD in technology infrastructure spending across data centre fit-outs.

But the infrastructure required to power, build and operate these facilities has become as complex as the technology they support. Grid connections can take a decade. Equipment lead-times stretch to seven years. A single 100 MW facility can cost more than \$5 billion. Sustainability requirements are tightening while uptime standards remain absolute.

This white paper examines the six challenges shaping data centre electrification today: grid capacity constraints, equipment lead-times, fossil fuel dependency, capital cost, workforce shortages and sustainability pressure. It outlines how developers and operators are responding, and how Cyient's engineering expertise supports clients navigating this landscape.

The Scale of the Opportunity

In Europe, the market was valued at USD 47.23 billion in 2024, with some projections placing it at USD 97.30 billion by 2030, a CAGR of 12.80%. The FLAP-D markets, Frankfurt, London, Amsterdam, Paris and Dublin, continue to lead the sector, while emerging markets such as Spain and Italy are expanding, driven by space availability and lower operating costs. The Nordic countries remain attractive for investment, with cool climates, established renewable energy infrastructure and supportive government incentives all contributing to their appeal.


Europe's ambitions are significant, but the growth expected in the United States remains in a different category. US power demand from data centres is projected to grow by 20–25% annually through to 2030.



Grid Capacity and Availability

Many regions lack the grid infrastructure to support large or rapidly expanding data centres, with connection timescales now measured in years rather than months. The table below summarises current conditions across key markets.

The table below summarises current grid connection conditions across key markets.



Marsden Hanna, Google's Global Head of Sustainability and Climate Policy, has publicly described transmission barriers as the biggest challenge to expanding data centre infrastructure, noting that one utility quoted a twelve-year timeline simply to study the interconnection process. Approvals are increasingly conditioned on infrastructure that must be designed, financed and built before a single server goes live.

In response, data centres are turning in growing numbers to microgrids and hybrid power systems. In practice, this means developers building and operating their own power stations, transmission infrastructure and battery storage, bypassing grid connection delays at the cost of significantly higher capital expenditure and project complexity.






Long Lead Times on Critical Equipment

Equipment lead times across the critical categories that data centre electrification depends on have extended significantly, with gas turbines now reaching seven years. The table below outlines current conditions across the three primary equipment categories.

Equipment Type: Gas Turbines Typical Lead Time: 5–7 years Key Context: Orders hit a twenty-year high in 2024 with 14 GW booked. OEMs are at full capacity across the board.	Equipment Type: Reciprocating Engines Typical Lead Time: 2–4 years Key Context: Drawn to data centres for fuel efficiency, fast start-up and fuel flexibility. Shorter than turbines but still a significant schedule risk at scale.	Equipment Type: Transformers (above 5 MVA) Typical Lead Time: 2–4 years Key Context: Many OEMs at full capacity. One major OEM is investing €220 million in a new facility to create around 350 jobs, but it will not come online until 2028.
---	--	---

Siemens Energy reported orders nearly 20% higher than the prior year on a comparable basis, with revenue rising 15.2% to €39.1 billion and a backlog of €136 billion at the time of reporting. The causes are not limited to demand alone. Limited manufacturing capacity and a shortage of skilled labour are compounding the problem, and these extraordinary lead times effectively remove many high-quality OEMs from consideration on large-scale projects.

With energy supply now on the critical path of most major projects, OEM engagement at the planning stage is essential. As the industry shifts from pure data centre construction to data centre and energy production, some developers have been slow to involve OEMs early, relying on generalisations for permitting. In countries with stringent planning legislation for large, energy-intensive developments, this creates significant problems when design or layout changes are needed later in the process. Options to mitigate long lead times include phased construction and modular approaches. Schedule overlap between design, procurement and construction is being used with increasing frequency, though this carries risk and requires disciplined programme management.



Meeting Tier IV Requirements Without Fossil Fuels

Data centre developers now carry significant pressure: the need to reduce reliance on fossil fuels and demonstrate credible sustainability credentials, while meeting some of the most demanding availability standards of any sector. The industry standard for Tier IV requires near-continuous operation at 99.995% uptime, equivalent to a little over 26 minutes of downtime per year. Some operator requirements go further still, specifying 99.999% availability where only seconds of unplanned downtime per year are acceptable before penalties apply.

Due to the intermittent nature of wind and solar, these sources alone cannot meet that standard. In the past, smaller data centres with straightforward grid connections relied on diesel generators as backup. At hyperscale, stricter emissions legislation has pushed the industry away from diesel. The table below outlines the fuel options currently available and the trade-offs each presents.

FUEL SOURCE	CHARACTERISTICS AND CURRENT STATUS
Natural Gas	Role: Now the primary replacement for diesel at hyperscale. Gas turbines are the standard, though modern reciprocating engines offer greater fuel flexibility, capable of running on gaseous, liquid or sustainable fuel sources. Status: Widely adopted. Considered cleaner than diesel but not zero-carbon. Stricter legislation is beginning to constrain its long-term role in some markets.
HVO (Hydrotreated Vegetable Oil)	Role: A drop-in replacement for diesel in reciprocating engines, compatible with existing fuel handling and storage infrastructure. Emissions: A reduction of up to 90% in CO2 compared with diesel, alongside lower NOx and carbon monoxide emissions. Recent projects in Ireland have been designed to include HVO specifically for this reason. Supply challenge: Waste oils such as used cooking oil are a finite resource and lack scalability beyond certain thresholds. Greater reliance on crop-based oils grown specifically for HVO production would reduce arable land available for food and could accelerate deforestation.
Wind and Solar + Battery Storage	Role: A viable generation component, but not a standalone solution for data centres with Tier IV requirements. Limitation: The inconsistent nature of wind and solar means they cannot power a data centre alone. Battery energy storage systems are developing rapidly, but operating a hyperscale facility from a BESS for any significant duration remains a major engineering and commercial challenge. Status: Increasingly paired with gas or other dispatchable sources to provide a credible sustainability layer.
Hydrogen	Role: Hydrogen is emerging as a potential energy solution for data centers, mainly to reduce carbon emissions and not as a primary power source Limitation: Hydrogen is not yet widely used in data centre power supply due to high production costs, limited transmission & storage networks, efficiency losses when compared to direct electricity supply and also for safety and handling considerations. Status: Hyperscalers have tested hydrogen fuels cells as alternatives to back up diesel generators and a <u>1MW data centre in the US</u> is currently operating on 100% hydrogen power.
Small Modular Reactors (Nuclear)	Role: A stable, weather-independent clean power source capable of delivering constant, large-scale baseload energy for data centres. Limitation: Most SMRs are not expected to be operational until the 2030s. Regulatory frameworks in many locations either do not exist or are limited to large-scale nuclear facilities, and public perception of nuclear energy adds further permitting complexity. High initial capex, limited availability of trained construction and operations staff, facility security requirements, and unresolved long-term spent fuel storage all remain significant barriers. Status: A promising longer-term solution, particularly as AI-driven demand accelerates, but not yet a near-term option for data centre power supply.

Fossil fuels will remain a significant part of data centre electrification in the short to medium term. The engineering challenge is to reduce that dependence progressively while maintaining the availability levels that operators and customers require.

The Cost of Building at Scale

The scale of investment required to build and operate a modern hyperscale data centre is redefining what large capital expenditure looks like. Some 100 MW facilities are now costing in the range of \$3.4 to \$5.5 billion. Hyperscalers including Amazon, Microsoft, Google, Meta and Oracle are projected to spend approximately \$300 billion in combined capital expenditure through 2025 and 2026. With market capitalisations running into hundreds of billions and trillions of dollars, cost overruns at individual facilities are manageable for companies of this scale, and self-financing remains a realistic option.

For the rest of the industry, the numbers are harder to absorb. According to Allianz Commercial construction experts, average-sized facilities now cost between \$500 million and \$2 billion. These figures put data centre development well beyond the reach of most individual companies acting alone. While traditional debt financing from large financial institutions remains common, funding from infrastructure funds, REITs, venture capital and sovereign wealth funds is playing an increasingly significant role. Analysts at Morgan Stanley and JPMorgan have estimated \$1.5 trillion in additional borrowing by technology companies in the near term for data centre expansion.


That level of debt is drawing scrutiny from financial analysts. Credit default swaps on at least one major hyperscaler have widened to multi-year highs, with some analysts advising clients to seek protection.

Tranched Financing and Utility Joint Ventures

Challenges with traditional funding and electrification timelines are pushing operators toward alternative development approaches. Joint ventures with utility companies are one emerging route. In Mainz, roughly 30km from Frankfurt, Green Mountain has partnered with German utility KMW to build a 54MW data centre in an industrial area adjacent to KMW's existing power plants. Proximity to the River Rhine enables water-based cooling, while waste heat feeds directly into the city's district heating network.

Financing structures are also evolving. Upfront capex can make lenders reluctant to commit full funding at once. Phased financing, where capital is released against construction milestones on fast-track schedules, is increasingly common. Vantage Data Centers is among those reported to have adopted this model. Secure tranches (Class A) are often paired with higher-yield, riskier tranches (Class B) to attract different investor classes. Utility companies are also entering as financial participants, with joint venture structures emerging across multiple markets.

Payback periods of ten to twenty years make it impossible to forecast energy prices or interest rates with confidence. Rapid technological change creates a genuine risk that facilities could face obsolescence well before being fully paid for. Global construction volume also raises the possibility of regional market saturation, with new facilities struggling to fill capacity at projected rates.



The Workforce Challenge

The pace of expansion is creating a serious skills shortage across the industry. The roles in highest demand, electrical and mechanical engineers, control system specialists and building management system (BMS) engineers, require hands-on experience that takes years to build. Tight project schedules leave little room for training candidates with limited backgrounds, so companies end up competing for a small and finite pool of qualified professionals. Data centre construction also competes directly with the technology and semiconductor sectors, which often offer higher salaries and carry greater brand appeal for the candidates both are trying to recruit.

The problem extends beyond recruitment. Rapid technological change means existing staff need continuous upskilling, but fast-moving construction schedules leave little room for structured training. Developers are frequently locating data centres in more remote areas outside cities and established industrial zones, where land is available and costs are lower. That geography makes it harder to find local staff and increases the challenge of retention, as workers seek opportunities closer to home. Near-zero downtime tolerance creates a work environment that is inherently high-pressure, a factor that compounds turnover risk further.

The mix of permanent employees and subcontractors that characterises most large-scale projects introduces additional friction. Integrating teams with different systems, processes and working practices is a recurring challenge. When construction peaks pass, contract staff leave, taking significant specialist knowledge with them. Project budgets rarely support retention without a clear pipeline of follow-on work, and the knowledge loss at each transition is a real and underestimated cost.

One structural response to the staffing challenge is the use of specialist engineering services partners. Working with a dedicated provider allows organisations to access experienced teams on a flexible basis, maintain continuity of knowledge across project phases and retain access to skilled professionals between projects, without carrying the fixed cost of permanent headcount. This model is becoming increasingly relevant as projects grow in complexity and the gap between available skills and required expertise continues to widen.



Sustainability in the Age of Hyperscale

Large AI-focused hyperscale data centres operating at 100 MW and above can require roughly the equivalent power of 100,000 homes. That scale of energy consumption creates real public perception challenges and is a recurring issue in planning and permitting processes across multiple markets.

Water consumption adds another dimension. Cooling requirements can consume substantial volumes, with some estimates suggesting up to 25.5 million litres of water annually per MW of capacity. For facilities in water-stressed regions, this is a material planning and operating consideration.

Geography plays a significant role in managing both issues. Cool climates, such as those in the Nordic countries and parts of Canada, reduce the volume of active cooling required and lower the overall environmental footprint of a facility. Warmer, sunnier regions present a different balance, with stronger solar generation potential but higher cooling demands.

Waste Heat Recovery

Waste heat is an underutilised opportunity in many markets. Data centres generate substantial heat that can be fed directly into district heating networks. This model is already operating in the Nordic countries, France and Germany, and works best where district heating infrastructure already exists. Where it is not available, waste heat can serve adjacent buildings such as offices, hospitals and university campuses, industrial facilities in the vicinity, or public amenities including swimming pools.

Tracking Sustainability Performance

Measuring and reporting sustainability performance across facilities requires a consistent set of metrics. The three most widely used in the industry are:

Power Usage Effectiveness (PUE)

Total facility energy divided by IT equipment energy. The closer to 1.0, the more efficient the facility.



Carbon Usage Effectiveness (CUE)

Total CO2 equivalent emissions from data centre energy consumption divided by IT equipment energy. Essential for Scope 1 and Scope 2 emissions reporting.



Water Usage Effectiveness (WUE)

Annual site water usage divided by IT equipment energy. Increasingly scrutinised in water-stressed markets and by ESG investors.



Facilitating the Transformation with Cyient

The challenges covered in this paper are interconnected and, at their core, engineering problems. Addressing them requires domain experience applied consistently across project phases, not just technical capability in isolation. Cyient is a global lifecycle engineering company, partnering with over 300 customers, including 40% of the top 100 global innovators. We power mission-critical industries across the full engineering value chain, combining domain knowledge, human expertise, and AI to deliver Intelligent Engineering and Technology Solutions.

We have a track record spanning more than 1,000 engineered power plants, including data centre power infrastructure projects in Ireland and Northern Europe. Working with a specialist engineering partner also directly addresses the workforce continuity challenge described in Section 7. Clients can draw on experienced teams flexibly across the project lifecycle, scale capacity to project demand, and retain specialist knowledge between phases, without carrying permanent headcount across every discipline. Our capabilities include:



Power systems and microgrid engineering

End-to-end electrical and power infrastructure design, including on-site generation, microgrid architecture, hybrid power configurations, and battery storage integration.



OEM engagement and procurement support

Engineering specification development to support early OEM engagement, reducing the risk of design changes that cause schedule and cost disruption later.



Resilience and fuel system engineering

Design capability across natural gas, HVO, hydrogen-ready, and hybrid renewable configurations, supporting Tier IV compliance with reduced fossil fuel dependency.



Digital engineering and BIM coordination

Multidisciplinary design delivery incorporating BIM, clash detection, and coordinated design management, critical on fast-track programmes where design and construction overlap.



Asset management and digital twins

Predictive maintenance frameworks and digital twin models that support near-zero unplanned downtime requirements across operations and lifecycle management.



Sustainability performance monitoring

Measurement frameworks for PUE, CUE, and WUE, baseline assessment, and reporting support aligned to regulatory and investor requirements.



Programme and cost risk management

Disciplined engineering delivery across phased and fast-track programmes, reducing rework, supporting schedule overlap between design, procurement and construction, and helping clients manage capital risk on facilities where cost overruns run into the hundreds of millions.

Conclusion

Data centre electrification is no longer a future ambition. It is a present-day constraint on the pace of digital growth. The challenges outlined in this paper, power availability, grid capacity and timelines, cost volatility, staffing, sustainability requirements and rapid technological change, are deeply interconnected and increasingly difficult to address in isolation. As data centres scale to support AI, cloud and high-density computing, they are becoming as much energy infrastructure as digital infrastructure.

Overcoming these challenges will require closer coordination between data centre developers, utilities, regulators and OEMs, alongside long-term planning and significant capital investment. Ultimately, the ability to navigate them will determine not only how and where data centres are built, but how fast the digital economy itself can grow.

References

1. **JLL:** [2026 Global Data Center Outlook.](#)
2. **Network World / Reuters :** [Google warns transmission delays are the biggest threat to data centre expansion \(January 2026\).](#)
3. **Siemens Energy:** [Q3 FY2025 Earnings Release.](#)
4. **Allianz Commercial:** [The Data Center Construction Boom \(2025\).](#)

ABOUT THE AUTHOR



Jay Gauld

Delivery Project Manager – Energy, Cyient

Jay Gauld is a Delivery Project Manager at Cyient specializing in Tier IV data centers and microgrid solutions. He brings a strong track record in delivering complex energy projects on schedule, with expertise in managing multiple priorities and leading high-performing teams.



About Cyient

Cyient is a global lifecycle engineering company powering mission-critical industries from design to aftermarket, across products, plants, and networks. We partner with over 300 global customers across 30 countries, supported by more than 14,000 associates. Recognized among the Top 10 pure-play global engineering services providers, we reported \$658M in annual revenue for FY26. Across Aerospace, Rail, Automotive and Mobility, Connectivity, Utilities, Mining, Energy, Healthcare and Life Sciences, and Spatial Intelligence, we combine deep domain expertise, human insight, and AI to deliver intelligent engineering solutions.

For more information, please visit www.cyient.com



Contact Us

North America Headquarters

Cyient, Inc.
99 East River Drive
5th Floor
East Hartford, CT 06108
USA
T: +1 860 528 5430
F: +1 860 528 5873

Europe, Middle East, and Africa Headquarters

Cyient Europe Limited
Apex, Forbury Road,
Reading
RG1 1AX
UK
T: +44 118 3043720

Asia Pacific Headquarters

Cyient Limited
L14 3 Parramatta Square,
153 Macquarie St,
Parramatta NSW 2150,
Australia
T: +61 2 8887 8600

Global Headquarters

Cyient Limited
Plot No. 11
Software Units Layout
Infocity, Madhapur
Hyderabad - 500081
India
T: +91 40 6764 1000
F: +91 40 2311 0352

Follow us on:  