standard chartered



Can AI keep its cool?

The choices shaping sustainable AI growth

sc.com

Contents

Foreword 3

Executive summary 4

AI turns data centres into a global infrastructure super cycle 9
60% of data centre capacity was added in the past 5 years 10
Asia emerging as a major data centre growth frontier 13
AI may trigger investments of up to USD9 trillion 16
Data centre-related funding needs until 2029 may reach USD1 trillion 19

Solving the data centre power challenge 22
AI puts a spotlight on power demand and the grid 23
The power challenge is about more than data centres 27
Solutions to further reduce data centre power consumption 28
Solution 1: Improve operational efficiency 29
Solution 2: Reuse waste heat to lower broader energy demand 31
Solution 3: Expand exposure to renewables further 32
Solution 4: Move to 800-volt DC distribution to reduce power loss 40
Solution 5: Develop battery storage and microgrid capacity 42
Integrating all solutions can turn data centres into flexible energy system participants 43
AI-enabled energy gains may outweigh its own power demand 44

Will low water data centres become the norm? 46
Direct data centre water use: perhaps less than expected 51
Data centre cooling developments 53
Data centre water efficiency: hyperscalers set the tone 55
Solutions to reduce data centre water use further 60
Solution 1: Maximise use of dry and immersion cooling 60
Solution 2: Allow data centres to run at higher temperatures 62
Solution 3: Integrate circularity solutions 64
Solution 4: Maximise the use of renewable power 66
Solution 5: Develop water replenishment programmes close to data centre locations 66
Solution 6: Regulation: a potential external driver? 71

What does good look like? 79
Summary 82

Appendix 83
Acknowledgement 83
About the author 83

References 84

Disclaimer 87

Foreword

Artificial intelligence is reshaping the global economy, and the data centres that enable it are becoming a defining infrastructure investment of this decade. Their rapid expansion brings significant opportunity, but also legitimate questions about electricity demand, water use, emissions and pressure on local infrastructure and communities. A credible response requires moving beyond headline claims to understand where these impacts arise, how they vary by location and operator, and which solutions can deliver the greatest improvement.

This report shows that data centre growth is neither environmentally benign nor inherently unsustainable. The outcome will depend on whether expansion is matched by efficient and well-utilised computing, low-carbon and low-water power, advanced cooling, circular design, heat reuse, stronger grids and active engagement with local stakeholders. It also highlights the importance of looking beyond portfolio averages: environmental performance is ultimately shaped at individual facilities, where local power systems, water availability and technology choices determine real-world impact.

The central message is constructive. Technology and operating practices are advancing quickly; many of the solutions needed to reduce power and water intensity already exist and compute power per unit of energy and water are following exponential Moore's Law-type efficiency curves with each successive generation. The challenge is to make "what good looks like" the standard across the sector: combining high efficiency with responsible location and resource choices, high-quality clean-energy procurement, circularity and transparent, comparable disclosure at data centre level.

We are grateful to Microsoft, NVIDIA and Trane Technologies for the insights, data and technical perspectives they contributed to this report. Their input helped deepen our understanding of the technologies, operational practices and infrastructure solutions shaping more sustainable data centre growth.

I hope this report supports a healthy, informed dialogue between operators, policymakers, financiers, customers and local communities and helps direct capital and policy towards the infrastructure capable of enabling AI-led growth within acceptable environmental and social constraints.

Marisa Drew
Chief Sustainability Officer



Executive summary

This report concludes that while AI data centres may create legitimate environmental pressures, the tools, technologies and operating practices available to address them are more advanced and scalable than recent headline concerns suggest. There is no doubt that the sector requires significant investment, substantially more electricity and stronger infrastructure which may create pressure where power and water systems are constrained. Yet the new class of data centres are highly efficient and the average data centre is already in the process of becoming so: power and water intensity are falling as operators optimise cooling, airflow, power chains and facility design and develop deep community engagement. The sector has the potential to outperform long-term environmental impact estimates by adopting efficiency enhancing and low-water solutions, clean energy power sourcing, integrating circular strategies and through the development of more transparent and comparable reporting frameworks.

AI is creating a global infrastructure super cycle

Artificial intelligence is transforming data centres from specialist digital infrastructure into a global investment super cycle. Around 60% of today's live capacity was added in the past five years, yet industry forecasts suggest the world may add roughly the same amount again by 2030, with Asia and other emerging markets becoming increasingly important growth hubs.

The scale of required investment is significant. Some estimates suggest that enabling the AI-driven build-out may require up to USD9 trillion across IT hardware, physical infrastructure, power generation, grid capacity, cooling systems and enabling supply chains. IT hardware remains the most immediate CAPEX driver, but growth is increasingly constrained by land, construction costs, specialised equipment, grid connections and reliable power access. Long lead times for key equipment mean infrastructure delivery may determine how quickly new capacity can be brought online.

Our analysis suggests that data centre-related external funding requirements could approach USD1 trillion by 2029, creating a substantial need for both public and private external capital. Investor concerns over rising sovereign debt levels and increased borrowing by data centre operators has caused yields to rise this year. Such an environment may benefit those data centre operators that focus on strengthening balance sheets and cash flows and that adopt transparent disclosure policies as outlined in this report. Sustainable finance is becoming a more important source of capital with Bloomberg estimating that data centre developers have issued approximately USD186 billion of green, sustainability-linked and other sustainable debt instruments since the launch of ChatGPT. The development of sustainability benchmarks for data centres may help unlock additional pools of sustainable capital, making sustainable finance an increasingly important funding source for future AI and data centre expansion.

Power is a constraint: efficiency and procurement can moderate it

Despite its rapid growth, the data centre sector makes up only a modest share of global electricity use when compared with traditional energy intensive sectors. However, without strong efficiency gains data centre electricity demand could double by 2030 implying that spatial concentration is a risk as gigawatt-scale facilities can create acute pressure on local generation, transmission, distribution and connection queues. Power availability is consequently becoming a strategic differentiator and may redirect investment towards markets able to expand energy infrastructure more quickly.

The power intensity of data centres is improving owing to the adoption of efficiency enhancing solutions. Better model architecture, more efficient chips, higher server utilisation, advanced thermal management, heat reuse and lower-loss power distribution can all help to further reduce the energy needed per unit of useful compute. Hyperscalers appear structurally advantaged through their scale, purpose-built facilities, proprietary hardware, higher utilisation and access to capital.



Energy procurement and diversification is another strategic option available to data centre operators to reduce power challenges. While some operators are expanding gas generating capacity, it is likely that the rapid expansion into renewable energy and energy storage capacity will continue at pace although there are reports that hyperscalers may face difficulties meeting their ambitious targets (Bloomberg, 2026). Nevertheless, hyperscalers currently lead from the front as Amazon, Google, Meta and Microsoft have together contracted more than 140GW of renewable power, including almost 35GW in 2025 alone. The quality of procurement matters: physical delivery, additionality, geography and hourly matching provide a more credible picture than a single annual renewable percentage.

AI's broader impact on the global energy system may be partly counterbalanced by AI-enabled efficiency gains across key sectors. While AI increases data centre power demand, if deployed effectively, AI-enabled optimisation also has the potential to help reduce energy consumption in larger end-use sectors such as buildings, industry, transport and power systems. This could offset a meaningful share of the electricity required to run AI infrastructure, although the net effect will depend on adoption rates and rebound effects.

Water intensity challenges can be mitigated, including through AI-enabled optimisation

The water debate is also more nuanced than headlines suggest. Although some facilities can exacerbate local water stress, particularly in arid regions, water use depends heavily on technology choices, climate and electricity sourcing. This distinction matters because, despite legitimate local concerns, data centres generally exhibit lower water intensity than many other industries. AI may help accelerate improvements because high-density workloads are increasingly difficult to support with conventional air and evaporative cooling. Direct-to-chip liquid cooling, immersion cooling, dry cooling, closed-loop systems and allowing data centre coolant to be run at higher temperatures can materially reduce cooling energy and, in suitable configurations, bring direct water consumption close to zero.

Using recycled and non-potable water, heat recovery and careful site selection can further reduce local impacts. The experience of some of the hyperscalers suggests that the development of water replenishment programmes in areas where data centres are operated have the potential to minimise their impact on freshwater availability. These solutions can also improve rack density and total cost of ownership, aligning sustainability and investment cases. Against common perception, low-water data centres could become increasingly common.

What good looks like

"Good" in relation to data centres must be judged as an integrated, facility-level outcome rather than through a single metric on a portfolio level. Power Use Effectiveness (PUE) is useful for measuring non-IT overhead and best-in-class operators now achieve a PUE of around 1.1. However, it does not capture IT productivity, carbon intensity, water stress, heat recovery or circularity. Water Use Effectiveness (WUE) is increasingly used as a measure for data centre water efficiency with sector-leading operators achieving a ratio of less than 0.25. However, WUE is location and climate dependent and maybe misleading when used in isolation. Leading facilities optimise energy, carbon and water together, taking account of local grid conditions, climate and water availability, while reporting both absolute impacts and intensity metrics.

Priority actions for data centre operators

- **Use less before procuring more:** improve model and software efficiency, deploy more efficient hardware, raise server utilisation and optimise workload scheduling where possible so that each unit of electricity and water produces more useful compute.
- **Design for low-resource operation:** use liquid, dry or closed-loop cooling where technically and climatically appropriate; allow facilities to operate at higher temperatures if possible; reduce dependency on potable water; and embed heat reuse and circular design from the outset.
- **Secure credible clean power:** prioritise additional, physically deliverable low-carbon electricity, improve geographic and hourly matching, and combine procurement with energy storage, microgrids and flexible demand to reduce pressure on constrained grids.
- **Make location a sustainability decision:** screen sites for grid readiness, carbon intensity, water stress, heat-reuse and water replenishment potential, water and power infrastructure resilience and community impact before committing capital.
- **Build local legitimacy:** engage regulators, utilities and communities early, set site-specific resource commitments and demonstrate how projects contribute to local infrastructure, water resilience and economic value.

Priority actions for policymakers and financiers

- Policymakers should link permitting and grid connections to credible power, water and community-impact plans; establish common facility-level measurement standards; and reward projects that add clean generation, flexibility, heat reuse and local water resilience projects.
- Financiers should move beyond data centre operators' portfolio averages and test each project against location-specific resource constraints, technology choices and transition plans, using performance-linked capital and covenants to accelerate measurable improvements.

Facility-level disclosure is the foundation for credible growth

This report shows that current operator disclosures are often insufficient for policymakers, financiers and local communities to assess the impact of individual data centres. Portfolio averages can conceal inefficient facilities, local water stress, constrained grid nodes and trade-offs between electricity and water use. Comprehensive disclosure that is standardised, independently assured and provided at data centre level wherever security and customer confidentiality permits, with absolute values and intensity metrics would help address this. Better data would also make facilities comparable, improve planning and approval decisions, direct capital towards stronger operators and allow environmental performance to be measured and financed with greater confidence.

Bottom line: AI-driven data centre growth is neither environmentally benign nor inherently unsustainable. Its outcome will depend on whether expansion is paired with timely low-carbon power, efficient and well-utilised computing, low-water cooling, circular design, local resource stewardship and facility-level transparency. If these conditions become standard practice, the sector has the potential to expand while reducing power and water intensity and may also enable wider energy savings across buildings, industry, transport and power systems. If they do not, local infrastructure, environmental and social constraints rather than demand for AI may become the binding limit on growth.

Key efficiency solutions for data centres

This report highlights various solutions or strategies that data centre companies can adopt to accelerate the improvement of their environmental footprint. Figure 1 provides a summary.

Figure 1

Summary of key solutions available to data centres that help improve their efficiency

Solution	How it reduces environmental footprint	Potential power / energy reduction	Potential water reduction
Optimise cooling technology			
Air-based / free-air cooling	Uses favourable ambient conditions and economisation to reduce compressor/chiller operation; dry air cooling can eliminate routine cooling-water consumption.	Strong reduction in cooling energy needs	High reduction in cooling water use
Hybrid cooling systems	Switches between dry and wet/adiabatic modes, using water only when temperature or load requires it.	Cooling related power demand can drop substantially	High reduction in cooling water use
Immersion / liquid cooling	Moves heat directly into liquid, reducing server-fan and room-cooling loads and enabling high-density compute with little evaporative cooling.	Meaningful reduction in total facility electricity	Almost full reduction in cooling water needs possible
Closed-loop / zero-water cooling	Recirculates coolant without evaporation, removing continuous cooling-water make-up while retaining heat-rejection equipment.	Cooling electricity demand can be reduced substantially	High reduction in cooling water use
Incorporate low carbon power solutions			
Renewable power	Replaces fossil/thermoelectric electricity with wind or solar; does not reduce IT demand but lowers emissions and the electricity supply's water footprint.	No reduction in data-centre electricity demand but sharp reduction in GHG emissions	Very high reduction in indirect generation water versus wet-cooled thermal supply; no direct onsite reduction
Nuclear	Provides firm low-carbon electricity but does not reduce data-centre power demand; cooling design determines the upstream water footprint.	Reduces transmission-related power loss through direct connection	SMRs with integrated hybrid systems cut cooling electricity and reduce indirect water consumption
Geothermal	Supplies firm, low-carbon electricity and can also provide direct geothermal cooling where suitable subsurface heat-exchange resources exist; it does not reduce IT demand itself.	No overall change in power demand but switch to cleaner power reduces GHG emissions	Allows for switch to non-potable water
Other (biofuels / fuel cells)	Biofuels such as HVO can replace fossil diesel in backup generators, while fuel cells can provide lower-emission onsite or backup power; neither reduces computing electricity demand.	No computing-electricity reduction; backup generation only; Reduced GHG profile	No standard reduction; fuel-cell water impact depends on fuel and design
Adopt circular solutions			
Recycled / reclaimed water	Substitutes treated municipal wastewater for potable cooling make-up; total cooling water may remain similar.	Treatment/pumping need may increase electricity use somewhat	Very high reduction in potable-water demand possible
Increase cycles of concentration	Keeps cooling-tower water in circulation longer, reducing blowdown and replacement make-up water.	No change in overall facility electricity use	Cooling tower and blowdown water needs can reduce substantially
Use rainwater, greywater or effluent	Captures local non-potable supplies for cooling make-up, irrigation or building uses, displacing municipal drinking water.	Small electricity increase from collection, treatment and pumping is possible	Very high reduction in potable-water demand possible

Solution	How it reduces environmental footprint	Potential power / energy reduction	Potential water reduction
Adopt circular solutions			
Saltwater / alternative thermodynamic systems	Uses seawater, brackish water or naturally cold water as a heat sink, reducing freshwater demand and mechanical chilling.	Substantial reduction in cooling electricity is possible	Very high reduction in freshwater withdrawal
Advanced water treatment	Uses filtration, reverse osmosis or blowdown recovery to enable more reuse and higher cycles of concentration.	Small electricity increase from collection, treatment and pumping is possible	Waste water recovery reduces freshwater demand
Waste heat recovery	Exports server heat to district heating, buildings or industrial users, avoiding separate heating fuel; may require heat pumps.	Can reduce power consumption	Benefits mainly through reduced indirect water consumption
Re-use and refurbishment of equipment	Repairs, redeploys and harvests parts from servers and racks, avoiding new manufacturing and e-waste.	Benefit mainly through reduced indirect power	Reduced manufacturing delivers indirect water benefits
Use circular construction materials	Uses lower-carbon/recycled concrete, steel, metals and plastics and designs structures/ components for recovery.	No facility energy impact but reduced GHG emissions through embodied-carbon reuse	Reduced need for virgin manufacturing reduces indirect water consumptions
HVAC condensate	Collects water condensed by air-conditioning or air-handling equipment and reuses it for cooling-tower make-up, humidification, cleaning or irrigation.	Spraying reused condensate can lower temperatures and reduce cooling power needs somewhat	Reusing captured air-conditioning condensate reduces potable-water need
Mine water, quarry lakes or deep surface-water reservoirs	Uses stable-temperature mine pools, quarry lakes or deep reservoirs as closed-loop heat sinks, reducing reliance on chillers and evaporative cooling.	Can cut facility cooling electricity by 30-90% depending on baseline and climate	Potentially very high reduction in potable cooling water with a closed-loop heat exchanger
Optimise design and engagement strategies			
Location optimisation	Sites facilities in cooler climates, low-water-stress basins and areas with lower-water electricity and access to free cooling.	Location and facility design optimisation can reduce power consumption substantially	Cooling and location optimisation can meaningfully reduce direct water footprint
Run DC at higher temperatures	Raises server inlet/set-point temperatures, expanding economiser hours and delaying evaporative or mechanical cooling.	Project-specific; can remove most mechanical chilling in favourable climates, but no robust comparable total-facility percentage is established	High reduction in facility cooling-water use with validated closed-loop dry-cooler designs
Optimise DC-management using AI	Uses real-time controls and forecasting to optimise pumps, fans, chillers, temperature, humidity and workload placement.	AI-driven cooling control software can enable significant energy savings	Optimised cooling can reduce direct water needs while reduced power needs also lower indirect water footprint
Invest in local water infrastructure	Funds leakage reduction, recycled-water pipelines, treatment capacity or storage that improves shared basin supply.	No direct facility impact	Improves local water security and availability through reduced leakage
Develop water replenishment projects	Restores watersheds, aquifers or community supply to balance or exceed operational water consumption.	No direct facility impact	No direct-use reduction; potentially 100%+ basin replenishment against a defined accounting volume
Switch from 400V AC to 800V DC power	This reduces the number of power conversion steps which reduces heat, power loss, cooling demand and copper requirements all with a power and water intensity.	Meaningful improvement in end-to-end power efficiency	No direct water impact; indirect benefit from improved efficiency and reduced copper requirement

Source: Standard Chartered Bank



AI turns data centres into a global infrastructure super cycle

Data centres have moved from a niche digital infrastructure asset class to a fast-accelerating, AI-led global investment cycle. Around 60% of today's live capacity was added in just the past five years, yet industry forecasts suggest the world may need to add roughly the same amount again by 2030. Growth is rapidly broadening geographically across Asia and other emerging markets, where countries such as India and Malaysia could become materially important capacity hubs. This expansion is creating a multi-trillion-dollar CAPEX opportunity across IT hardware, physical infrastructure, power generation, grids and cooling but also exposes potential bottlenecks: construction costs vary sharply by location, inflationary pressure is building, and critical equipment such as transformers, generators and switchgear can take more than a year to procure.

60% of data centre capacity was added in the past 5 years

Strong growth in digital and internet services has been a key driver for the development of data centres (DC) as demand for data storage and processing accelerated (Shehabi & et al, 2024). Some of the current factors increasing the need for data centre capacity include the development of the Internet-of-Things (IoT), cloud-based computing, the development of Artificial Intelligence (AI) and a growing desire for onshore data storage capabilities to help reduce data security concerns (Figure 2).



Figure 2

Key drivers of data centre demand

IoT	Rapid growth in the number of internet-based connected devices, systems and services drives incremental data centre demand.
Cloud	Corporates are increasingly storing data in the cloud rather than on onsite, physical, servers. Data centre demand further increased by growing cyber security needs and digitisation of economic services.
Connectivity	Expansion of internet and mobile connectivity across developing countries creates demand for more online services and data centre capacity needs.
Localisation	Concerns over data security drives legislative changes and need for onshore data storage capabilities.
AI	Artificial intelligence requires significant data storage and computing capacity. Accelerating AI capabilities including generative AI likely to be a key future driver for data centre capacity.

Source: (Deloitte, 2024), Standard Chartered Bank

The different types of data centres

The changing nature of data storage and processing requirements has meant that the structure of the data centre market has shifted during the past 25 years.

- **Onsite:** Development of computer networks, decentralised data processing allowed companies to start managing data more efficiently. The result of this was a development of company-owned data centres. Onsite DCs accounted for well over 90% of the entire DC market in 2000 (Synergy Research Group, 2025).
- **Co-location:** The commercialisation of internet-based services increased demand across companies large and small for data centres that could support e-mail services, e-commerce and online applications. To allow smaller companies access to these services too, co-location data centres were created where companies can lease capacity for these services.
- **Hyperscalers:** The development of cloud-based computing services allowed corporates to further streamline their IT operations and offered some of the largest technology companies with the potential to offer large scale on-demand cloud services through hyperscale data centres. At present hyperscaler-related DC capacity makes up more than 50% of the global market (Figure 3).
- **Neoclouds**: The scarcity of high-end computing power especially because of the development of (generative) AI models has allowed more than 100 Neocloud companies to develop. They offer a Bare-metal-as-a-service (BMaaS) cloud computing model that allows vendors to lease physical servers for their computing needs. Neoclouds typically lease data centre capacity, however, this may change especially if they pivot towards sovereign clients who value independence from hyperscalers (Mazza & et al, 2025).

The gradual shift from smaller more decentralised onsite data storage facilities to larger hyperscaler and neocloud facilities has meant that power and water demand has become more centralised and help create related environment challenges.

Data centre capacity growth is accelerating

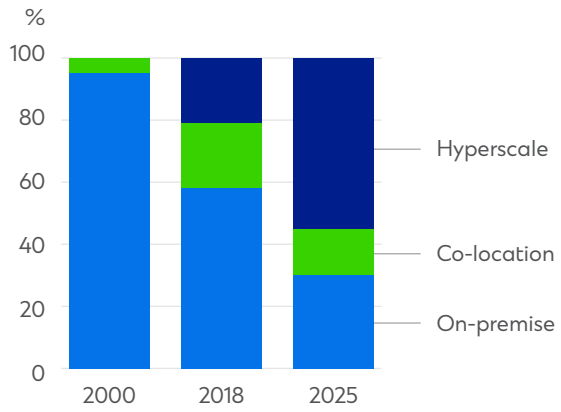
Data from BloombergNEF for data centre capacity shows that this has started to accelerate during the past few years (Figure 5 and 6). Total headline capacity stood at more than 106GW at the end of 2025 while live capacity reached more than 60GW. Annual growth in live capacity and the amount of capacity under construction has rapidly increased since 2022 (Figure 6). Bloomberg's data shows that c. 60% of current data centre capacity globally was only installed in the last five years.

Bloomberg's data lists more than 2000 companies globally with data centre capacity, however, data centre capacity is highly concentrated as just 15 companies account for c. 50% of global capacity, with Amazon, Meta, Microsoft and Google together accounting for almost 26% (Figure 4).

Figure 3

The changing data centre landscape

% of total data centre capacity

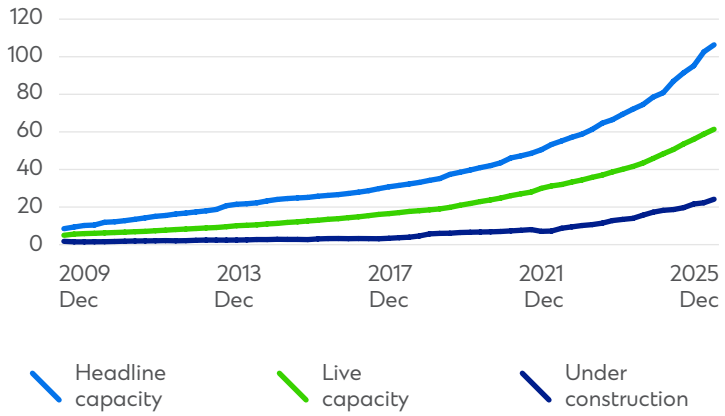


Source: SRG Research, Standard Chartered Bank

Figure 5

Total data centre capacity

GW

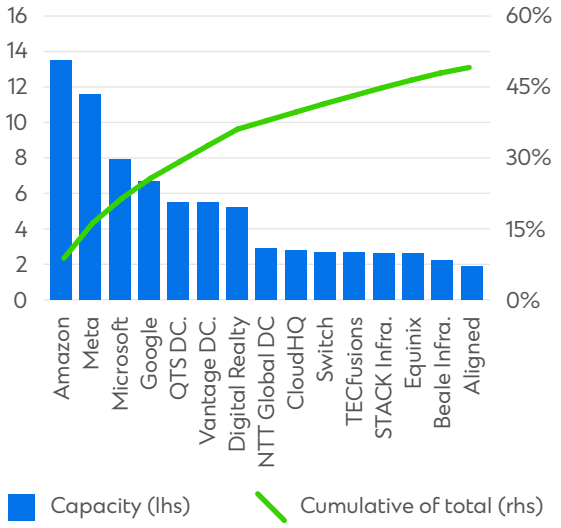


Source: SRG Research, Standard Chartered Bank

Figure 4

Largest data centre owners

GW, 2025

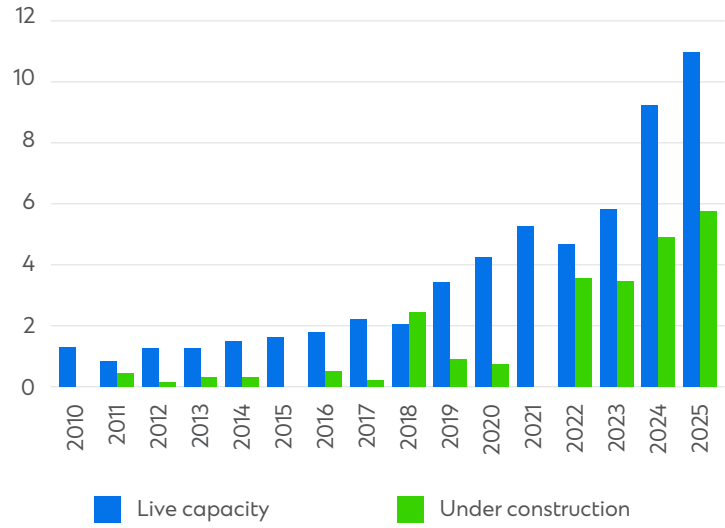


Source: SRG Research, Standard Chartered Bank

Figure 6

Annual data centre capacity growth

GW



Source: Bloomberg, Standard Chartered Bank

Data centre growth is a global rather than US-only topic

Historically, the US accounted for more than 40% of global data centre capacity, however, during the past few years this has accelerated to more than 50% as hyperscalers such as Amazon, Microsoft, Google and Meta started to increase investments (Figure 7).

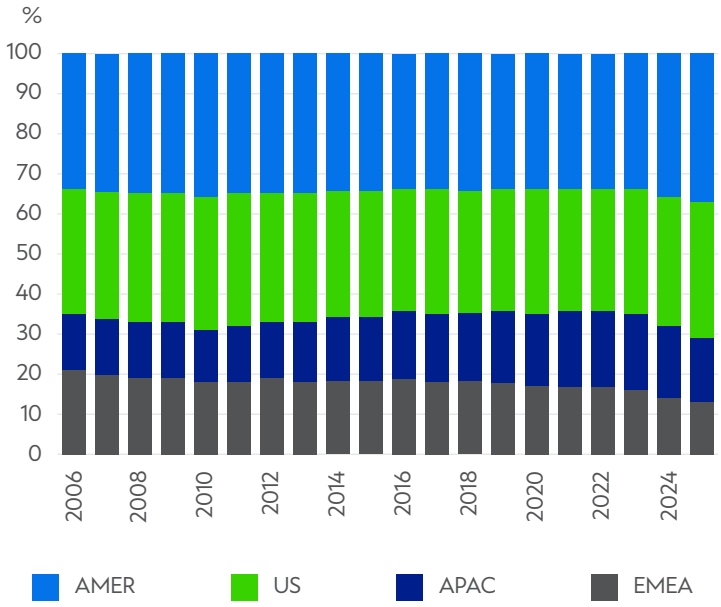
Notwithstanding the dominance by US data centre companies it is worth noting that outside of the US data centres are becoming a strong growth story too. Countries across the Global South with the greatest amount of capacity are China, India, Indonesia and Brazil (Figure 8).



Figure 7

Live data centre capacity by region

2025, GW

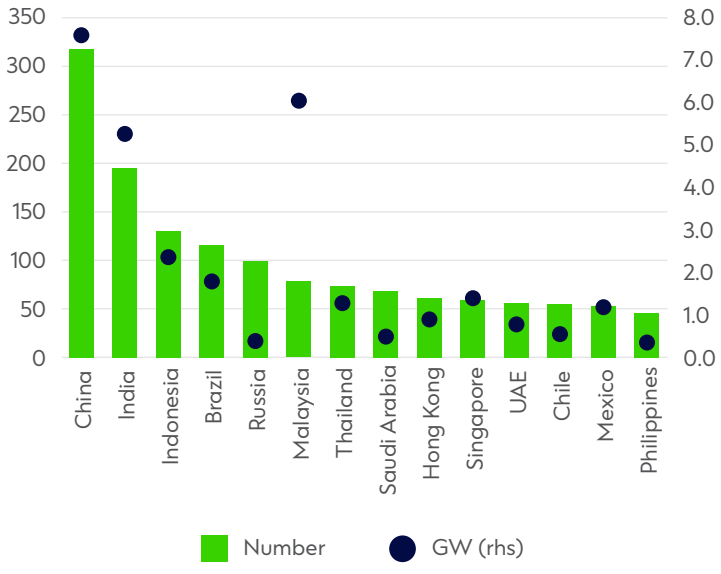


Source: Bloomberg, Standard Chartered Bank

Figure 8

Largest data centre presence outside Europe and North America

(2025)



Source: Bloomberg, Standard Chartered Bank

Asia emerging as a major data centre growth frontier

Data centre live capacity in the region has increased by c. 60GW between 2000 and 2025 (Figure 6). However, capacity may increase by the same amount again during the next five years if industry forecasts prove correct (McKinsey (a), 2025), (JLL, 2026) and (IEA, 2025) (Figure 9). AI is seen as the key driver for this continued acceleration in capacity growth which may increase AI's share of data centre capacity in Asia to 70% in 2030 (Figure 10).

The rapid increase in AI-related data centre capacity is seen as a tailwind for neocloud-related centres too. Industry estimates suggest that the number of neocloud data centres may increase more than 3-fold to more than 2,200 by 2035 (ABI Research (a), 2026). As far as capacity growth is concerned neoclouds may represent 36% of the total market only slightly behind the 39% share expected to be occupied by the Tier-1 hyperscalers such as Amazon, Google, Meta and Microsoft (ABI Research, 2026) (Figure 11).

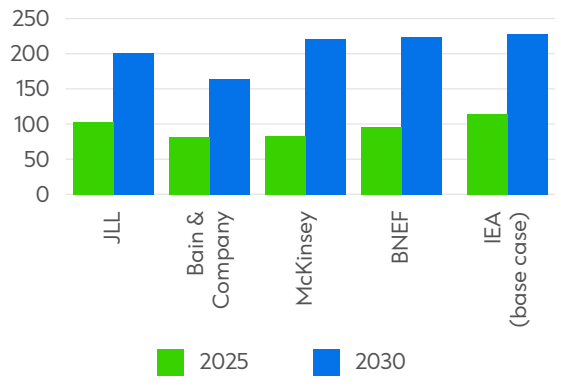
US dominance may continue but growth elsewhere could be stronger

Potential data centre growth by region shows that capacity expansion in the Americas may increase by 122% between 2025 and 2030 (JLL, 2026). The growth rate for Europe, the Middle East and Africa is estimated to reach almost 80% while growth across Asia and the Pacific is forecast at 62% (Figure 12).

Figure 9

Data centre capacity scenarios

GW

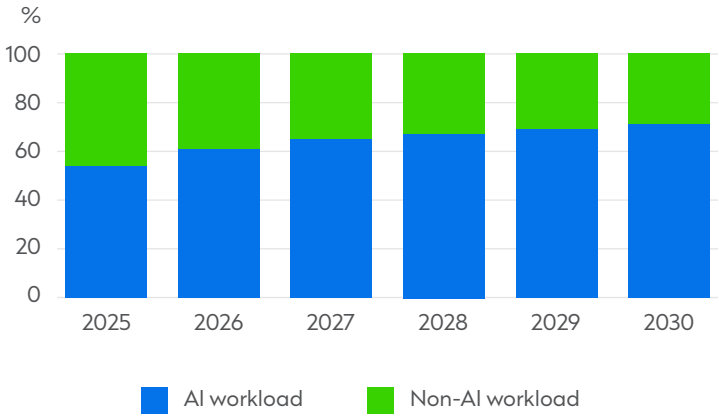


Source: McKinsey, Standard Chartered Research

Figure 10

AI is driving data centre capacity increases

% of data centre capacity requirement by activity

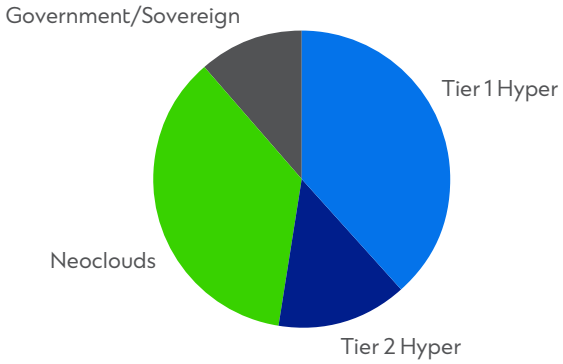


Source: (McKinsey (a), 2025), Standard Chartered Bank

Figure 11

Data centre capacity growth by user

GW added between 2025-2035

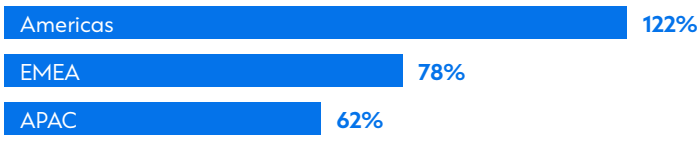


Source: (ABI Research (b), 2026), Standard Chartered Bank

Figure 12

Data centre capacity growth by region

(% change in GW, 2030 vs 2025)



Source: JLL, Standard Chartered Bank

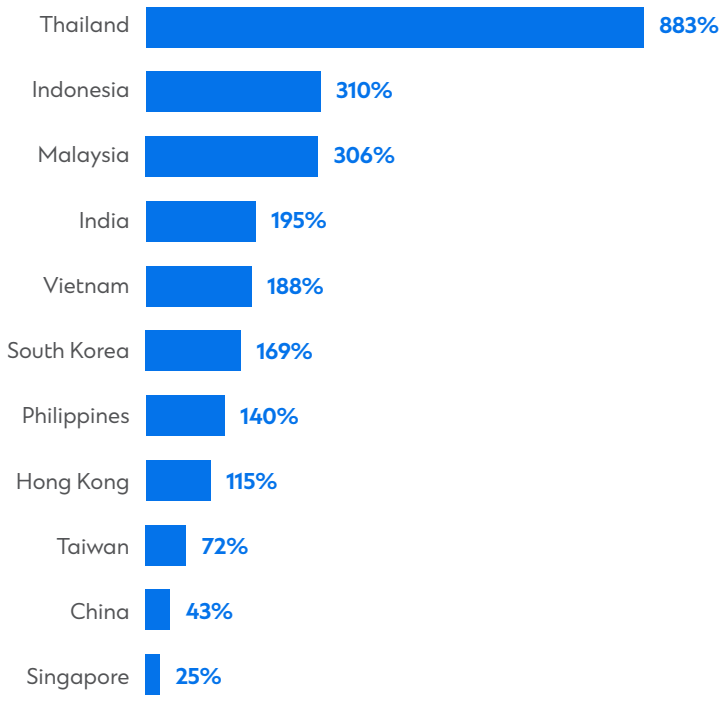
Asia is becoming the next major battleground for data centre growth

While headline growth for data centre capacity in APAC appears more moderate than for the Americas or EMEA, this masks the scale of expansion underway outside China (Figure 13 and 14). Capacity additions in markets such as Malaysia and India could narrow the gap with China by 2030, while many Asian countries are expected to more than double data centre capacity between 2025 and 2030 (Figure 14). This suggests data centres may move up the policy agenda across Asia, with governments facing pressure to balance digital infrastructure ambitions with power, water and land constraints.

One of the reasons that many countries across the Asian region may experience strong data centre growth during the next few years may be consumer engagement. The most recent Ipsos AI monitor poll (Ipsos, 2026) shows that consumers across Asia are more excited about the potential impact and benefits of AI than those located in other countries especially consumers across mature economies such as Canada, the UK and the USA (Figure 15).

Figure 14

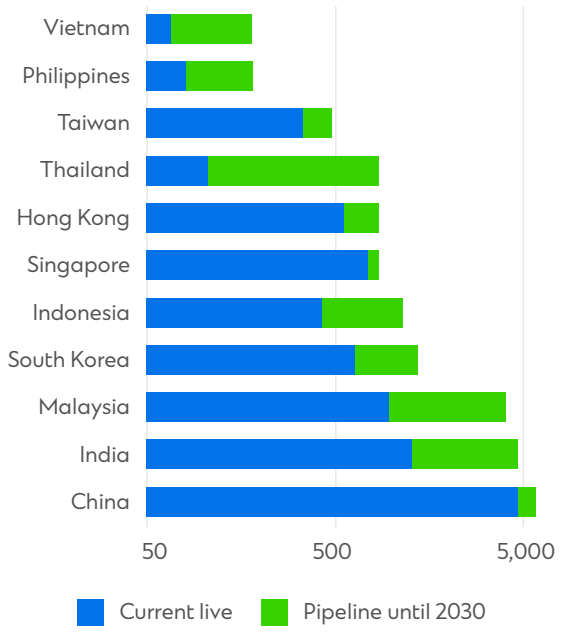
Data centre capacity growth scenario by Asian country
(% 2030 vs current)



Source: (Cushman & Wakefield, 2026), Standard Chartered Bank

Figure 13

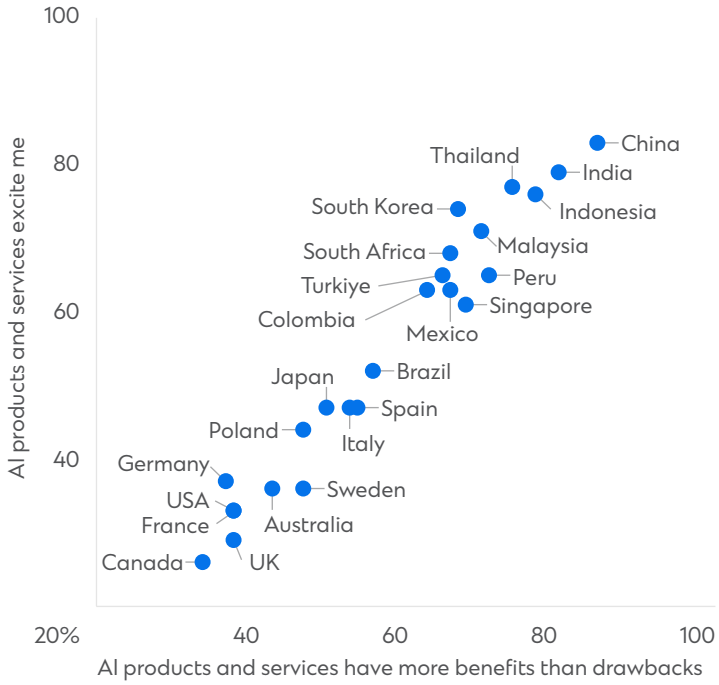
Data centre capacity scenario for countries across Asia
(MW)



Source: (Cushman & Wakefield, 2026), Standard Chartered Bank

Figure 15

Consumer excitement with AI
(% of surveyed consumers by country)



Source: Ipsos, Standard Chartered Bank

Data centre capacity is likely to become much larger

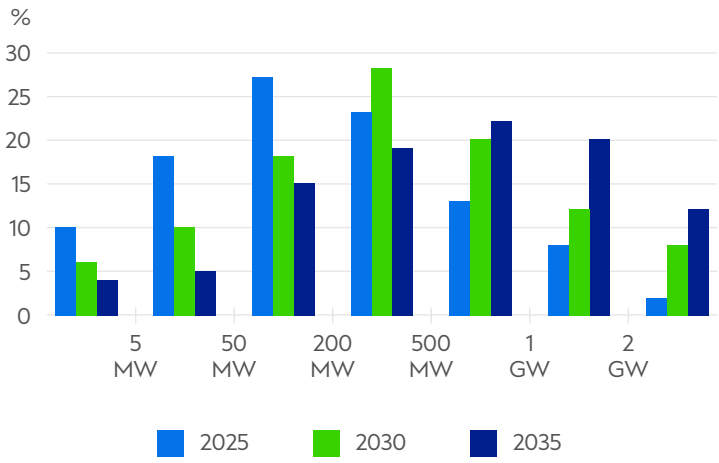
AI is changing the amount of computing power that data centres need. Training and running advanced AI models requires substantial parallel processing across dense clusters of GPUs, making speed, proximity and interconnect efficiency critical priorities. As a result, operators are likely to compress more compute power into each cabinet and rack, reducing latency and copper/fibre interconnect requirements but sharply increasing power density. The build-out of AI infrastructure is therefore set to drive a shift toward more energy-intensive and technically complex data centres (Datacenters.com, 2026).

The latest Data Centre Power Report by Bloom Energy puts this in perspective (Bloom Energy, 2025). It suggests that by 2030 about one in five data centre campuses are expected to exceed 1GW, rising to one in three by 2035 (Figure 16). This compares to just one in ten in 2025. By how much the capacity of individual data centres ultimately increases is not just a function of growing processing requirements but also depends on how efficiently these are being executed. An acceleration in data centre efficiency levels or potentially new technologies including quantum computing might limit the need for a strong rise in data centre power and capacity demand.

Figure 16

Data centre universe by size

(% of data centres with certain capacity)

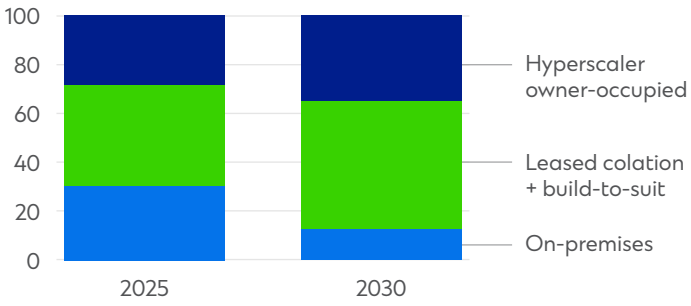


Source: (Bloom Energy, 2026), Standard Chartered Bank

Figure 17

Data centre capacity by owner

(% of total)



Source: JLL, Standard Chartered Bank

AI may trigger investments of up to USD9 trillion

Investment activity levels related to data centres have increased rapidly during the past 10 to 15 years. Capital expenditure by the largest data centre companies globally has increased c. 15-fold since 2010 to around USD400 billion in 2025 (Figure 18). Venture capital investments into data centre-related activities have increased more than 40-fold since 2012 to circa USD80 billion in 2025 (Figure 19). To facilitate the future expansion of AI requires investment levels to increase further.

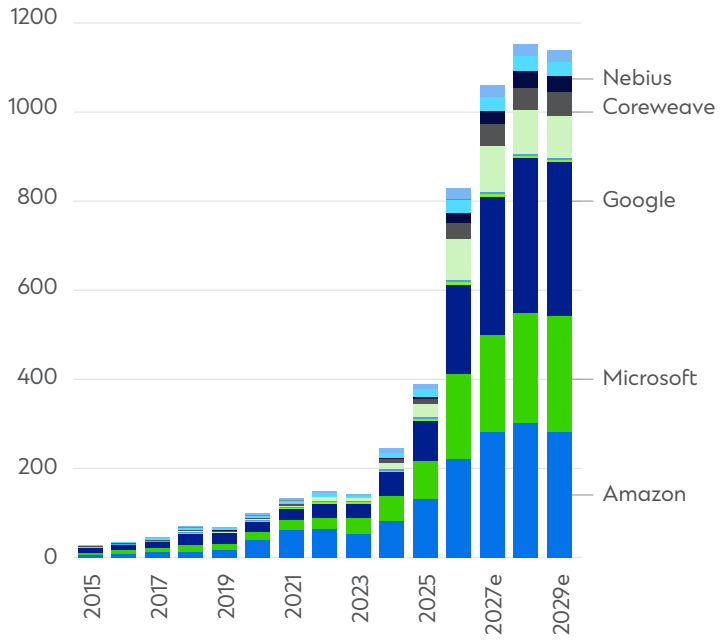
Consensus expectations for the companies in Figure 18 indicate that their cumulative capital investments for the 2026-2029 period alone may reach USD4.8 trillion or roughly three times the total amount invested during the previous 10 years. These estimates correspond to those from organisations such as McKinsey, the International Energy Agency (IEA) and JLL that show total data centre investments for the 2025-2030 period totalled between USD3-6.6 trillion (IEA, 2025), (JLL, 2026), (McKinsey (a), 2025). Wood Mackenzie estimates that data investments may reach USD9 trillion between now and 2040 (Wood Mackenzie, 2026).

Data centre-related investments can be grouped into three major pools:

- **IT hardware is the largest and most immediate investment driver.** Servers, networking equipment, memory and storage remain the core of data centre spend, with the IEA estimating that IT-related hardware accounted for c. 70% of total data centre investment in 2025 (IEA, 2025). This makes semiconductor, server and networking supply chains the first-order beneficiaries of AI-driven capacity growth (Figure 21).
- **Physical infrastructure is becoming increasingly complex and capital intensive.** Beyond the IT layer, data centres require significant investment in the building shell and supporting mechanical and electrical systems, including cooling, transformers and uninterruptible power supplies (Figure 21). These components can account for close to 60% of non-IT equipment spend, highlighting the growing importance of electrical, cooling and construction supply chains as facilities become larger and more power dense (McKinsey (c), 2025).

Figure 18

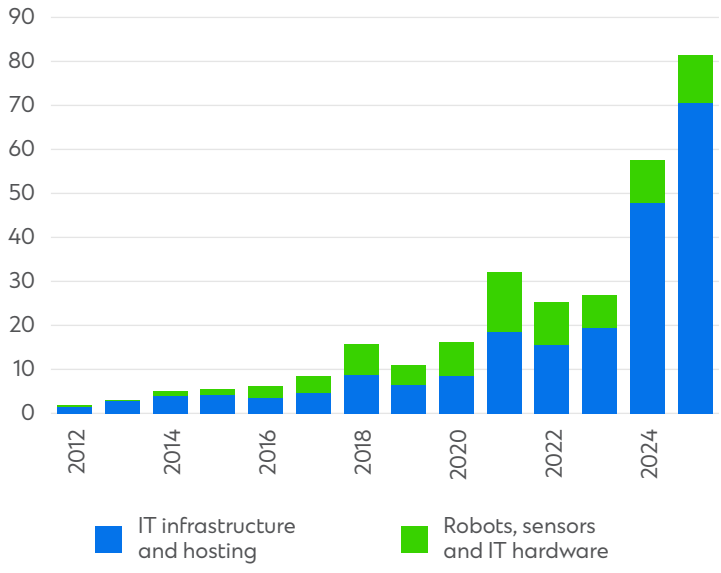
Capital expenditure levels by key data centre companies
(USD billion)



Source: Companies mentioned, Bloomberg, Standard Chartered Bank

Figure 19

Venture capital investments in data centre-related activities
(USD billion)



Source: (OECD.AI, 2026), Standard Chartered Bank

- **Energy infrastructure is emerging as a critical constraint.** Additional data centre load requires substantial investment in power generation, battery storage, transmission and distribution. The IEA estimates data centre-related energy investment needs at close to USD500 billion, while BloombergNEF puts associated grid-related investment at c. USD480 billion (IEA, 2025); (BloombergNEF, 2025). This suggests the data centre build-out is increasingly becoming a power-sector investment story as much as a digital infrastructure one.

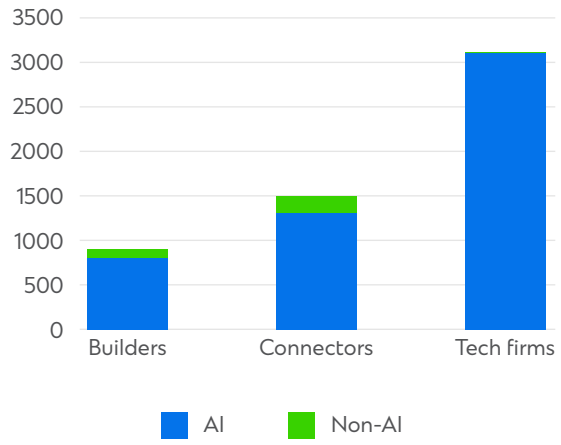
Data centre construction costs vary sharply by location, creating a possible competitive advantage for some markets

Cost benchmarks from (Turner and Townsend, 2025), (JLL, 2026) and (Cushman & Wakefield, 2026) suggest that new data centre capacity in China and India can be built for less than USD7 mn/MW, compared with USD12.5-15.6mn/MW in higher-cost markets such as the UK, Singapore and Japan (Figure 22). This spread matters: as AI-driven demand accelerates, construction cost differentials could influence where incremental capacity is built, particularly where power, land and permitting constraints are also manageable.

Figure 20

Data centre CAPEX by enabling sector

(2025-2030, USD billion)

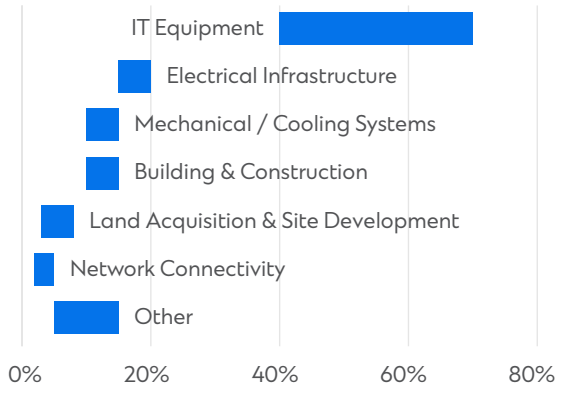


Source: McKinsey, Standard Chartered Bank

Figure 21

Cost drivers for data centres

% of construction costs

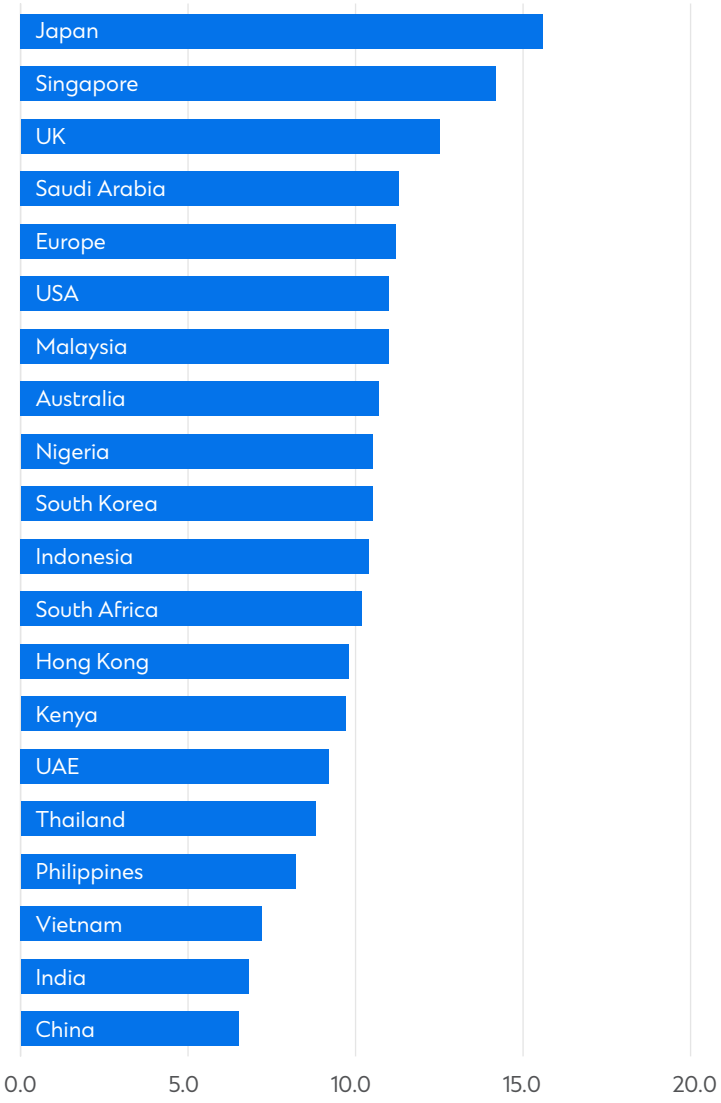


Source: Uptime Institute, McKinsey, CBRE, JLL, Standard Chartered Bank

Figure 22

Data centre construction costs by location

(USD million / MW)



Source: (Turner and Townsend, 2025), (JLL, 2026) and (Cushman & Wakefield, 2026) Standard Chartered Bank

Cost inflation is becoming an additional constraint on the data centre build-out

Strong demand for new capacity is already pushing up construction costs, and survey evidence suggests these pressures are unlikely to ease soon. Turner & Townsend's latest data centre cost survey found that 47% of operators saw costs rise by 6-15% last year, while 21% experienced increases of more than 15%. Looking ahead, 60% expect costs to rise by a further 5-15% over the next 12 months, and 21% expect increases above 15% (Figure 23). This points to a market where demand growth is not only raising the scale of required investment but also increasing the unit cost of delivery.

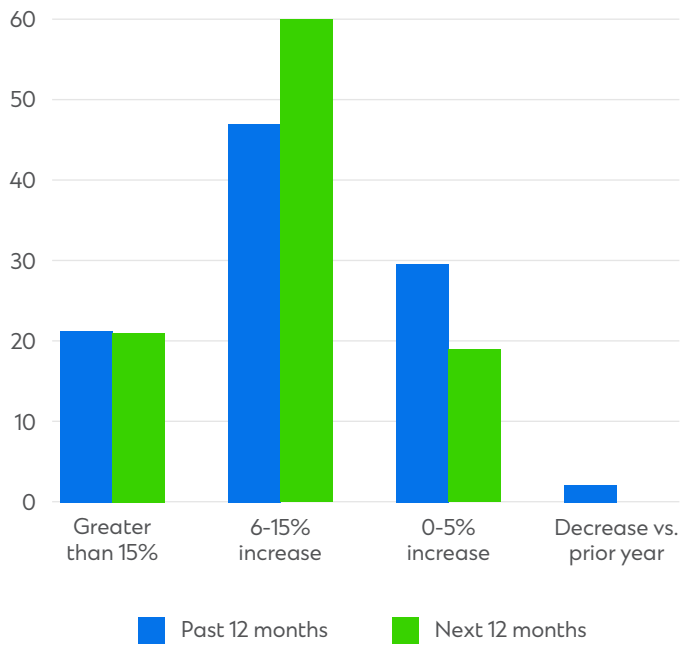
Equipment availability may become a bottleneck for data centre delivery

The build-out of new capacity increasingly depends on securing specialised electrical and mechanical equipment, yet supply chains are struggling to keep pace with demand. JLL estimates that average global lead times for electrical and mechanical equipment are now around 33 weeks, while critical components such as transformers, generators and switchgear can take more than a year to procure (2026 Global Data Center Outlook). Even with developers ordering materials up to two years in advance, 57% of data centre projects still experienced construction delays of at least three months in 2025. This suggests that equipment scarcity may become a constraint on how quickly AI-driven data centre capacity can be brought online.

Figure 23

Data centre construction cost change

% of construction costs



Source: (Turner and Townsend (b), 2025), Standard Chartered Bank

Data centre-related funding needs until 2029 may reach USD1 trillion

The rapid expansion of data centre infrastructure is already driving a rise in external financing needs. According to S&P Global, data centre-related debt issuance almost doubled from USD92 billion in 2024 to USD182 billion in 2025 (CNBC, 2025). Momentum has continued in 2026, with Reuters estimating that hyperscalers alone have raised approximately USD250 billion in debt year-to-date (Reuters, 2026a). As investment in AI and data centre capacity accelerates, funding requirements are likely to rise further.

Consensus forecasts for nearly 90 publicly traded companies across the data centre and AI value chain suggest that capital expenditure growth is set to outpace internally generated cash flows over the coming years. Using analyst expectations from Bloomberg for revenues and capital expenditure, and assuming companies maintain current operating cash flow conversion rates, we estimate a cumulative funding shortfall for this group of data centre exposed companies alone of more than USD660 billion between 2027 and 2029. This gap would need to be financed through external sources, primarily debt markets. Annual funding requirements could reach almost USD330 billion in 2027, USD250 billion in 2028, and nearly USD90 billion in 2029.

The funding requirements associated with the AI and data centre build-out are likely to be significantly larger than those implied by listed data centre companies alone. The estimates above exclude the capital needs of privately held AI platforms such as OpenAI, Anthropic and xAI, as well as the investment required to expand power generation, transmission and distribution infrastructure needed to support rising data centre electricity demand.

Public disclosures suggest that frontier AI companies alone may require substantial capital with OpenAI forecasting cumulative negative free cash flow of approximately USD278 billion between 2026 and 2030 (Financial Times, 2026) while Anthropic has disclosed long-term data centre and general AI-related infrastructure spending commitments that may exceed USD800 billion (Financial Times). The power sector will also require substantial investment to support AI-driven electricity demand growth. The International Energy Agency estimates that incremental energy infrastructure investment associated with rising data centre electricity consumption could approach USD500 billion by 2030 (IEA, 2025).



Taken together, these estimates suggest that funding requirements linked to the broader AI and data centre investment theme could exceed USD1 trillion by the end of the decade. Given the scale of capital required and the long-lived nature of the underlying assets, debt markets are likely to play an increasingly important role in financing the next phase of AI and digital infrastructure growth.

The evolving landscape of data centre funding structures

The financing mix supporting AI and data centre expansion changes depending on the maturity of development. Construction loans, project finance and private credit are typically relied upon during the development or construction phase, while corporate bonds, infrastructure debt and revolving credit facilities are most often used when facilities or assets achieve stable cash flows. Operators with larger pools of data centres may during these more stable phases also explore securitisation products such as asset-backed securities (ABS) and commercial mortgage-backed securities (CMBS) too. Securitisation can help to broaden the investor base and diversify funding sources and in doing so may become a more important source of long-term institutional capital as operators seek to refinance their assets (Structured Finance Association, 2026).

Sustainable finance as a growing source of funds

This report shows that power consumption, water use and local environmental aspects are important considerations for data centre and AI development. Against that backdrop sustainable finance can become a suitable and growing funding channel for AI and data centre infrastructure.

Bloomberg estimates that data centre developers have issued approximately USD186 billion of sustainable debt since the launch of ChatGPT in late 2022, underscoring the growing role of green bonds, sustainability-linked loans and other labelled financing instruments in supporting AI and data centre expansion (OneStop ESG, 2026). Sustainable Fitch notes that sustainable debt issuance linked to data centres has increased more than ten-fold since 2022, reflecting both rising investor demand for climate-aligned infrastructure and growing issuer interest in sustainable financing frameworks (Sustainable Fitch, 2026).

Several sustainable finance instruments are particularly relevant to the sector:

- **Green bonds and green loans** are well suited to funding renewable energy procurement, on-site and off-site clean power generation, battery storage, energy-efficient cooling systems, heat recovery infrastructure and water efficiency projects.
- **Sustainability-linked loans (SLLs)** provide an additional financing tool by linking borrowing costs to sustainability performance indicators such as renewable energy usage, power usage effectiveness (PUE), water usage effectiveness (WUE), greenhouse gas emissions intensity and other operational efficiency metrics for data centre facilities.
- **Sustainability-linked bonds (SLBs)** may also become increasingly relevant as operators seek to align financing costs with long-term decarbonisation, water stewardship and energy efficiency targets.

The development of industry-wide sustainability standards could further accelerate the use of sustainable finance within the sector. In April 2026, nine leading built-environment and sustainable finance organisations launched the Greening AI Data Centres Coalition (GADCC) to establish transparent and credible benchmarks defining what constitutes a sustainable or "green" data centre (Climate Bonds, 2026). The initiative aims to improve consistency, reduce greenwashing concerns and provide investors with greater confidence when allocating capital to digital infrastructure projects. The coalition specifically highlights the need for clear sustainability definitions as AI-driven demand for data centres continues to accelerate.

Elsewhere in this report, we identify several characteristics that may determine whether a data centre can be considered environmentally sustainable, including renewable energy sourcing, energy efficiency, water stewardship and broader environmental performance. As reporting standards, sustainability benchmarks and disclosure frameworks continue to mature, sustainable finance instruments are likely to play an increasingly important role in funding the next generation of AI and data centre infrastructure.

Aggregator platforms as an additional provider of capital

The large scale of investment required to support AI and data centre expansion may create funding challenges, particularly for smaller operators and emerging infrastructure providers seeking access to long-term capital. In response, a new generation of specialist financing and capital aggregation platforms is beginning to emerge. One example of this is the AI infrastructure financing initiative established by NVIDIA in cooperation with Apollo, BlackRock, Blackstone, Brookfield, Goldman Sachs and KKR (NVIDIA, 2026a). The aim of the platform is to mobilise over USD500 billion of third-party capital for the build-out of AI infrastructure.

The significance of this initiative extends beyond its scale. By treating AI compute infrastructure as an investable asset class, the platform seeks to connect institutional investors directly with the rapidly growing demand for AI infrastructure. This could broaden access to funding, lower financing costs for some operators and create new channels through which pension funds, insurers, infrastructure investors and private credit providers can participate in AI-related growth.

Rising borrowing costs may favour higher-quality, more transparent operators

Data centre expansion is highly capital intensive and, like other infrastructure sectors, depends heavily on external debt financing. This leaves the sector sensitive to changes in interest rates and bond market conditions. Since the start of the year, government bond yields have risen across most markets amid rising inflation, broader macroeconomic uncertainty and heightened geopolitical concerns.

Higher government bond yields have increased borrowing costs for infrastructure projects, including data centres. To assess the implications for data centre operators and hyperscalers, we analysed more than 60 corporate bond, asset-backed securities (ABS) and commercial mortgage-backed securities (CMBS) issuances across 26 issuers. Two key conclusions emerge:

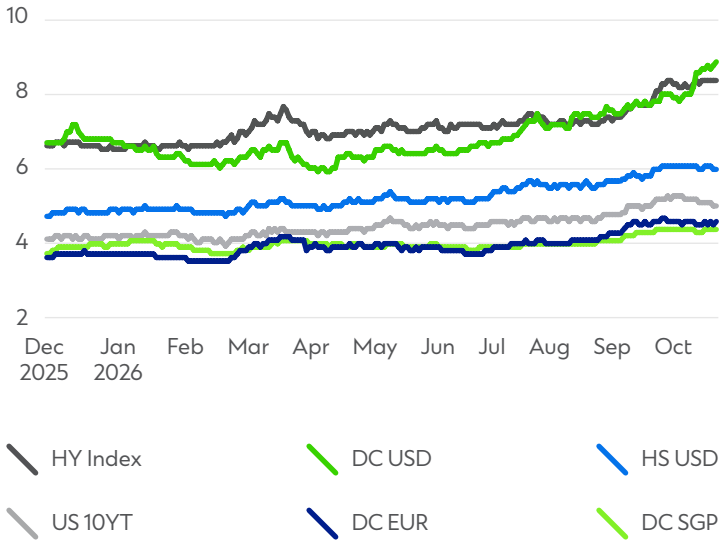
- **Issuer-specific credit risk is a key differentiator.** Bond yields have increased this year for both data centre operators and hyperscalers, but the rise has been more modest for larger, often better-capitalised hyperscalers (Figure 24). Yields on bonds issued by smaller data centre operators and neoclouds have risen more sharply, moving more closely in line with the broader US high-yield bond market (Figure 25). This may reflect their often high-yield-like credit ratings (Forbes) and growing investor concerns about potential increases in leverage. Current market conditions therefore appear to favour larger, higher-quality borrowers.
- **Currency diversification may offer benefits.** Our analysis suggests that euro- and Singapore dollar-denominated bonds have been less affected by current market conditions than US dollar-denominated bonds. This pattern holds both on average and for individual issuers with bonds in multiple currencies. All else being equal, these findings suggest that access to funding in multiple currencies may benefit data centre borrowers financing their expansion plans.

Investor concerns about rising sovereign debt levels and increased borrowing by data centre operators may persist. Operators seeking funding in this environment may benefit from strengthening their balance sheets and cash flows, alongside adopting the transparency practices outlined elsewhere in this report. These measures could help address investor concerns and support confidence in their credit quality.

Figure 24

Yield development

Average by currency, DC = data centre operators, HS = Hyper scalers, HY = US High Yield Index

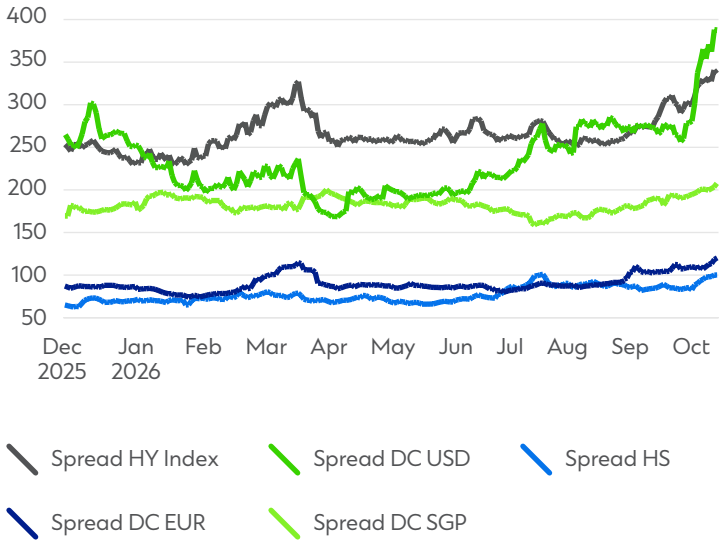


Source:, Bloomberg Standard Chartered Bank

Figure 25

Yield difference with sovereign benchmarks

Basis points, DC = data centre operators, HS = Hyper scalers, HY = US High Yield Index



Source: Bloomberg, Standard Chartered Bank



Solving the data centre power challenge

Power may become a defining constraint for the growth of the AI sector as data centre electricity demand could rise from almost 500TWh in 2025 to c. 1,100TWh by 2030, with AI's share increasing from 31% to almost 50% (Figure 27). The challenge is not only how to meet this rising demand, but how to support digital growth without placing unnecessary pressure on global and local electricity systems, energy costs or decarbonisation efforts.

This chapter highlights five potential solutions that could help reduce the power required per unit of compute and reposition leading data centres as flexible participants in the energy system, rather than largely passive electricity users. These are improving operational efficiency, reusing waste heat for residential or industrial purposes, expand exposure to renewable energy power sources, adopt high voltage direct current electricity distribution and develop battery storage and microgrid capacity. Scaling these solutions and deploying AI effectively across different industries could meaningfully reduce the impact of AI and data centre development on net electricity demand.

AI puts a spotlight on power demand and the grid

Data centres require power not only to run servers and networking equipment, but also to support cooling, space conditioning and auxiliary systems such as lighting. As data centre capacity has expanded and internet-based services have become increasingly embedded in the global economy, electricity consumption has risen sharply. Although data gaps remain regarding data centre electricity use (Masanet & et al, 2024), (Mytton & Ashtine, 2022) the IEA estimates that data centre power demand has more than doubled since 2010, reaching almost 500TWh in 2025 (IEA, 2026). While this still represents only c. 1.5% of global electricity demand, up from 1.4% in 2024, the direction of travel is clear: data centres are moving from a relatively small load category to an increasingly important factor in power-system planning (Ember, 2026).

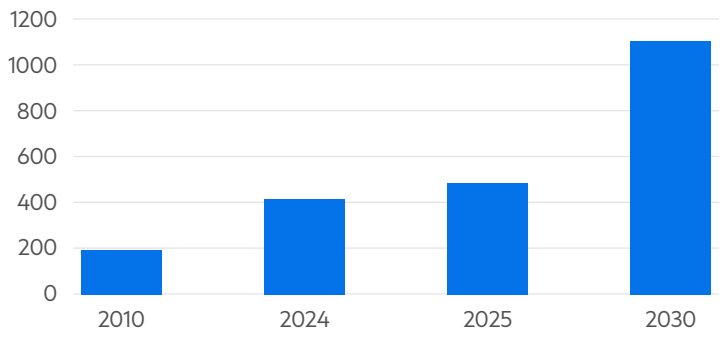
If industry forecasts prove correct and global data centre capacity roughly doubles between 2025 and 2030, electricity consumption may follow a similar trajectory unless efficiency gains can be achieved. Electricity demand is not just driven by the amount of data centre capacity but also how this capacity is used. With some estimates suggesting an AI query can require up to ten times more electricity than a traditional search engine request, rising AI adoption could materially increase the energy intensity of data centre services (AI Kez & et al, 2025).

Averaging forecasts from the IMF, IEA, Gartner and McKinsey suggests data centre power demand could reach c. 1,100TWh by 2030 — around twice 2025 levels (Figure 26). AI is expected to be the main accelerator of this growth: IEA estimates imply AI could account for almost half of total data centre electricity demand by 2030, up from 31% in 2025 (Figure 27). This makes AI not only a catalyst for data centre capacity growth, but a structural force reshaping electricity demand, grid planning and energy infrastructure investment.

Figure 26

Data centre power demand

TWh

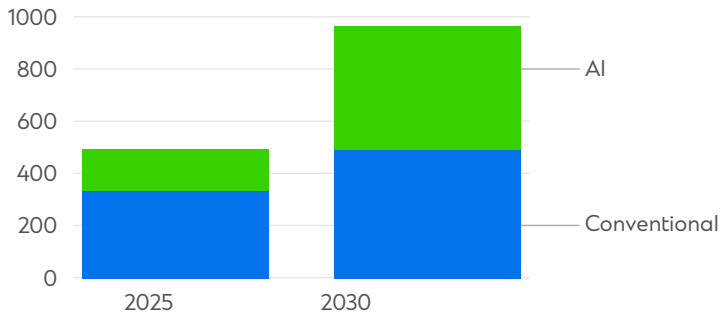


Source: (IEA, 2026), (McKinsey (b), 2025), (Deloitte, 2024), (IMF, 2025), (Gartner, 2025), Bloomberg, Standard Chartered Bank

Figure 27

AI set to drive data centre power demand

TWh



Source: IEA, Standard Chartered Bank

Data centre power demand can become material locally

IEA estimates suggest that data centres account for less than 10% of global electricity demand growth, yet this global figure masks the pressure that data centre capacity additions can place on local power systems. Modelling work from Bloomberg suggests that the power share of data centres in the US and Europe may rise to around 8% of total electricity usage by 2035 while the Lawrence Berkeley National Laboratory suggests that data centre power usage in the US may reach 11.8% by 2030 (Smith, Hubbard, & et al, 2026). Asian markets may experience supply-demand constraints too as they see faster growth in data centre power albeit from a lower base (Figure 28 and 29).

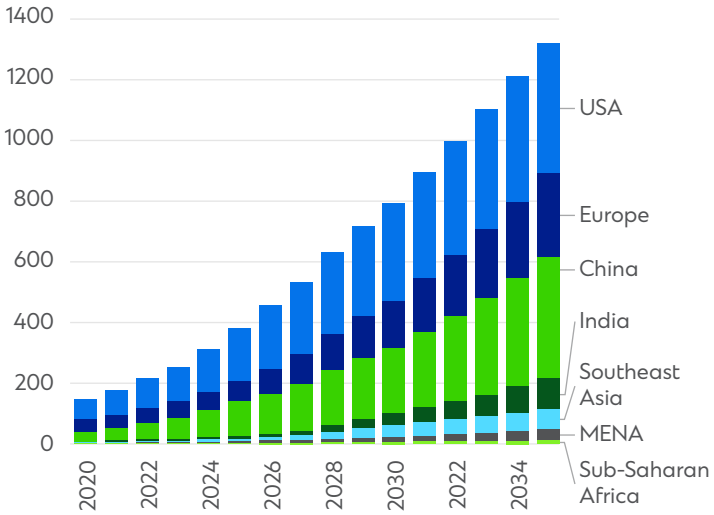
As data centre growth broadens across more markets, supply-demand constraints are likely to become increasingly jurisdiction-specific, shaped by local grid capacity, generation availability and permitting timelines. Securing grid access and a supply of sufficient power are key challenges faced by data centres.



Figure 28

Data centre power demand likely to accelerate across all regions

TWh

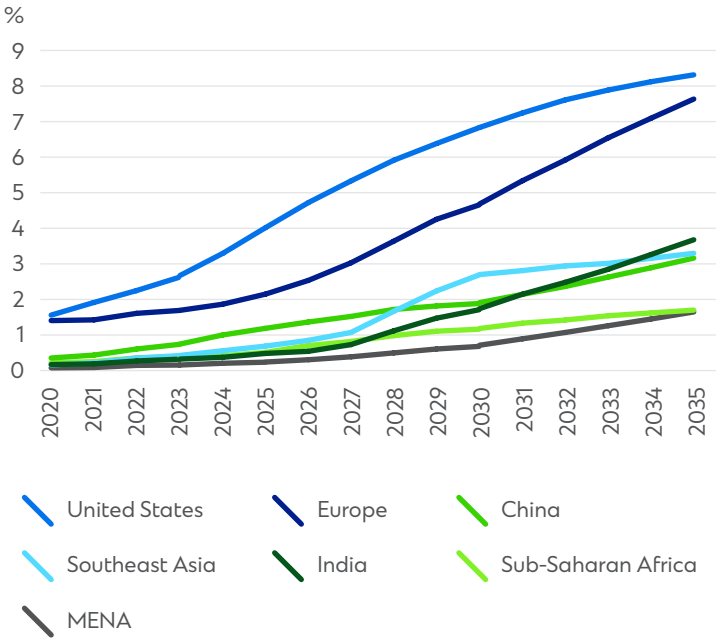


Source: Bloomberg, Standard Chartered Bank

Figure 29

Data centre share of total power demand

% by market



Source: Bloomberg, Standard Chartered Bank

Meeting growing power demand requires a stronger grid

Securing timely, reliable and affordable power, while minimising the impact on electricity availability for local economies, is emerging as one of the data centre sector's most critical strategic challenges. How effectively operators address this constraint will likely play a major role in determining the pace and scale of future industry growth. The latest survey from the European Data Centre Association (EDCA) underlines the point: almost 70% of European data centre operators identified access to power as one of their biggest issues (EDCA, 2026) (Figure 30). The constraint is increasingly visible globally. The IEA estimates that around 20% of planned data centre capacity additions could be at risk of delay because of grid constraints, while its data shows that connection-delay risk is becoming a material drag on build-out plans through 2030 (World Energy Outlook, 2025).

Grid access may also start to shape location decisions: connection timelines can range from around 10 years in markets such as Amsterdam and Tokyo to roughly two years in Mumbai, Thailand and Malaysia (Figure 31). In the UK, Ofgem – the independent regulator for the gas and electricity markets – has warned that demand-connection queues have grown rapidly, with data centres a major contributor, prompting reforms to filter speculative projects and accelerate viable connections (Ofgem, 2025). In the US, Bloom Energy's 2026 survey found that utility delivery timelines are c. 1.5-2 years longer than hyperscalers and colocation providers expect, and that power availability is already reallocating growth toward more power-advantaged regions such as Texas and Georgia (Bloom Energy, 2026). These examples suggest that power access is moving from an operational constraint to a core determinant of data centre competitiveness, site selection and market growth.

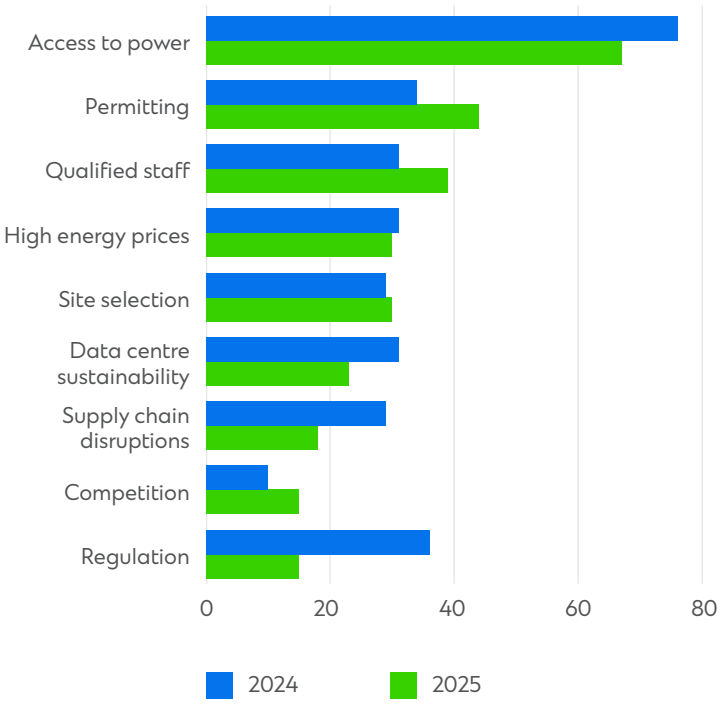
Securing power for data centres is broader than obtaining a grid connection. Other potential challenges include whether the wider power system can generate, transmit and distribute enough reliable electricity to large, concentrated loads.

- **Generation adequacy:** a grid connection only matters if there is enough firm power behind it. AI data centres create large, concentrated and high-load-factor demand that must be served around the clock. To put this in context the IEA's estimate for the incremental increase in electricity demand from data centres during the next five years is equivalent to the annual electricity consumption of Germany or France.

Figure 30

Biggest issues for data centre operators

European operators

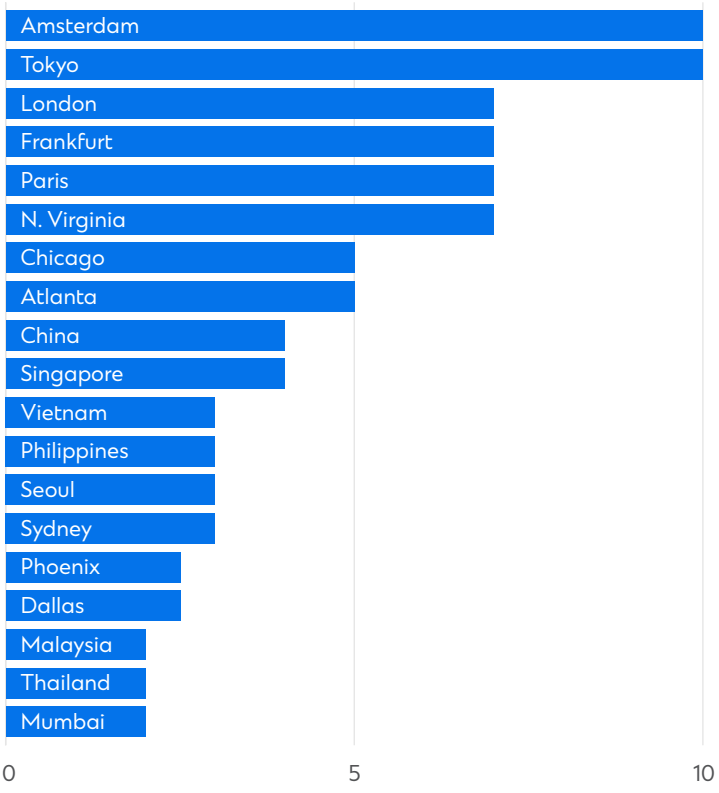


Source: EDCA, Standard Chartered Bank

Figure 31

Connection times for data centres

Years



Source: JLL, Cushman - Wakefiled, Standard Chartered Bank

- **Transmission capacity:** sufficient generation at national level does not guarantee power deliverability at the data centre site. S&P Global notes that transmission and distribution constraints remain a key issue as data centre load growth stresses the system, with interconnection queues still overcrowded and large-load facilities facing longer timelines (S&P Global, 2025). BloombergNEF estimates that global grid investment may have exceeded USD470 billion in 2025, that annual grid investment needs for the 2025-2029 period are c. 60% higher than the 2020-2024 average and that a total of USD15.8 trillion needs to be invested in the global electricity grid between now and 2050. These estimates reflect the scale of investment needed to address general electricity network bottlenecks which impacts data centres too (BloombergNEF, 2025).
- **Local distribution capacity:** data centres connect to specific grid nodes, and those nodes may not have enough headroom. This is why constraints are often local even when national power balances appear manageable. Deloitte notes that some US grid connection requests already face waits of up to seven years (Deloitte, 2025), while CBRE's UK market review highlights limited power availability across London (CBRE, 2025). These examples suggest local substations, feeders and distribution circuits may become binding constraints even before broader generation adequacy is exhausted.

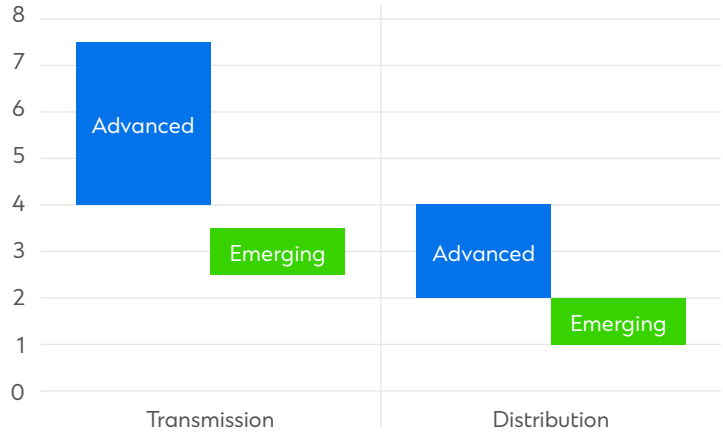
Could grid constraints benefit some developing markets?

While emerging economies often have less mature electricity networks, the lead times for new grid infrastructure can be materially shorter than in advanced economies — in some cases around half as long (Figure 32). This matters for data centre developers because the speed at which power can be made available is becoming as important as land, fibre connectivity and construction costs. Markets that can expand generation, transmission and distribution capacity more quickly may therefore become more attractive locations for incremental data centre capacity, even if their existing grid infrastructure is less developed. Conversely, advanced economies with longer permitting timelines, congested connection queues and slower transmission build-out may risk losing growth to markets where grid upgrades can be delivered faster.

Figure 32

Grid infrastructure lead times

Years

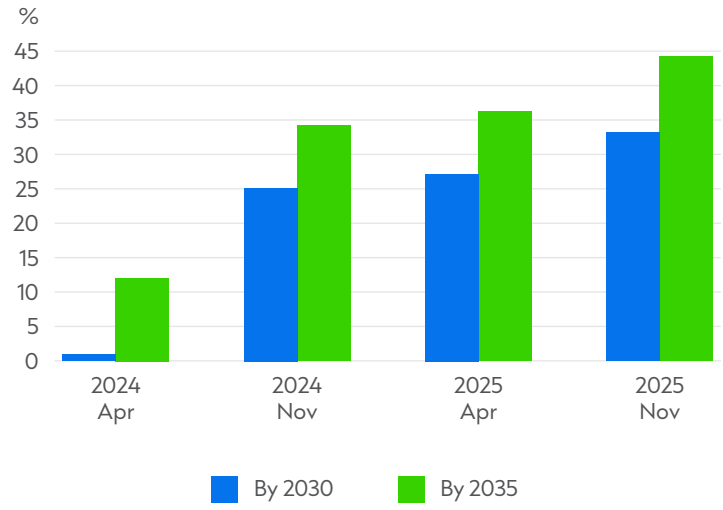


Source: IEA, Standard Chartered Bank

Figure 33

Popularity of onsite power generation is rising

% of data centres with 100% onsite power generation



Source: (Bloom Energy, 2026), Standard Chartered Bank

The power challenge is about more than data centres

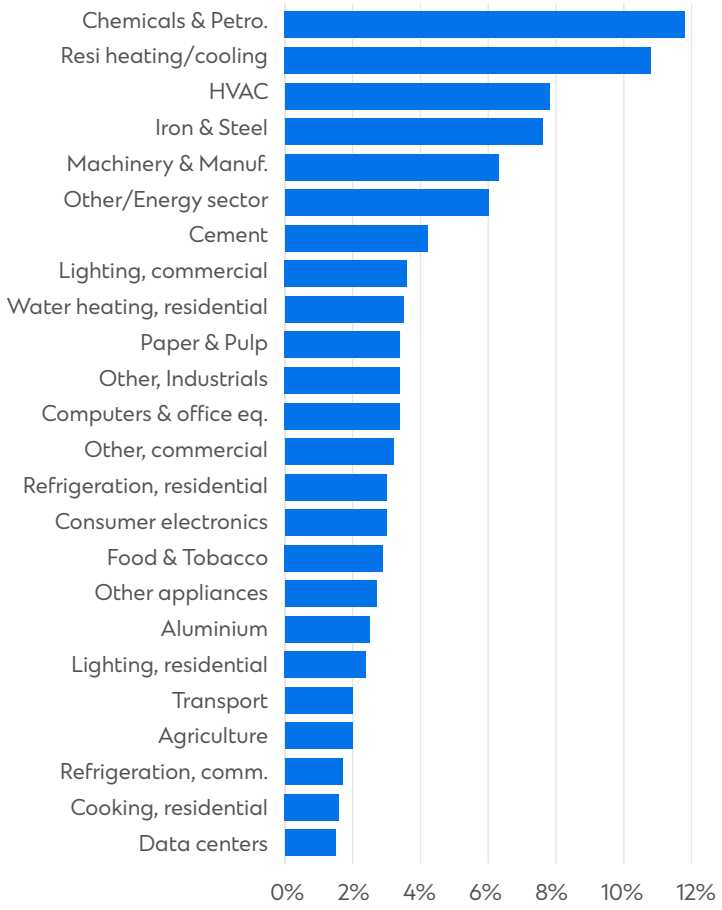
Despite the rapid growth of the data centre market, it is important to keep the sector's electricity use in perspective. In 2023, data centres accounted for around 1.5% of global electricity consumption — still well below many major end-use sectors (Global Electricity, 2026). By comparison, the chemicals and petrochemicals sector represented almost 12% of global electricity use, while total residential and household-related demand accounted for 27% (Figure 34). This suggests that, although data centre power demand is rising quickly and can create acute local grid pressures, the sector remains a relatively modest contributor to global electricity consumption today.

Future power supply challenges should be viewed in the context of economy-wide electrification, not data centres in isolation. As transport, industry, buildings and other sectors increasingly switch to electricity, competition for power is set to intensify across the economy. Electric vehicles (EVs) are a useful example. Estimates from the (IEA, 2026) and (BloombergNEF, 2025) suggest that EV-related electricity demand could grow 2.1-3.7 times faster than data centre electricity demand and ultimately exceed the sector's total power consumption (Figure 35). This highlights that while data centres are an important driver of future electricity demand growth, they are only one part of a much broader electrification trend that will shape future power-system investment, grid expansion and energy security requirements.

Figure 34

Share of global electricity consumption

2023

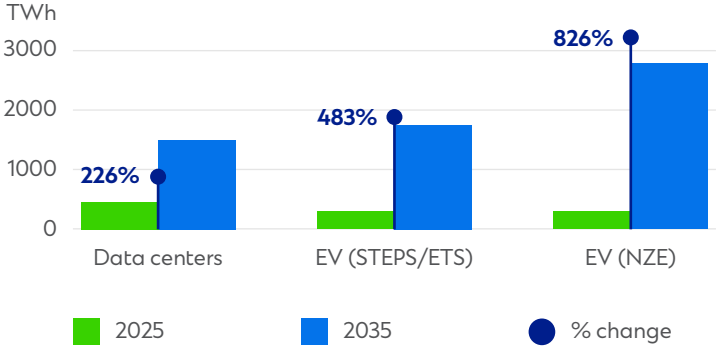


Source: (Global Electricity, 2026), Standard Chartered Bank

Figure 35

Electricity consumption: data centres vs EV

STEPS = Stated Policy Scenario, ETS = Economic Transition Scenario, NZE = Net Zero Energy Emissions scenario



Source: BloombergNEF, IEA, Standard Chartered Bank

Solutions to further reduce data centre power consumption

A data centre's power consumption is primarily driven by its IT equipment including servers, storage systems and networking hardware and by the systems needed to cool this. Additional power is used for back-up generation systems and supporting infrastructure such as lighting or general office equipment. Estimates from the IEA suggest that approximately 70% of data centre power use is driven by its IT equipment (IEA, 2025). The remainder is mostly related to powering cooling equipment (Figure 36).

Power Usage Effectiveness (PUE) ratio: not a long-term useful measure

The mix of power consumption by data centres shown in Figure 36 suggests that to reduce the pressure that data centre growth puts on power systems requires improvements not just in how much power is needed to run IT equipment but also in the power used by cooling and supportive infrastructure.

The Power Usage Effectiveness (PUE) of a data centre is an often-used ratio to measure the power efficiency of a data centre. It is defined as the ratio of the total energy used by a data centre to the amount that was needed to power the IT equipment (Lei & Masanet, 2022). PUE data shows that the data centre sector has been able to significantly reduce the average PUE during the past 10 years (Figure 37). The average power effectiveness for US data centres has improved by 15% between 2015 and 2024. Reasons for this improvement include better air management, cooling optimisation and power chain improvements (Smith, Hubbard, & et al, 2026). Company-level data suggests that hyperscalers achieve structurally lower PUEs than other data centre operators (Figures 38). Reasons for this may be greater scale, having purpose-built facilities, higher utilisation of their IT equipment, the use of more advanced cooling solutions, proprietary more efficient hardware and deployment of AI-enabled optimisation.

Although the PUE ratio is widely used by the data centre sector and policymakers as a measure of a data centre's efficiency, there are reasons why it may not be the best ratio for this.

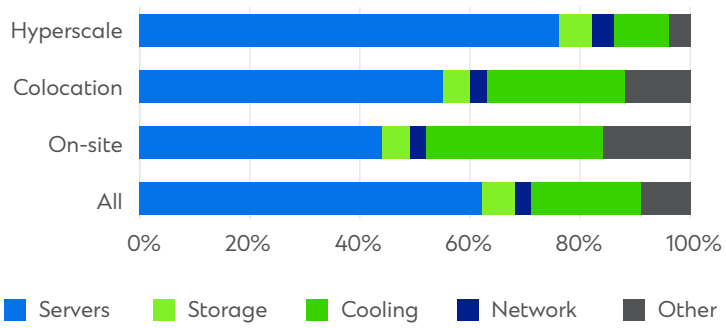
- The PUE ratio only shows how much power a data centre needs for cooling and other non-IT equipment relative to what is needed to run its IT systems. It provides no standalone insight into the efficiency of its IT equipment and services or its cooling equipment.
- As shown elsewhere in this report, solutions exist that reduce the amount of power needed for cooling. As these solutions are adopted, the PUE ratio will organically reduce closer to its minimum of 1 by which time the ratio provides stakeholders with no insight into the power efficiency of a data centre at all.
- The PUE ratio is ignorant to the type of power and the carbon content of this that a data centre uses. It therefore is not a good indicator of the environmental power efficiency of a data centre.

Various strategies and solutions are available that can help to further reduce the overall power intensity of data centres and ease the electricity-supply challenges that future capacity growth may create.

Figure 36

Share of data centre electricity consumption by type and equipment

2024

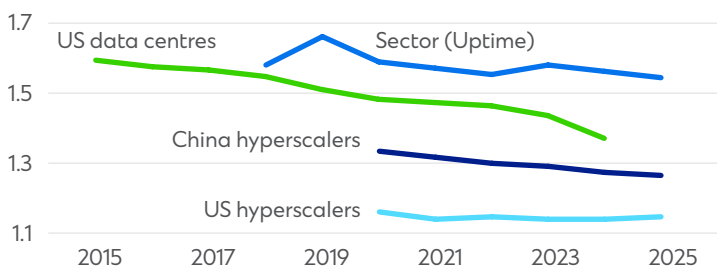


Source: (IEA, 2025), Standard Chartered Bank

Figure 37

Power Usage Effectiveness of data centres

(x)



Source: Companies mentioned, Standard Chartered Bank.
* 2024, others 2025

Solution 1

Improve operational efficiency

Data centre operators together with IT hardware and software suppliers are aggressively pursuing various strategies to help minimise IT-related power requirements.

1. Software efficiency gains are a low-cost lever

Improving model architecture and algorithm efficiency can reduce the compute required to deliver the same level of AI performance without the need for large-scale physical infrastructure investment. Evidence of progress is significant: analysis of more than 200 language models between 2012 and 2023 suggests that the compute needed to reach a given performance threshold has halved roughly every eight months (Ho & et al, 2024). More recent models also show how quickly efficiency can improve, with some achieving comparable performance to older models while using less than 1% of the parameters (Maslej & et al, 2025). Data from NVIDIA also shows the benefits of IT and software optimisation as the company managed to quadruple the output of its GB200 NVL72 supercomputing system between October 2025 and January 2026 by delivering 38 optimisations (Figure 39).

Software gains help lower data centre electricity consumption growth. However, they may also trigger an unwanted negative feedback loop that could reduce these gains. As AI becomes cheaper, faster and easier to deploy, demand for it could rise materially creating a Jevons paradox where higher usage offsets part, or even all, of the per-query efficiency gains (Jegham & et al, 2025).

2. Improved hardware efficiency further reduces power demand growth

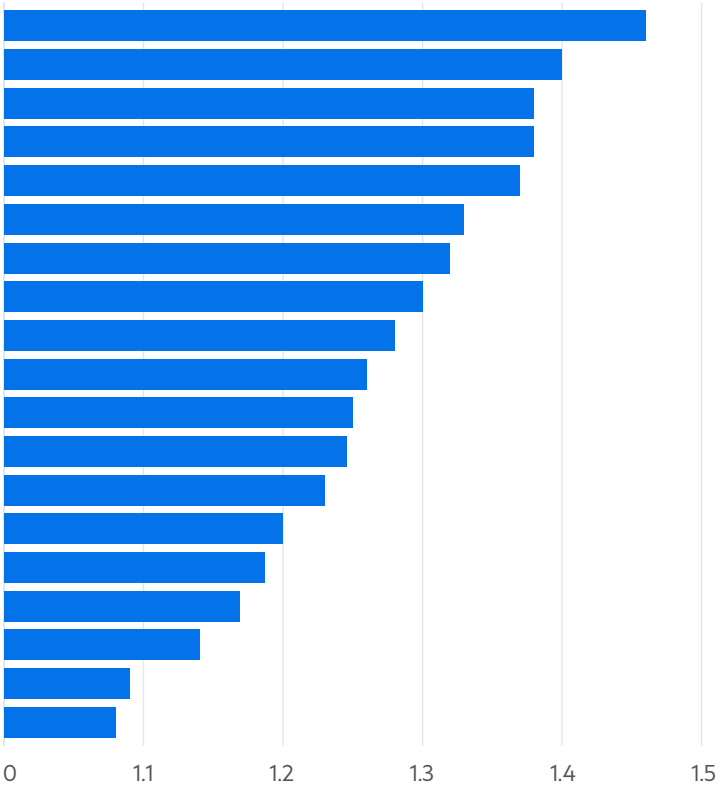
Hardware efficiency is another powerful lever for slowing data centre electricity growth because it increases the amount of compute delivered per watt. Ratios that provide a good indication of hardware efficiency include Performance per watt, Tokens per watt or Floating-point operations per second per watt.

Ratios such as these show that strong hardware efficiency gains have been achieved during the past ten years as the average machine-learning hardware developer has managed to increase

Figure 38

PUE for key data centre companies

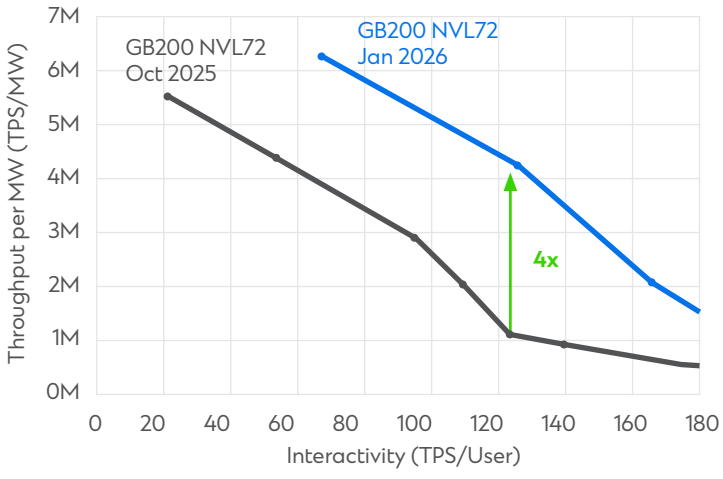
(x)



Source: Sustainability reports of companies mentioned, Standard Chartered Bank.
* 2024, others 2025

Figure 39

Optimisation of NVIDIA's Blackwell



Source: NVIDIA

the number of operations per watt almost 100-fold (Figure 40). Results from some of the leading AI manufacturers are even more impressive. The efficiency of NVIDIA's latest Graphical Processing Unit (GPU) "Rubin" as measured by tokens per megajoule is more than 1 million times greater than what its "Kepler" GPU managed to achieve in 2012 (Figure 41).

3. Increase the workload of data centre hardware

Optimising hardware utilisation is another critical lever to help reduce electricity supply-demand challenges for those operators that have control over data centre workloads and therefore over capacity utilisation. Low server utilisation means data centres continue to consume power even when computing resources are underused. Virtualisation, workload consolidation, right-sizing and intelligent workload scheduling can raise utilisation rates and reduce the need for redundant physical servers. ENERGY STAR notes that virtualisation can consolidate multiple workloads onto fewer physical servers (Energy Star, accessed 2026), while Uptime Institute analysis suggests that doubling processor utilisation from low levels can materially improve server efficiency and that power management can add a further 10-15% efficiency gain (Uptime Institute, 2023). Work from Ambienta (Ranghino, 2024) reinforces this point by showing that increased IT hardware efficiency and higher capacity utilisation can generate significant gains in data centre computing efficiency (Figure 42).

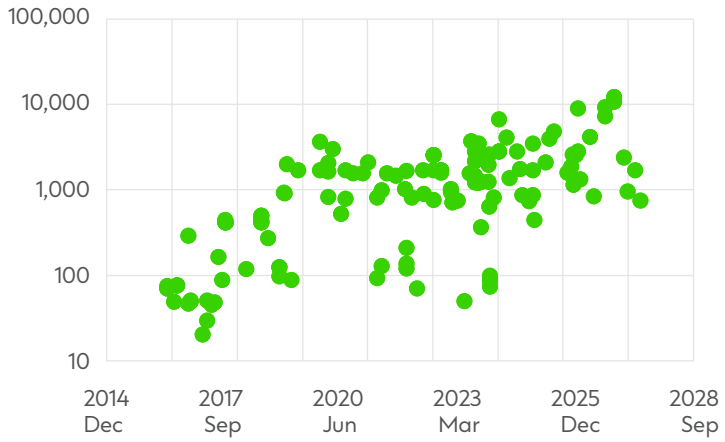
4. Thermal management of data centres: Industry collaboration

The rising power density of AI data centres is increasing thermal management challenges. To reduce associated power and cooling requirements, companies across the data centre value chain are increasingly collaborating. One example is the partnership between NVIDIA and Trane Technologies, which late last year led to the launch of the industry's first comprehensive thermal-management reference design for gigawatt-scale AI data centres using NVIDIA's latest AI infrastructure (Trane Technologies, 2025). The design aims to optimise power, water and land use across high-density AI facilities. Building on this foundation, Trane has since introduced enhanced designs for large-scale AI deployment, including a solution that supports free cooling and integrated heat recovery. Trane says this can improve thermal-management efficiency by 14%, with 10% of the heat-rejection load redirected to heat recovery (Trane Technologies, 2026). These types of partnerships suggest that closer cooperation across the data centre value chain can improve operational efficiency and help reduce both power and water intensity.

Figure 40

Energy efficiency of machine learning hardware

(GFLOPS* per watt)

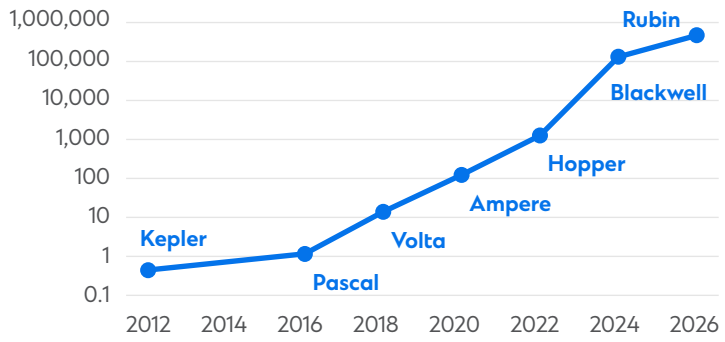


Source: EPOCH AI, Standard Chartered Bank
* = Billion floating point operations per second

Figure 41

Data centre efficiency improvement

Tokens per Megajoule

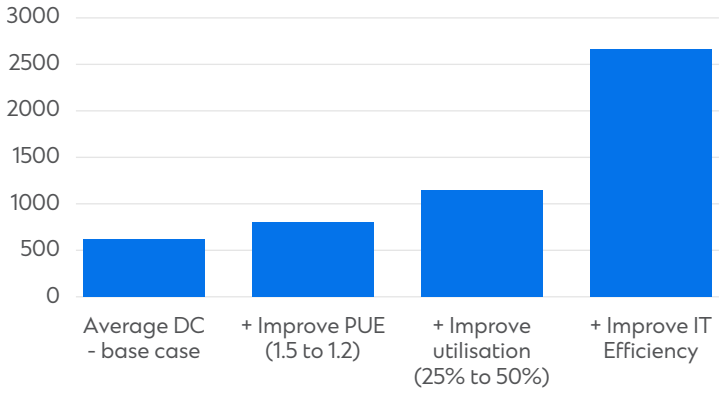


Source: NVIDIA

Figure 42

Data centre efficiency improvement

Transactions per second per watt



Source: Ambienta, Standard Chartered Bank

Solution 2

Reuse waste heat to lower broader energy demand

Improving equipment efficiency and operational performance are not the only ways data centre operators can reduce their overall energy needs and environmental footprint. Another increasingly relevant solution is heat reuse. This captures the waste heat generated by data centre equipment and uses it as a source of useful thermal energy.

The benefits of reusing a data centre's heat

Heat reuse involves capturing the heat produced by data centre equipment and redirecting it for other purposes, such as heating homes, commercial buildings or industrial facilities. Historically, most data centre heat has been rejected into the atmosphere through cooling systems. Reusing this heat can reduce overall energy demand by displacing separate heating fuels, such as natural gas. In effect, the electricity used to power digital infrastructure can also help meet local heating needs.

The benefits are greatest when heat reuse is combined with efficient liquid cooling and heat pumps. In these systems, recovered heat can be upgraded to a temperature suitable for district heating, while the heat pump process can also support data centre cooling. This can reduce cooling-related power consumption, thereby improving PUE and potentially water use intensity too and lower emissions where recovered heat replaces fossil-fuel-based heating. The emissions benefit is strongest when the data centre itself is powered by low-carbon electricity.

Case study: Trane Technologies

Trane Technologies and data centres

Trane Technologies is one of the global leading providers of thermal management systems, solutions and services and has been rapidly expanding its capabilities to serve this market. Earlier this year, Trane Technologies acquired LiquidStack, a global leader in liquid cooling technology for data centres while late last year it acquired Stellar Energy Digital Business, a provider of turnkey modular data centre cooling solutions.

Trane heat pumps support the reuse of data centre waste heat

Heat pumps form a key component of a data centre heat reuse system. Two of Trane's so-called water-to-water heat pumps were used as part of Swiss cloud and internet services company Infomaniak's fully underground D4 data centre in Geneva, Switzerland (D4 project). The heat reuse set up allows up to 100% of waste heat generated by Infomaniak's data centre to be converted and reused.

The heat generated by the data centre's equipment is captured at 40-45°C and then increased by heat pumps to 67°C in the summer and up to 85°C in the winter after which it is fed into the local district heating network. In parallel, the refrigerant expansion in the heat pumps cools a second water flow to about 28°C which is then used to cool the servers. This approach therefore allows not only to reuse 100% of the waste heat generated by a data centre but the cooling capabilities of the heat pumps also mean that the facility does not require additional water or air conditioning.

The heat reuse process installed for Infomaniak's data centre helps to avoid the combustion of 3,600 tCO2eq of natural gas per year and heat 6,000 households per year (Infomaniak, 2025). It also allowed the data centre to achieve a PUE of 1.1 which is among the lowest reported for data centres globally. The Trane/Infomaniak collaboration therefore demonstrates how heat reuse can become an additional efficiency lever: it can reduce the net energy demand of local heating systems, lower emissions and improve the overall environmental profile of high-density data centre operations.

Energy Reuse Factor (ERF)

While IT efficiency can be measured using tokens per watt or floating-point operations per second per watt, the extent to which a data centre captures and reuses waste heat to reduce power requirements can be measured by the Energy Reuse Factor (ERF). This ratio represents the proportion of a data centre's total power consumption that is reused.

Solution 3

Expand exposure to renewables further

The potential growth in electricity demand from data centres coupled with the earlier-mentioned power access challenges are reasons why an increasing number of data centre operators are developing alternative power sources to ensure consistent and predictable power supply and minimise the impact on power supply or prices for local communities and businesses. Hyperscalers in particular take a portfolio approach to power supply and often base decisions on the mix of power sources on reliability, grid challenges, carbon and water implications and local community impact assessments. Sources that data centre operators have started to pivot towards include renewable power, the development of onsite or 'behind the meter' power generation, nuclear power and geothermal capabilities.

Expansion of gas fired generation: Environmental impact may be lower than expected

One source of potential power supply that has received significant media attention is gas power plants. Recent data from Global Energy Monitor (GEM) suggests that the amount of gas-fired electricity capacity for data centres that has been announced or is in pre-construction or construction phase of development doubled between 2025 and the first half of 2026 to 189GW (GEM, 2026).

The potential emission impact of this could be significant

If all 189GW were genuinely incremental, fully built and operated at an 80% capacity factor, the emissions impact could be substantial. Natural gas without carbon capture has an average lifecycle emissions intensity of around 490gCO2e/kWh. This is well below coal at around 820gCO2e/kWh, but materially above nuclear, wind, solar and hydropower at around 11-48gCO2e/kWh (World Nuclear Association, 2026).

Applying EIA heat-rate and combustion-factor estimates suggests annual emissions could exceed 600MtCO2, equivalent to roughly 40% of US electric-power-sector CO2 emissions in 2025 (EIA, 2025) (EIA, 2024) (EIA, 2026).

The actual build-out of gas power plants is likely to be lower than suggested

Several factors suggest that only a portion of the announced pipeline will become operational.

- **Demand may be overstated:** The 189GW total includes projects at announced, pre-construction and construction stages, many of which remain speculative or insufficiently defined. More than half reportedly lack a named turbine or engine manufacturer, while almost one quarter have no stated start year. The figure may also contain duplicated or inflated demand where developers seek capacity from multiple utilities, reserve more power than they ultimately use, or advance data centre projects that are delayed, downsized or never built.
- **A lack of equipment:** Gas-turbine order books are already constrained, with lead times of more than three years cited by manufacturers such as Siemens (Utility Dive, 2026). Developers may turn to modular engines rather than combined cycle turbines, but doing so reduces efficiency, increases local emissions and may complicate permitting processes.
- **Gas supply could be a constraint:** A full buildout of the 189GW could require 30 billion cubic feet of gas per day if this were run at 80% capacity. This would suggest that this incremental data centre gas power capacity would consume c. 25% of the US' current total natural gas production which would therefore not be available for other domestic demand source or exports (EIA, 2026b). Even if enough gas can be procured to power these plants it would also require gas pipelines to deliver it. This itself can be challenging owing to approval, right-of-way and construction struggles.
- **Economics could become a challenge:** Building 189GW of power plants can cost almost USD650 billion according to Global Energy Monitor (GEM) (American Public Power Association, 2026). This is sufficiently high to ensure that data centre operators will have to weigh the possibility that other power sources including a renewable with storage combination is not cheaper. Another key financial risk is that a strong increase in demand for gas to power these plants will drive up the price of gas which would further undermine the economics of investing in gas power plant capacity relative to other power sources.

- **Permitting and general opposition:** Finally, it is relevant to point out that developing new gas power plants at scale may result in public and political opposition as the environmental or financial implications become clear. The risk of permitting delays can therefore also not be discounted.
- **Continued improvement in capabilities of alternative sources:** The intermittency of renewable power is a key reason why data centre operators need to explore other sources of electricity. Further improvements in battery technologies and storage investments will help address this intermittency problem and likely reduce investment activity in less environmentally-friendly power sources such as gas power plants.

Focus on expanding renewable energy exposure likely to continue

Given the challenges described above, the amount of incremental gas-fired capacity ultimately built to serve data centres is likely to be lower than the announced pipeline.

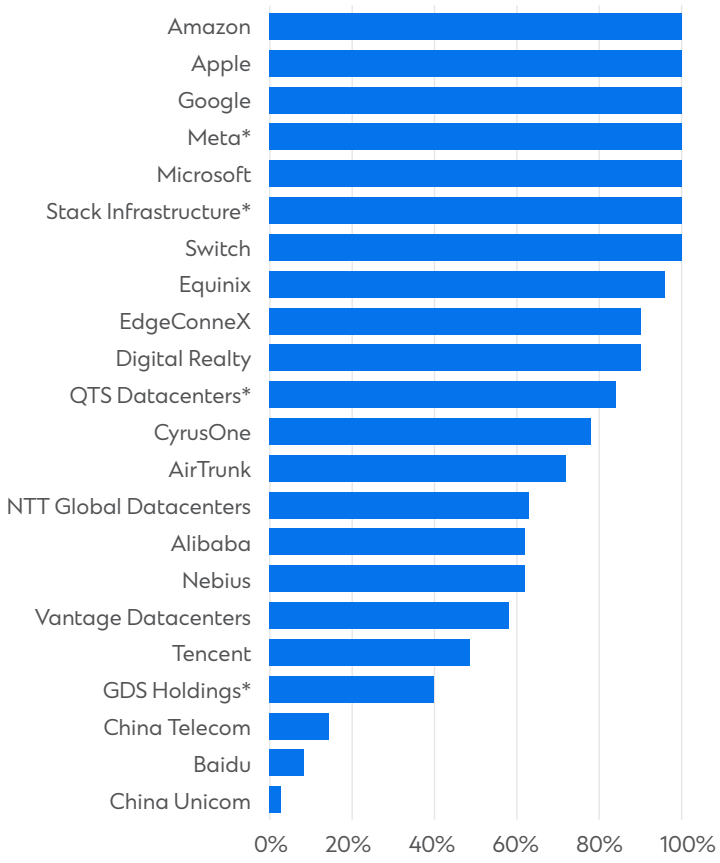
Company disclosures for some of the largest data centre operators suggest that contracted renewable electricity already covers well over half of their total power consumption while in some cases it appears to match or exceed their entire reported electricity use (Figure 43).

The relationship also appears to strengthen with scale: plotting renewable power coverage against data centre capacity suggests that larger operators are more active in renewable procurement (Figure 44). This is logical, as the biggest platforms face the greatest challenge in securing large volumes of reliable power while simultaneously scaling energy-intensive AI workloads. For these operators, renewable power access is becoming a competitive necessity — helping reduce emissions exposure, strengthen long-term power security and support continued capacity growth in increasingly constrained electricity markets.

Figure 43

Renewable power as % of electricity usage

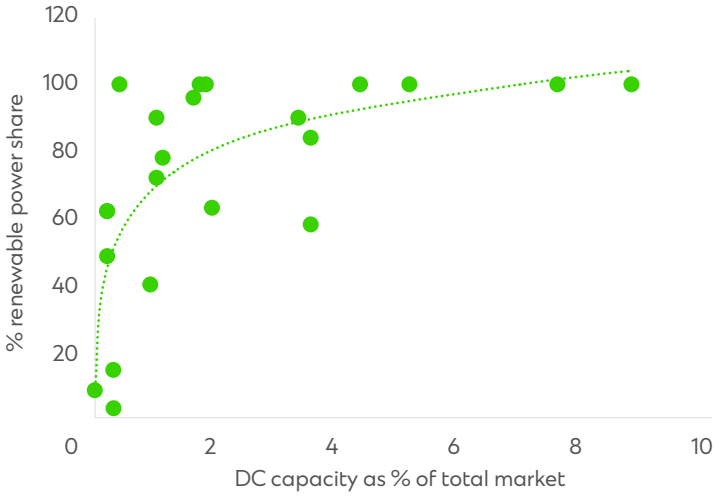
2025, * = 2024



Source: Sustainability reports for companies mentioned, Standard Chartered Bank

Figure 44

Renewable coverage rises with DC size



Source: Data centre companies, Bloomberg, Standard Chartered Bank

Data centre renewable energy procurement strategies

Data centres can increase their exposure to clean energy either through the development of onsite generation capacity or through external renewable power projects. At present on-site generation is very low for most data centre companies, however, this may change significantly. Bloom Energy survey work shows that around 35% of data centre operators expect to have 100% onsite power generation by 2030, up from just 1% in April 2024. Exposure to renewable energy from external projects is mostly done through the following structures:

Power purchase agreements (PPAs) are most popular with DC operators

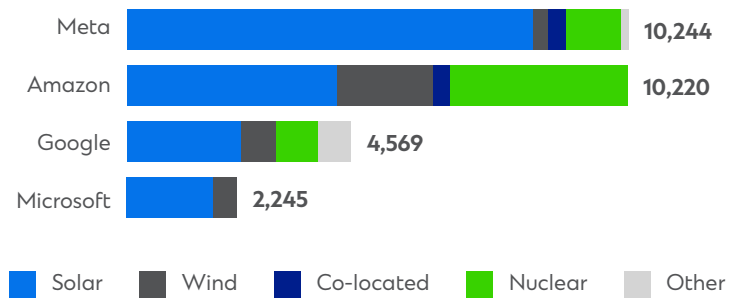
PPAs are long-term contracts under which a data centre or corporate buyer purchases electricity directly from a renewable project. PPAs are the main vehicle through which data centres procure their energy. Company filings show that Amazon, Google, Meta and Microsoft combined have contracted more than 140GW of renewable power to date and almost 35GW in 2025 alone (Figure 45). The majority of contracted renewable power from these companies was done via PPAs.

Data centre companies that use clean energy contracting through PPAs may encounter issues that impact their emission goals if the electricity that is contracted is not generated close to the data centres or if it is not generated at the time when the data centre needs it. In both cases data centre companies would use electricity from their local grid which may be produced using fossil fuel and therefore have higher carbon emission intensities. To avoid these issues some of the larger data centre companies such as Microsoft and Google are moving to hourly matching of energy demand with low-emissions electricity. In its 2025 environmental report Google showed that most of its owned and operated data centres already match energy usage with carbon-free energy for more than 60% of the time (Figure 46).

Figure 45

Corporate clean energy purchase

MW, 2025

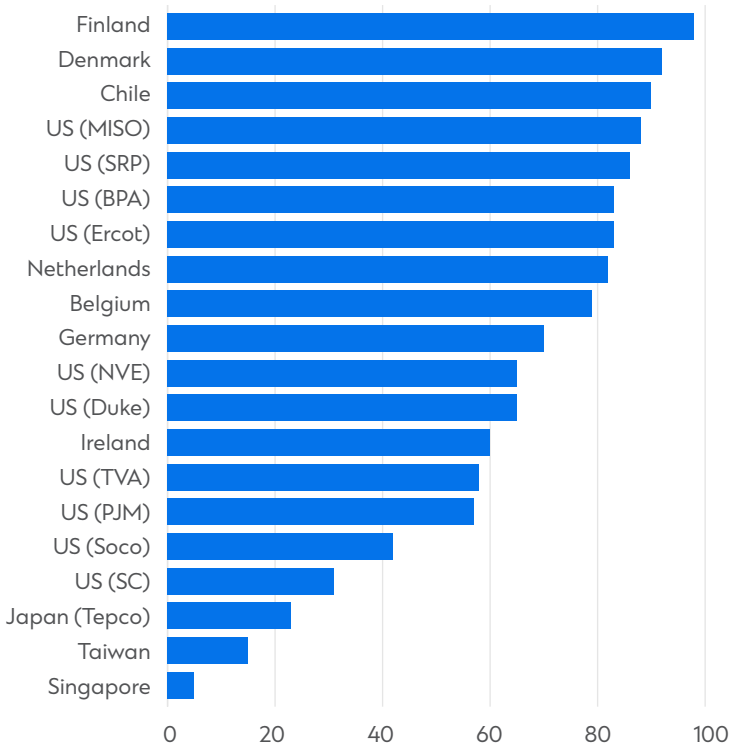


Source: BloombergNEF, Standard Chartered Bank

Figure 46

Google data centres

% hourly matched with carbon-free energy, 2025



Source: Google, Standard Chartered Bank

Figure 47

Comparison between PPAs and EACs

Solution	Data centre receives	Physical electricity delivered?	Main benefits	Main challenges	Best suited for
Physical PPA	Electricity plus associated EACs from specific renewable projects	Yes	Strong project-specific link; long-term price visibility; supports new renewable capacity; strong procurement narrative	Complex to negotiate; market/geographic constraints; may require deregulated or direct-access markets	Large data centres with concentrated load in markets where direct renewable supply is feasible
EACs / RECs / GOs	Renewable attribute certificates only, usually one certificate per MWh	No	Simple, flexible, quick to procure; useful for residual emissions coverage; works in markets where PPAs are unavailable	Price hedging; additionality and double counting concerns if unbundled; quality varies by vintage, geography and certification rules	Covering residual load, early-stage procurement, or markets with limited PPA availability

Source: Standard Chartered Bank

Energy Attribute Certificates (EACs)

EACs, which are called Renewable Energy Certificates (RECs) in North America and Guarantees of Origin (GOO) in Europe, are tradable certificates representing the environmental attributes of one MWh of renewable electricity generation. EACs may create uncertainty over perceived additionality if they are unbundled from new project development. Buying EACs without actual energy supply might allow data centres to offset non-renewable energy use but it does not help to meet growing energy demand. EACs might also create a risk of double counting in the absence of strong regulatory frameworks and monitoring solutions. Figure 47 highlights some of the differences between PPAs and EACs.

Data centre operators help build the next generation of renewable energy.

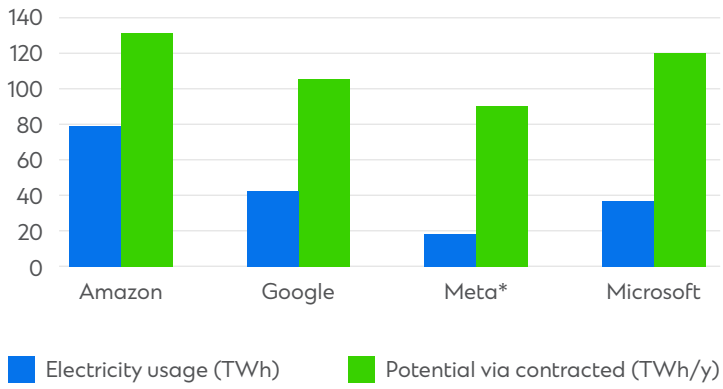
The renewable power capacity developed or supported by hyperscalers appears capable of generating materially more electricity than they currently consume (Figure 48). Amazon's 2025 sustainability report states that its contracted clean-energy projects generated 131TWh or around 70% more than the company's electricity consumption. Based on Amazon's contracted capacity, this is equivalent to approximately 3.1TWh of annual generation per gigawatt. This estimate is broadly consistent with the c3.3TWh that 1GW of a representative mix of solar, onshore wind and offshore wind could generate each year, based on typical capacity factors of 25%, 38% and 50%, respectively (Center for sustainable systems (a), 2026), (Center for sustainable systems, 2026). Applying a ratio of 3TWh per GW to the other hyperscalers suggests that they may also have secured enough renewable generation capacity to cover their current electricity consumption several times over (Figure 48).

Although these ratios should be interpreted carefully given timing differences and the use of 2024 consumption data, for some companies the data does show that data centres are moving beyond annual renewable matching toward strategic energy procurement at scale. They are positioning themselves less as passive electricity consumers and more as anchor buyers and in some cases effective builders of clean energy infrastructure. This strategy not only helps contribute to the global energy transition but may also help reduce some of the power supply challenges that a further built out of data centre capacity may create.

Figure 48

Hyperscalers might be ready for the growth in electricity demand

TWh, 2025 * = 2024



Source: Sustainability reports for companies mentioned, Standard Chartered Bank

Figure 49

Climate and renewable energy targets for data centre operators and industry body engagement

Company	Carbon neutral (Sc. 1/2)	Year	Net zero (Sc 1/2)	Year	100% renewable target	Year
AirTrunk			✔	2030	✔	2030
Alibaba	✔	2030			nd	
Aligned	nd				✔	nd
Amazon			✔	2040	✔	Achieved
Apple	✔	2030			✔	2030
Baidu	✔	2030			nd	
Beale Infrastructure	nd				✔	nd
China Telecom	✔	2060			nd	
China Unicom	✔	2060			nd	
CloudHQ			✔	2040	nd	
CyrusOne	✔	2030			✔	Achieved
DayOne			✔	2040	✔	2030
Digital Realty	-42% vs '23				✔	
EdgeConneX	-50% vs '21				✔	2030
Equinix			✔	2040	✔	2030
GDS Holdings	✔	2030			✔	2030
Google			✔	2030	✔	24/7 by 2030
Meta			✔	Achieved	✔	Achieved
Microsoft			✔	Negative by 2030	✔	24/7 by 2030 on same grid
Nebius						
NTT Global Data Centers			✔	2040	✔	2030
QTS Data Centers	nd				nd	
Rowan Digital Infrastructure	nd				nd	
Stack Infrastructure			✔	2050	✔	Achieved
Switch			✔	Achieved	✔	Achieved
TECfusions	nd				nd	
Tencent	✔	2030			✔	2030
Vantage Data Centers			✔	2030	✔	2030

Company	RE100	24/7 Carbon-Free Energy (CFE)	Clean Energy Buyers Association	Climate Neutral Data Centre Pact	The iMasons Climate Accord	Coalition for Sustainable AI	EPA s Green Power Partnership
AirTrunk	✔	✔					
Alibaba							
Aligned			✔				✔
Amazon				✔	✔		
Apple	✔						
Baidu						✔	
Beale Infrastructure			✔		✔		
China Telecom							
China Unicom							
CloudHQ							
CyrusOne			✔	✔			
DayOne							
Digital Realty			✔	✔	✔		✔
EdgeConneX	✔		✔	✔	✔		
Equinix			✔	✔	✔	✔	✔
GDS Holdings							
Google	✔	✔		✔	✔		✔
Meta				✔	✔		
Microsoft	✔	✔		✔	✔		✔
Nebius							
NTT Global Data Centers				✔			
QTS Data Centers	✔		✔				✔
Rowan Digital Infrastructure			✔				
Stack Infrastructure			✔	✔	✔		
Switch		✔		✔	✔	✔	✔
TECfusions							
Tencent							
Vantage Data Centers			✔	✔	✔		

Source: Sustainability reports for companies mentioned and websites of organisations highlighted, Standard Chartered Bank, nd = not disclosed

Data centres lead on engagement with renewable energy

The strong commitment to renewable energy among data centre operators, particularly hyperscalers, reflects a broader strategic effort to align long-term growth with ambitious climate objectives. Some data centre operators have recently suggested that their climate targets had become harder to achieve because of the rapid growth in AI (Google, 2026), Nevertheless a review of climate and energy commitments across the sector shows that most leading operators aim to meet or match 100% of their electricity consumption with renewable energy by 2030. Many have also adopted ambitious decarbonisation targets, with carbon-neutrality goals for Scope 1 and 2 emissions typically targeted for 2030 and broader net-zero commitments extending into the 2030-2040 period (Figure 49).

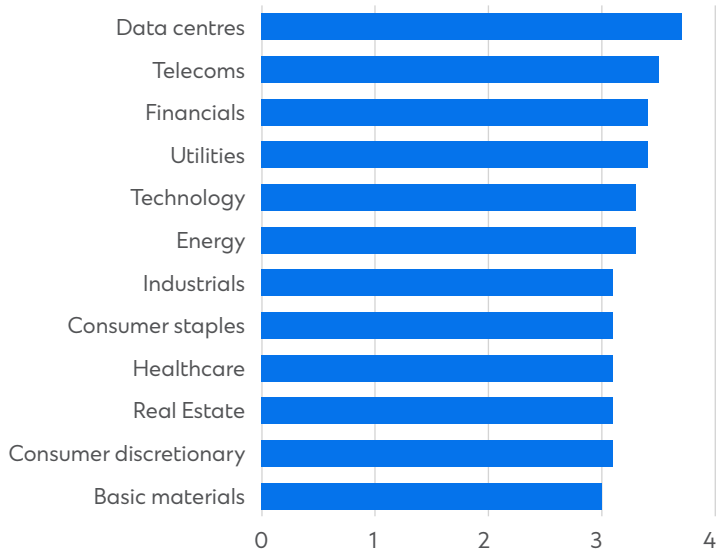
Beyond formal targets, the sector's sustainability credentials are reinforced by their widespread participation in various climate, renewable energy and clean-energy procurement initiatives as also shown in figure 49. Taken together, these commitments suggest that leading data centre operators increasingly view access to clean energy not only as a sustainability objective, but as a strategic prerequisite for securing future growth, reducing emissions and strengthening long-term energy resilience.

The Transition Pathway Initiative (TPI) assesses the extent to which a company's management team has integrated climate objectives and strategies into their operations (Transition Pathway Initiative, 2026). The TPI ranks the quality of a company's transition strategy from 0 "Unaware if climate change as a business issue" to 5 "Transition planning and implementation is established". The ten data centre operators that have been assessed by TPI received an average score of 3.7 (Figure 50). This is higher than the average score achieved by each of the 11 major industries covered by TPI, implying that data centre operators have integrated climate and transition-related policies more centrally into their business operations than most companies covered by TPI.

Figure 50

Average management quality score for transition pathway by industry

Ranging from 0 (weakest) to 5 (strongest)



Source: TPI, Standard Chartered Bank



A comparison of the energy intensity of different industries and data centre operators also suggests that US hyperscalers have adopted renewable energy at a scale that appears to exceed most other US sectors (Figure 51). While these large data centre operators exhibit relatively high gross electricity intensity when measured as electricity consumption per dollar of revenue, their extensive renewable energy procurement has significantly reduced the emissions associated with that electricity use. Adjusting for renewable energy exposure shows that hyperscalers have among the lowest net emissions-adjusted electricity intensities of any sector. This highlights how scale can create a sustainability advantage.

Hyperscalers also focus on carbon removal to reduce their emissions

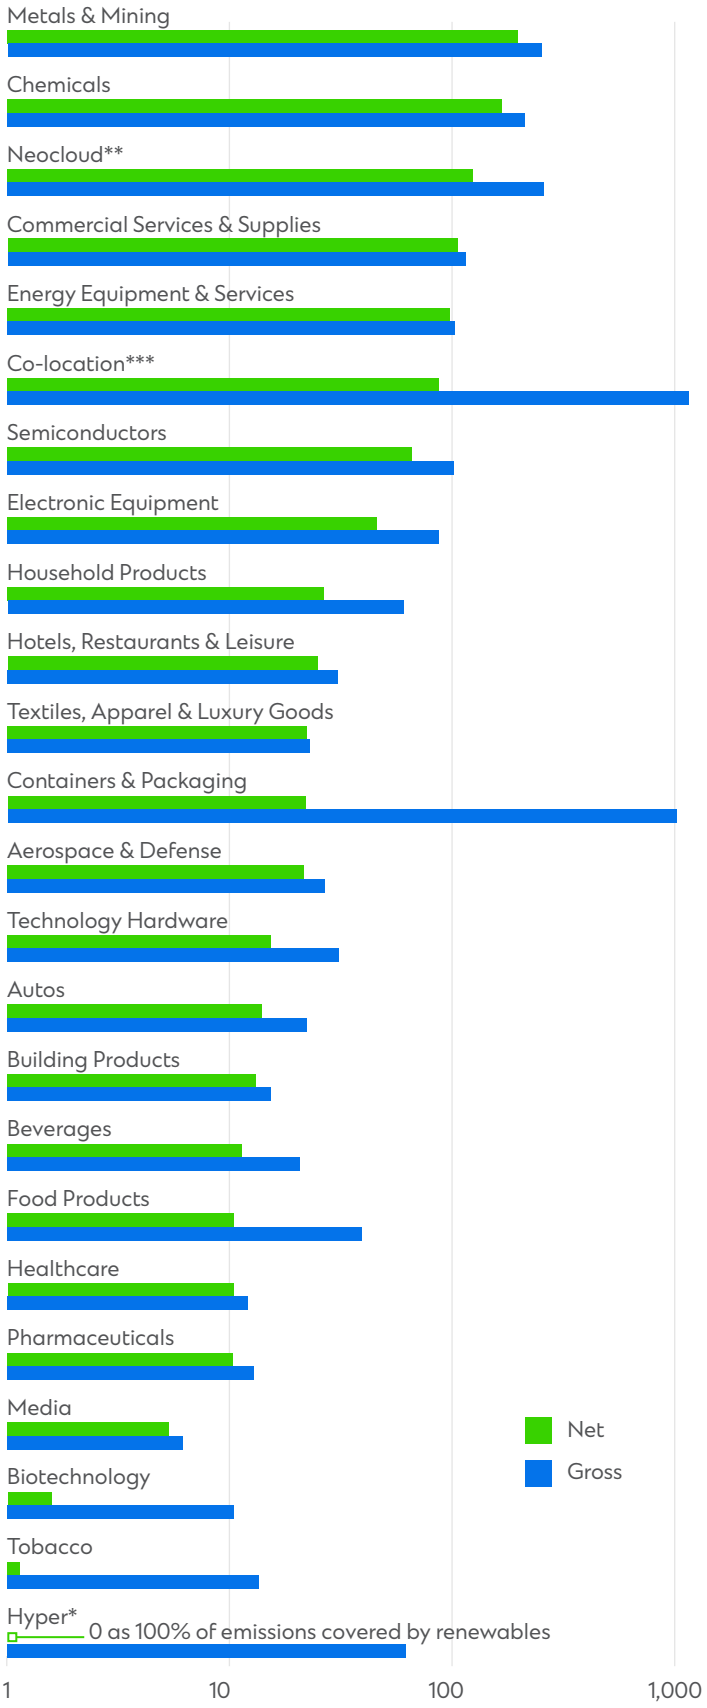
As part of their broader decarbonisation strategies the leading hyperscalers are complementing renewable energy procurement with direct investment in carbon dioxide removal (CDR), recognising that residual emissions cannot be eliminated through operational efficiency and clean energy alone. Microsoft, Google, Meta and Amazon have all established carbon removal strategies spanning technologies such as direct air capture, biochar, enhanced weathering, BECCS and high-integrity nature-based solutions.

These commitments strengthen the credibility of their sustainability strategies by demonstrating a willingness to address hard-to-abate emissions while helping to scale the wider carbon removal market. As hyperscalers continue to expand rapidly to support AI workloads, their investments in removals provide an important signal that long-term decarbonisation will require not only low-carbon energy and efficient data centres, but also permanent pathways for removing carbon already present in the atmosphere.

Figure 51

Electricity intensity comparison for US sectors

kWh/USD revenue, 2025, Gross is total electricity usage, Net is after deducting renewable-based electricity



Source: BloombergNEF, Standard Chartered Bank
* = Amazon, Microsoft, Meta, Google, ** Nebius Group, *** - Equinix, Digital Realty

Solution 4

Move to 800-volt DC distribution to reduce power loss

Facilitating the growth in AI-related computing requires not only a more powerful and efficient IT infrastructure but also a more efficient way to carry power from the grid to the servers. At present, electricity is delivered mostly as an alternating current (AC) at between 400 and 480 volts. Before it reaches the server, that power is converted multiple times, adding complexity and creating avoidable energy losses.

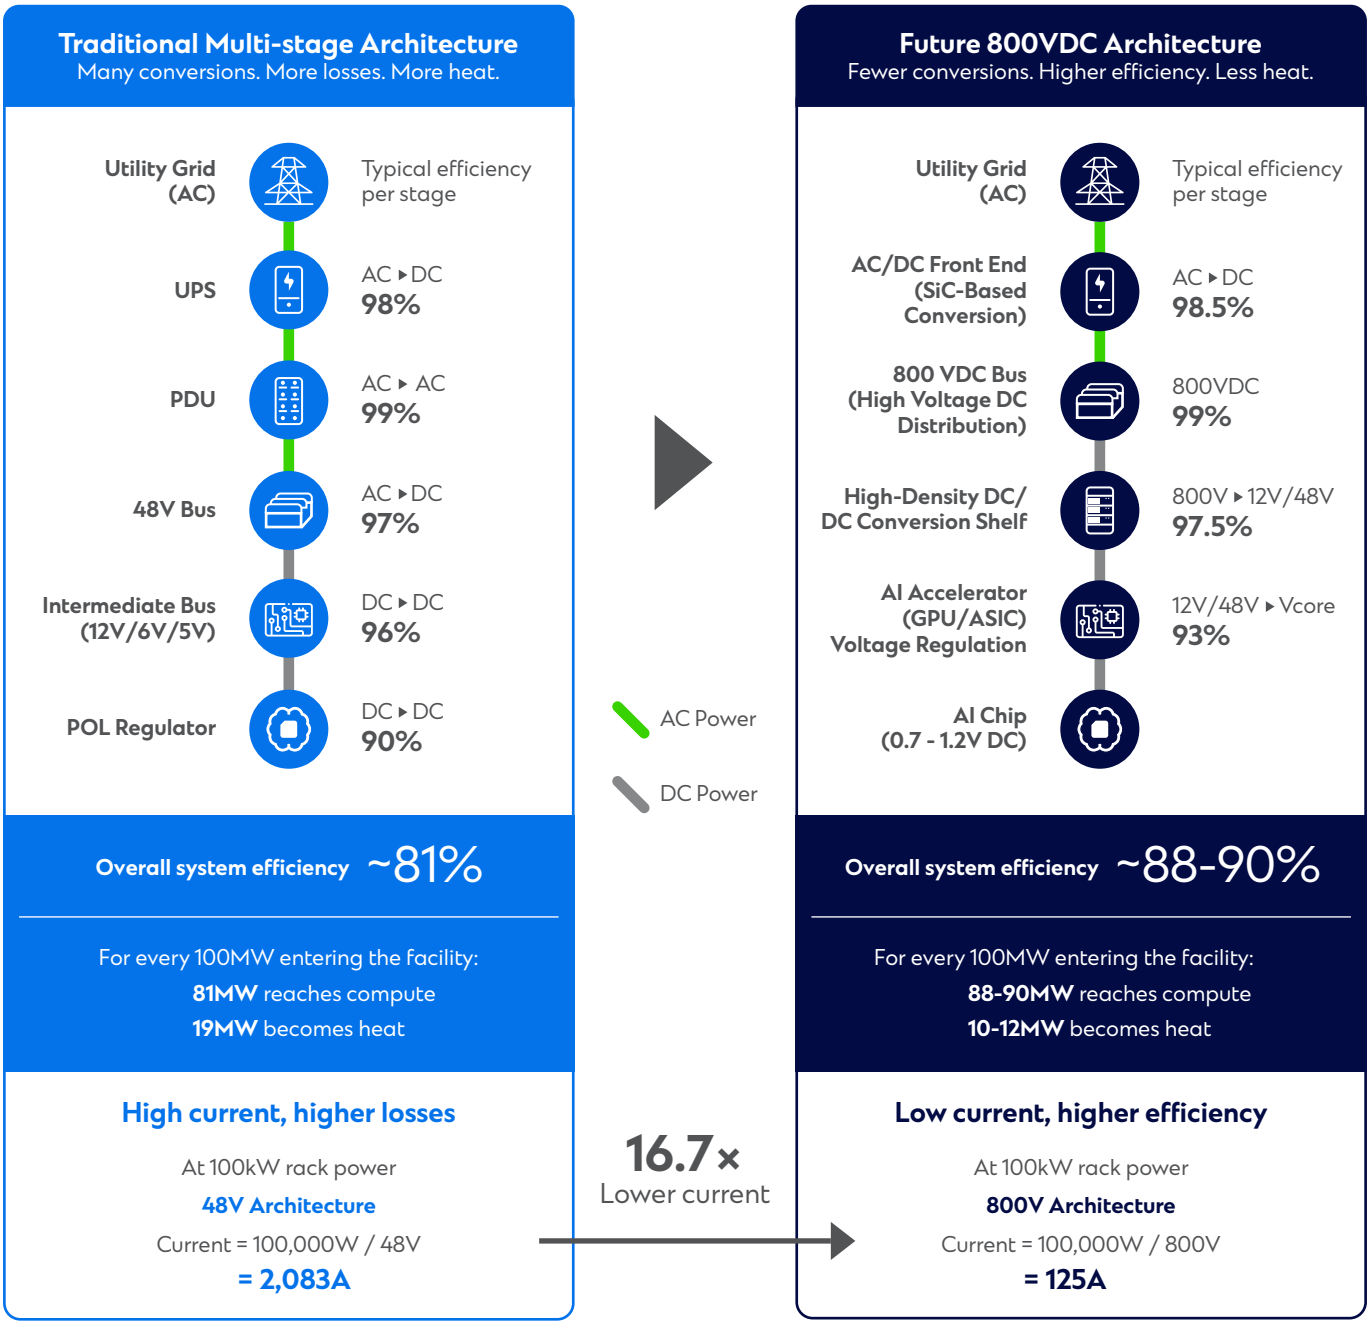
Some of the industry's main players are therefore working together and have developed an 800-volt direct current (DC) distribution architecture for data centres. Moving from 400 to 800 volt cuts the current in half for a given level of power which lowers resistive heat losses and enables lighter, thinner copper cabling. It not only simplifies power distribution but also increases energy efficiency (Open Compute, 2026). Google, Microsoft and NVIDIA have been developing the 800 VDC infrastructure through the Open Compute Project (OCP) and their work has already attracted more than 80 equipment manufacturers and infrastructure companies that are building products aligned with the 800 VDC specifications (NVIDIA, 2026b).

The proposed 800 VDC infrastructure not only improves power efficiency or reduces power loss (Figure 52) but it also supports using high density modular power blocks for rapid deployment and forms the backbone for optional Battery Energy Storage System (BESS) integration and future DC microgrids. Calculations from Enverus, NVIDIA and Semi Analysis suggest that adopting the 800 VDC infrastructure can improve facility efficiency levels by 5-14% (Semianalysis, 2026). Enverus also estimates that the amount of copper mass falls by up to 60% per MW and that an 800 VDC architecture cuts electrical capital expenditure requirements by 13% relative to current c. 400-volt AC architectures (Enverus, 2026). NVIDIA's calculations suggest that an AI factory using an integrated architecture of 800 VDC combined with the use of dry coolers and a closed loop system could achieve a PUE of 1.1 compared with around 1.4 for conventional data centres using evaporative cooling.



Figure 52

Power architecture evolution in AI data centres



Benefits of fewer conversion stages

- **Higher efficiency** — More power delivered to compute
- **Lower thermal load** — Reduced cooling requirements
- **Lower cost** — Less power hardware and copper
- **Smaller footprint** — More power density per rack
- **Improved reliability** — Fewer components, fewer failure points

Source: Semi Engineering, Standard Chartered Bank

Solution 5

Develop battery storage and microgrid capacity

Battery storage will likely take on greater prominence if renewable power is to become a driving factor in supporting data centre capacity expansion as it addresses the intermittency of solar and wind, and can help maximise the use of onsite renewable power generation (Schneider Electric, 2023). A microgrid combines local power generation with energy storage which allows a data centre facility to operate if needed independently of the main grid. By combining batteries with onsite solar, wind, fuel cells, gas generation or contracted renewable power, data centre operators can improve the resilience of their operations through an increase in self-generated renewable power, reduced curtailment risk and potentially provide services such as frequency response or reserve capacity back to the main grid. For data centres, this not only strengthens their resilience but it can also reduce their reliance on emission-intense diesel generators that might have been used for backup purposes previously (Vertiv, 2024).

AI data centres can be power dense and, when connected to the local grid, can create supply-demand stress periods. Even without onsite renewable power capacity, data centres may still want to install battery storage capacity as this has the potential to reduce the impact that data centres have on local grids during peak periods. Batteries can be charged when electricity is cheaper or supply is greater and be discharged during peak demand periods to help smooth load profiles on the electricity grid. This so-called energy arbitrage also allows data centre operators to lower their energy costs in a meaningful way. Calculations from Vertiv show that a 10MW Battery energy storage systems (BESS) system can save more than USD1 million per year on energy arbitrage alone, depending on the utility rate structure (Vertiv, 2025).

Current battery storage capabilities remain largely focused on supporting uninterrupted power supply (UPS) and help store renewable energy for a relatively short period of time. Potentially, battery energy storage systems can be used to provide the data centre value chain with as many as 13 different services (Schneider Electric, 2024).

Even though battery technology continues to develop and improve, the investment case for a battery and solar data centre does already seem to look promising. Calculations from Moduledge suggest that the payback period for a data centre in Southern Europe using a solar and battery set up for 50% of its electricity needs is less than 2 years (Moduledge, 2026). Calculations from Vertiv suggest that the payback period for a BESS system in the US is between 3 to 5 years and that this can have an internal rate of return of up to 18% (Vertiv, 2025).

Benefit to power and water intensity of data centres

As enabling solutions, battery storage and microgrids provide indirect benefits to a data centre's power and water intensity. By enabling a greater use of onsite renewable power, they reduce strains on the grid and when accompanied by a shift away from traditional power sources also reduce emission generation and water consumption.

Integrating all solutions can turn data centres into flexible energy system participants

Taken together, the five solutions outlined in this chapter can reposition leading data centres from large, largely passive electricity users into more efficient and flexible participants in the energy system. Operational improvements across software, hardware, utilisation and thermal management reduce the power required for each unit of useful compute, while heat reuse converts an otherwise wasted output into a source of energy for nearby buildings or industry. Renewable procurement expands the supply of low-carbon electricity, 800-volt DC distribution reduces conversion and transmission losses within the facility, and battery storage and microgrids add the capacity to manage when and how grid electricity is used.

An integrated approach of all solutions can help manage peak demand without compromising the availability of mission-critical services. Coordinated through intelligent energy-management systems, these solutions can smooth the data centre's load profile, reduce exposure to price volatility and make better use of existing generation, network and connection capacity.

By reducing the power required for each unit of compute, helping finance new clean generation, absorbing surplus renewable output and reducing grid draw during constrained periods, data centres can support better utilisation of existing generation, network and connection capacity. Data centre operators' contribution should therefore be assessed not only by the electricity they consume, but also by the demand they avoid, the clean power they help bring online, the useful heat they supply, and the flexibility and resilience services they provide to the wider energy system.



AI-enabled energy gains may outweigh its own power demand

The rise in AI-driven data centre electricity demand has raised concerns that large digital loads could tighten power supply or increase electricity costs for households and other businesses. The solutions highlighted on the previous pages can help reduce some of these concerns, however, the broader impact of AI on the global energy system is often overlooked, including the extent to which this could be counterbalanced by AI-enabled efficiency gains across key sectors.

If AI is used to optimise energy consumption in buildings, industry, transport and power systems, the resulting efficiency gains in those sectors may perhaps offset part of the additional electricity demand from data centres. Under some estimates it could even contribute to a net reduction in overall energy use.

- **Buildings:** The strongest argument is that AI can reduce energy demand in much larger end-use sectors than data centres. Buildings are a good example: the IEA notes that AI can optimise heating, cooling, lighting and load management, and that widespread adoption of AI-driven building optimisation could deliver around 300TWh of annual global electricity savings (IEA, 2025). This is almost half of the incremental electricity demand associated with the AI scenarios between 2025 and 2030. Additional savings potential across other industries are highlighted in Figure 53.
- **Transport and Industry:** Other analysis from the IEA suggests that full scale adoption of AI by the energy sector has the potential to boost the efficiency of the transport and industry sectors by up to 10% by 2035 (IEA, 2026). This, according to the IEA, would equate to energy savings of 13.5EJ which is roughly 3,750TWh or more than seven times the total global level of electricity consumed by the data centre sector today.

Figure 53 provides an overview of some of the energy-saving potential that AI deployment can achieve across key sectors. Although these estimates are significant, there are reasons to treat these with some caution. First, data centre load is concentrated or local while AI-enabled savings in sectors such as buildings, industry, transport and grids are more diffuse and depend on adoption rates. Second, some of the potential efficiency gains could be eroded by the Jevon's paradox as mentioned earlier if AI makes services cheaper, faster and more widely available and this triggers a rise in usage. Nevertheless, even a partial adoption of AI across other industries may already deliver energy and emission savings that offset a significant share of the energy associated with running AI and data centres.



Figure 53

Potential AI-enabled savings across sectors

Sector	AI-enabled saving mechanism	Indicative saving / impact	Source links
Buildings	AI optimisation of HVAC, lighting, occupancy and load management	IEA analysis suggests widespread AI-enabled building optimisation could deliver around 300TWh of annual electricity savings globally.	(IEA, 2025)
Data centres	AI optimisation of cooling systems	Google/DeepMind reported up to a 40% reduction in data-centre cooling energy. Other analysis suggests that 30-40% reduction in energy and water usage is possible.	(Google, 2016) (Ali, 2026)
Power grids / energy systems	AI-enabled optimisation of power plant operations, maintenance, dispatch and grid utilisation	IEA estimates AI in power plant operations and maintenance could yield up to USD110 bn/year in cost savings by 2035, and AI could unlock up to 175GW of additional transmission capacity from existing lines.	(IEA, 2025)
Energy systems economy-wide	AI-enabled asset optimisation, predictive maintenance, real-time balancing and energy-system management	Deloitte estimates AI could save almost 12,000TWh of energy by 2050 and unlock nearly USD500 billion in cost reductions; it also states annual AI-enabled energy savings by 2030 could be around three times AI's projected energy use.	(Deloitte, 2025)
Industry / manufacturing	Predictive maintenance, process optimisation, digital twins and real-time energy management	Schneider Electric highlights analysis showing that AI-enabled industrial processing plants have reported 10-15% production gains and 4-5% EBITA gains, with energy optimisation a key use case; other industrial AI studies find energy reductions of 25-35% in selected early-stage test cases.	(Schneider Electric, 2024) (Dos Santos Junior, 2025)
Transport / logistics	AI route optimisation, dynamic scheduling, traffic avoidance and load optimisation	Literature on AI-enabled logistics suggests route optimisation can reduce fuel consumption by up to 15-30% in specific scenarios; public-transit case studies show examples such as 12% fuel reduction and 500,000 gallons of diesel saved annually.	(Sarioguz, 2025) (Supplychain briefing, 2025)
Cross-economy modelling	Economy-wide AI-enabled efficiency improvements versus incremental data-centre demand	PwC modelling suggests that if AI improves energy efficiency across the economy at only one tenth of its adoption rate, it could offset the additional energy required by data centres, making AI's net energy and emissions impact broadly neutral.	(PWC, 2026)

Source: Standard Chartered Bank

In summary

This chapter shows that rising AI workloads could materially increase data centre electricity demand, creating pressure on grid connections, generation adequacy, transmission capacity and local distribution networks. However, the chapter also highlights that operators have multiple levers to reduce the scale of this challenge, including software and hardware efficiency gains, higher server utilisation, improved thermal management, waste-heat reuse, clean-energy procurement and more efficient power-distribution architectures such as 800VDC.

Hyperscalers currently appear better positioned than smaller peers because of their scale, purpose-built facilities, superior reported PUE levels, advanced cooling capabilities and ability to procure or develop large volumes of clean power. To raise the sustainability performance of the broader sector, hyperscalers could help accelerate the dissemination of best practices, operational experience and proven technologies across the industry. The sustainability debate around data centres should focus not only on the sector's growing electricity demand, but also on how effectively operators combine efficiency, clean-energy procurement, grid engagement and new power architectures to reduce pressure on electricity systems and support broader decarbonisation.



Will low-water data centres become the norm?

Data centre water consumption is often viewed negatively, particularly where new facilities are planned in regions facing local water stress. This chapter shows that the water story is more nuanced than headline concerns often suggest and highlights several solutions that, if adopted, may cause future water use to be lower than is currently expected. These include wider adoption of liquid and dry cooling, higher-temperature server operation, better site selection, increased use of recycled or non-potable water, the development of water replenishment programmes, renewable power procurement and potentially stronger operator and regulatory targets. AI-optimisation may add further momentum by accelerating the shift toward high density infrastructure that is harder to cool with conventional evaporative systems and better suited to cooling technologies with a low- to virtually-zero water usage. As a result, low water data centres could become increasingly common, provided operators combine technical innovation with transparent disclosure, careful location choices and active engagement with local communities.

Drivers of the data centre water footprint

The rapid build-out of data centre capacity is intensifying scrutiny of the sector's environmental footprint. Rising electricity demand is one concern as discussed in the previous chapter but water is emerging as another area of focus.

From withdrawal to replenishment

The water intensity of a data centre ultimately depends on its net impact on local freshwater resources. Understanding this impact requires distinguishing between four key water metrics.

- **Water withdrawal** refers to the total volume of water taken from the environment for data centre operations.
- **Water consumption** represents the portion of that water that is not returned to the local ecosystem, typically because it is evaporated or otherwise lost during use.
- **Water discharge** is the volume of water returned to the environment after use.
- **Water replenishment** reflects efforts to restore water resources through activities such as watershed restoration, conservation projects and habitat rehabilitation.

Together, these metrics provide a more complete picture of a data centre's true impact on water availability and water security.

Three factors drive the data centre water footprint

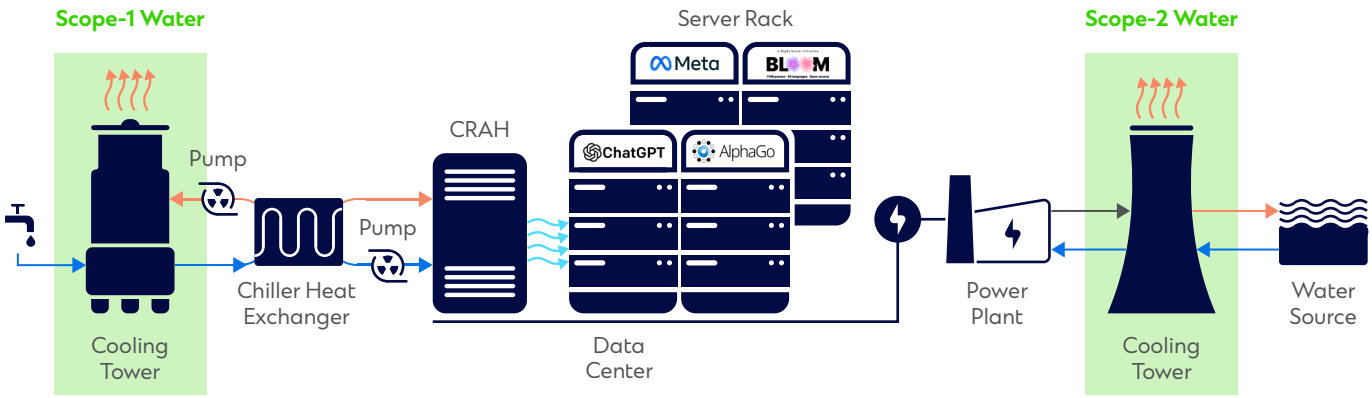
The water footprint of a data centre is often larger and more complex than commonly assumed. This footprint is made up of three distinct components:

- **Scope 1:**
  Direct water use for cooling and thermal management to maintain the reliability, performance and lifespan of increasingly power-dense IT equipment (Al Kez & et al, 2025). This depends on the type of cooling technology that is used by a data centre which, as shown elsewhere in this report, is partly dependent on local climate conditions.
- **Scope 2:**
  Indirect water consumption linked to the electricity used to operate data centres, reflecting the water intensity of different power-generation technologies.
- **Scope 3:**
  Embedded water use from the production of IT hardware and the broader supply chains that support data centre construction and operation.

Figure 54

Data centre water consumption

Scope 3 not included

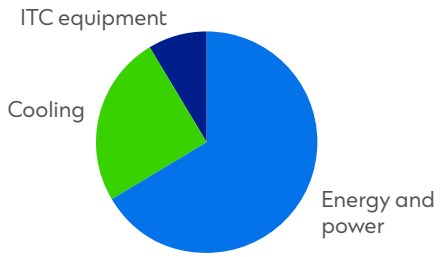


Source: OECD.AI, Standard Chartered Bank

Figure 55

Data centre water usage by scope

% of total, 2023



Source: (IEA 2025), Standard Chartered Bank



Scope 1: Water is used to cool data centres

Almost all electricity consumed by data centres is ultimately converted into heat. Unless this heat is removed, equipment might get damaged as temperatures rise too much. The acceleration in data centre computing intensity is increasing the need for effective cooling solutions. This trend is particularly pronounced for AI workloads, as AI training and inference can be significantly more computationally intensive than traditional digital workloads, generating substantially higher heat loads that must be managed to maintain performance and reliability (Al Kez & et al, 2025). Various cooling technologies exist that use water as part of the cooling process. As a result, c. 25% of total data centre water usage is driven by cooling requirements (Figure 55).

Scope 2: Indirect water use is the biggest driver of the water footprint

While water used for cooling is essential to data centre operations, the IEA estimates that direct cooling-related water use accounts for less than one-quarter of total water consumption (Figure 55). By contrast, approximately two-thirds of the sector's water footprint is linked to electricity generation, making the water intensity of power supply currently a more significant driver of overall data centre-related water consumption. This highlights that improving water stewardship for data centres is not solely a cooling challenge but also an energy sourcing challenge.

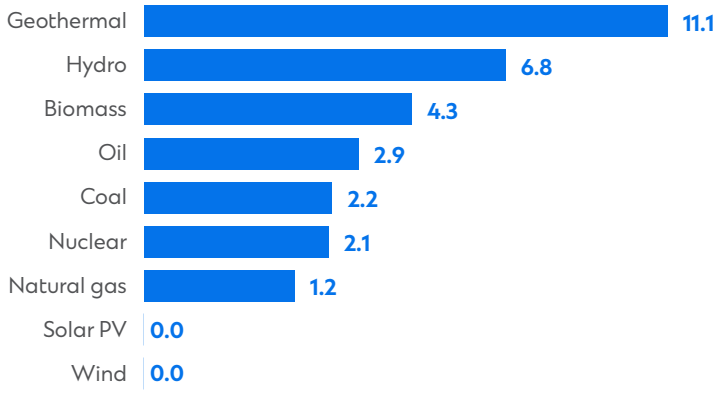
Location materially influences a data centre's water footprint and impact

Analysis from (Lei & et al) shows that geothermal power requires more than 11 litres of water per kilowatt-hour generated, while solar and wind power consume virtually no water during electricity generation (Figure 56). As a result, two data centres with identical cooling systems and electricity demand can have markedly different water footprints depending on the composition of their local power grid. This underscores the importance of renewable energy procurement not only as a decarbonisation strategy but also as a water-risk mitigation tool.

Figure 56

Water consumption factor by power source

l/kWh



Source: Lei et al, Standard Chartered Bank



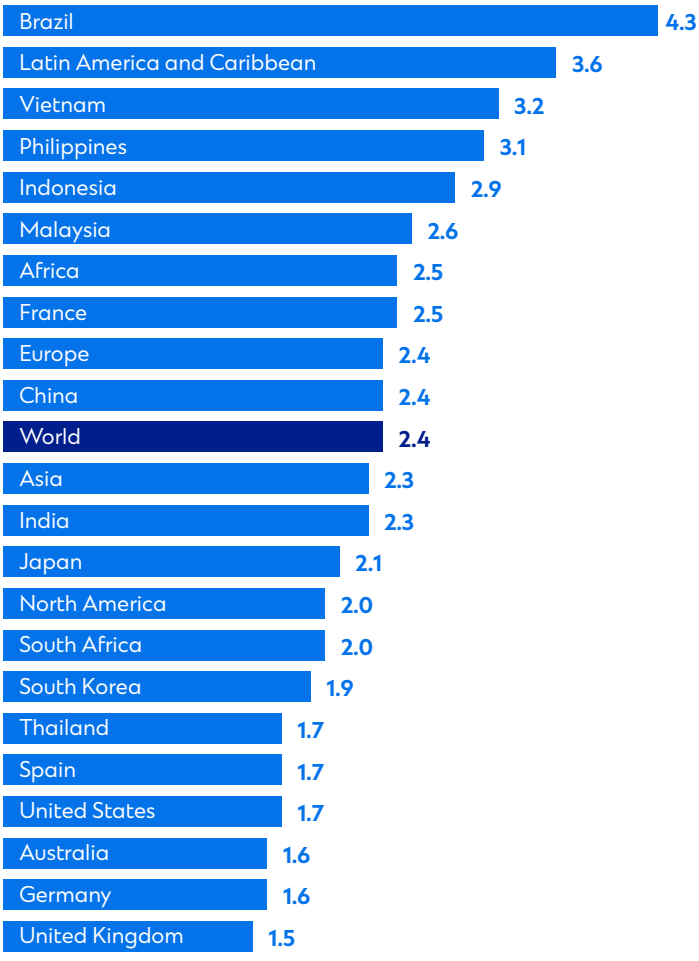
Combining country-level power generation data from EMBER with technology-specific water consumption factors suggests that the water intensity of electricity generation in markets such as Brazil, Vietnam, the Philippines and Indonesia can be more than double that of countries including the UK, Germany and Australia (Figure 57). This implies that the environmental impact of data centre expansion is highly location dependent. Analysis from (Li & et al.) provides further support for this as it estimated that the water intensity of training the GPT-3 large language model developed by Open-AI in Washington, Arizona or Sweden is substantially higher than if that training was done using data centres in Ireland, Texas or Georgia (Figure 58). Importantly the analysis also shows that the majority of water consumption is Scope 2 rather than Scope 1 related.

The location of data centres is also material if its direct and indirect water consumption amplifies water-related risks. Data from the World Resources Institute (WRI) shows that most existing data centres are situated in regions facing medium to high levels of water stress (Figure 59). Looking ahead, many of the markets expected to drive the next wave of data centre expansion also exhibit elevated water vulnerability. These include India, China, Vietnam, Thailand and the Middle East (Figure 60). This may create a growing tension between digital infrastructure ambitions and sustainable water management, suggesting that water availability, efficiency and resilience may become increasingly important factors in future data centre site selection, operating costs and long-term licence to operate.

Figure 57

Water intensity of electricity by region

l/kWh, weighted average using 2024 data



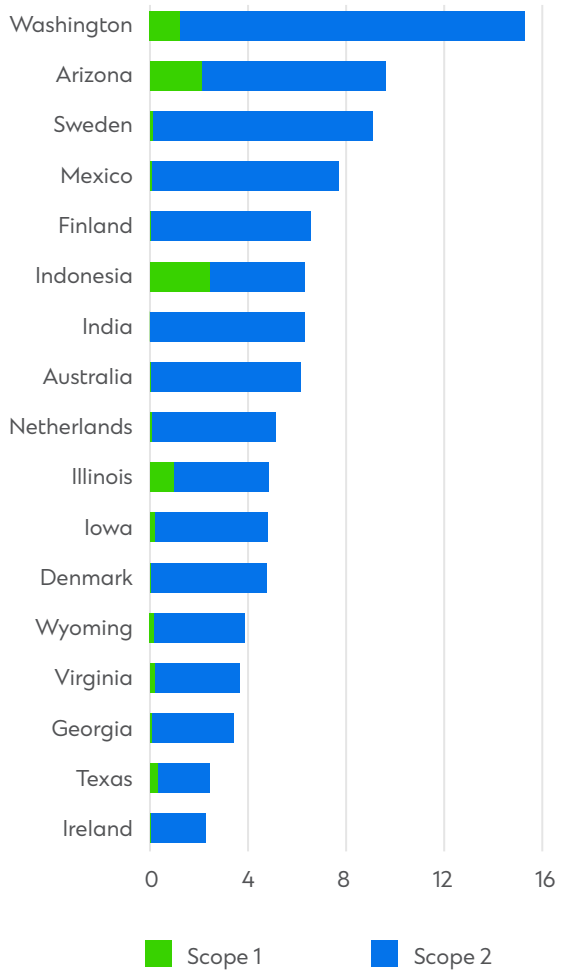
Source: Lei et al, EMBER, Standard Chartered Bank

For investors and operators, assessing water risk therefore requires a broader lens that incorporates not only cooling technologies and direct water withdrawals, but also the water intensity of local electricity systems, future power procurement strategies and the long-term sustainability of regional water resources.

Figure 58

Water used for training GPT-3

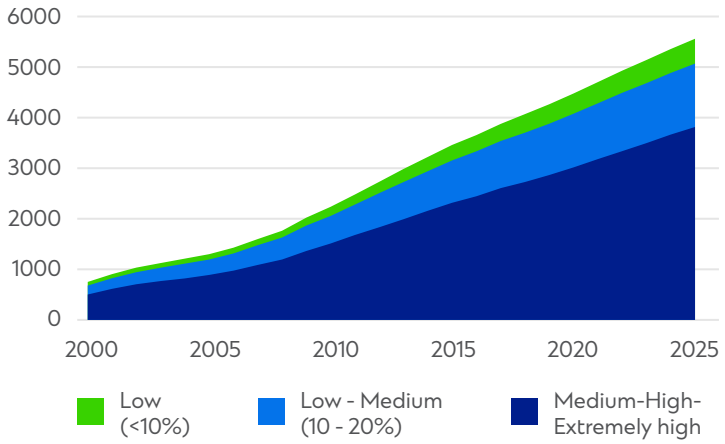
L bln



Source: Lei et al, Standard Chartered Bank

Figure 59

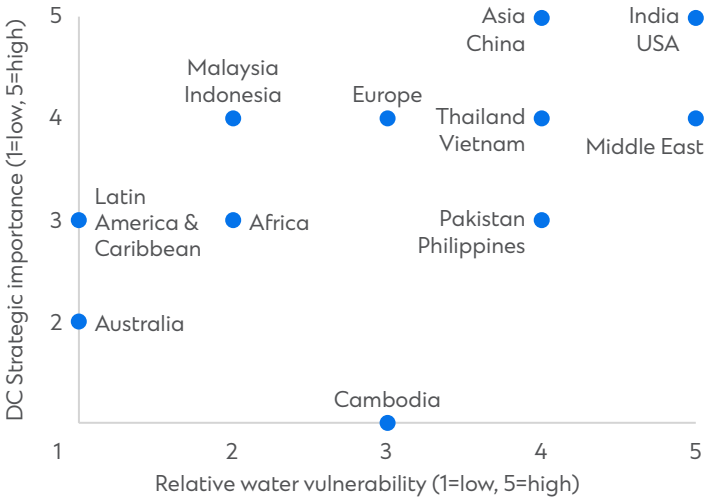
Number of data centres grouped by WRI 2030 water stress indicator



Source: Aquaduct database from WRI, Bloomberg, Standard Chartered Bank

Figure 60

Water vulnerability versus the strategic importance of data centres for key countries/regions



Source: WRI, Standard Chartered Bank

Direct data centre water use: perhaps less than expected

The rapid expansion of data centre capacity, particularly to support AI-related workloads, is likely to increase both the sector's direct and indirect water footprint unless meaningful action is taken across the value chain. Before assessing the strategies available to reduce water intensity, it is important to place the sector's current water footprint in context.

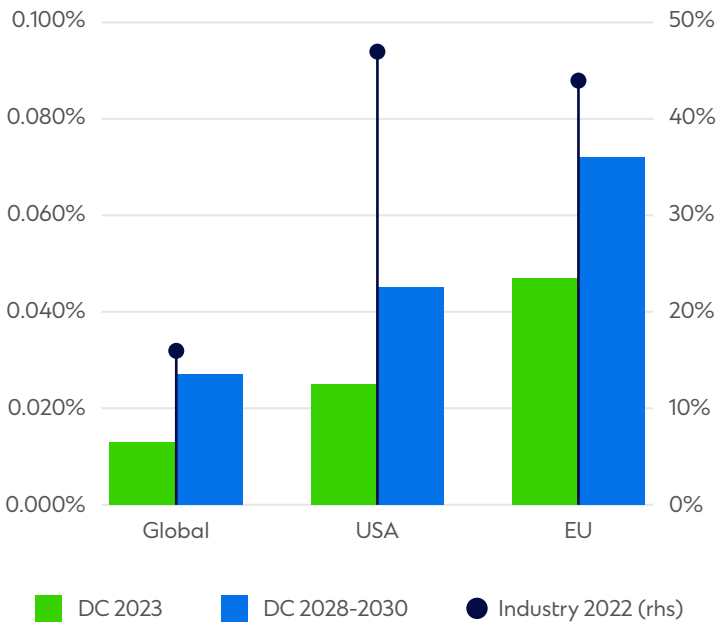
The IEA estimates that data centres consumed approximately 560 billion litres of water in 2023 of which 150 billion litres was associated with direct consumption (IEA, 2025). While substantial in absolute terms, this represents only around 0.01% of global freshwater consumption excluding agriculture. The volume of freshwater that is consumed directly by data centres in the USA and EU is also low as it represents just 0.025-0.047% of freshwater consumed by industry and municipalities combined for 2023 respectively (Figure 61).

Some medium-term data centre water consumption projections indicate that this could reach up to 120 billion litres for the US and 94 billion litres in the EU (Shehabi & et al, 2024) (Water Europe, 2024). Although these growth rates are high, the volume of data centre-related freshwater consumption would still be less than 0.1% of that consumed by industry and municipalities combined for both regions (Figure 61).

Figure 61

Water intensity comparison

% of total freshwater consumption



Source: World Bank, IEA, Water Europe, (Shehabi & et al, 2024), Standard Chartered Bank



Some long-term projections suggest that overall data centre-related water consumption could rapidly increase under business-as-usual growth assumptions (Herrera & et al, 2025) (IEA, 2025) (de Vries-Gao, 2025). Scenarios such as these should be interpreted with caution because they often assume limited improvements in water efficiency. As highlighted later in this report data centre operators are increasingly investing in advanced cooling technologies, water recycling and reuse systems, alternative water sources, and replenishment programmes, while also adopting more ambitious water-stewardship targets. As these measures scale, they have the potential to materially reduce future water consumption and weaken the link between data centre growth and water demand. Nevertheless, even if data centre water consumption reaches 10 trillion litres, this would still represent only around 0.23% of current global freshwater consumption, all else being equal.

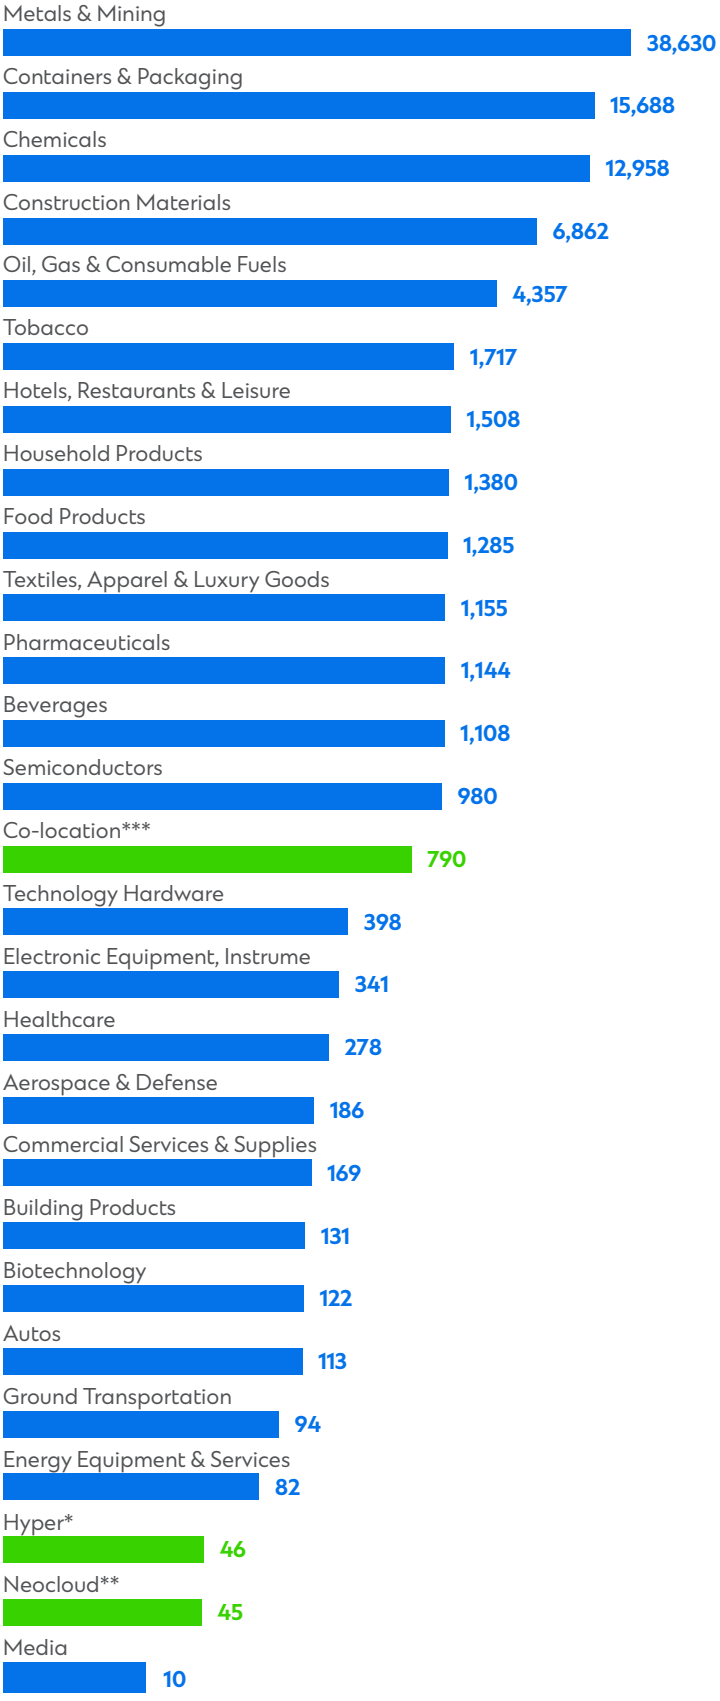
Data centre water intensity appears moderate relative to other sectors

Another way to assess the water footprint of data centres is by comparing the sector with other industries. Data for companies included in the S&P 500 suggest that the largest data centre operators have significantly lower water intensity, measured as water consumption per unit of revenue, than most other sectors (Figure 62). Although this comparison is primarily based on direct water use and therefore excludes the indirect water consumption associated with electricity generation, the overall conclusion is unlikely to change materially. All sectors rely on electricity and therefore have a Scope 2 water footprint. Given the substantial advantage that leading data centre operators already display in terms of direct water efficiency, they are likely to remain among the less water-intensive sectors even when indirect water consumption is taken into account.

Figure 62

Water intensity comparison

m3/USD mln revenue, S&P500 constituents by US sector



Source: Bloomberg, Data centre companies, Standard Chartered Bank
* = Microsoft, Google, Amazon and Meta,
** = Nebius Group,
*** = Digital Realty and Equinix

Data centre cooling developments

Cooling a data centre requires heat generated by the IT equipment to be transferred outside of the data centre either into the air or to be used in other processes for example for district heating. A data centre's cooling system consists of two elements (Figure 63).

- The technology cooling system (TCS) is directly connected to the servers and consists of a closed loop system that can circulate air or a liquid between the heat-producing IT equipment and a heat exchanger or that can use a hybrid structure of liquid and air. In case of an air-based TCS the heat exchanger will be a computer room air handler (CRAH) or a computer room air conditioner (CRAC). A liquid-based TCS will use a coolant distribution unit (CDU) as the heat exchanger. The two main types of liquid cooling solutions are direct-to-chip and immersion cooling.
- The heat rejection system (RHS) absorbs heat from the heat exchanger and transfers this outside the data centre. In contrast to the TCS, the RHS can be based on an open or closed loop structure. In case the data centre uses an open loop RHS structure, water will be lost to the atmosphere through evaporation which significantly increases total water use as this cannot be reused by the data centre. An RHS can include water- or air-cooled chillers which use cold air or water to extract heat, evaporative or dry cooling towers which transfer heat from chillers to the outside atmosphere or make use of air- or water-side economiser systems which operate a second cooling loop that can bypass chillers and thereby reduce energy consumption.

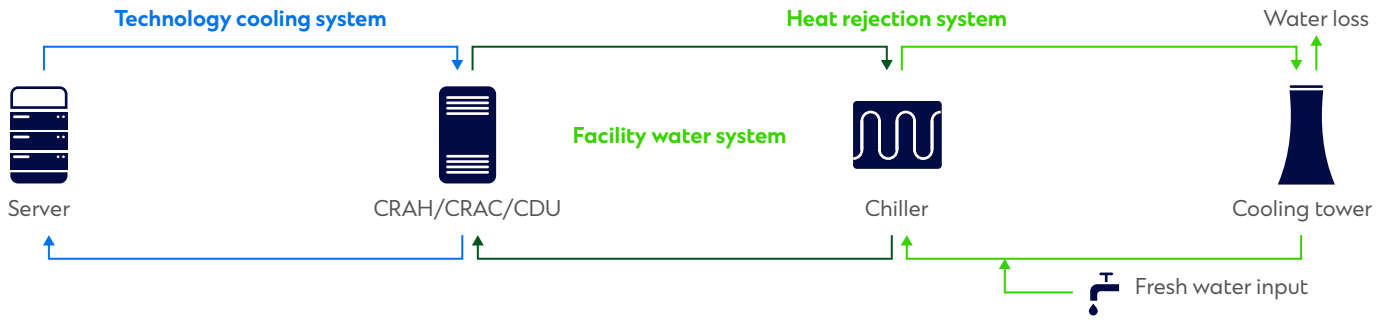
Historically, most facilities have relied on air-based cooling solutions as these are simple and relatively cost effective to install. However, the increasing power density of AI-related infrastructure is pushing the limits of conventional air cooling. Air has a relatively low heat capacity suggesting that to transfer increasing amounts of it requires data centres to run the TCS cooling fans and the RHS chillers more intensely which requires more energy.

Liquid-based cooling solutions have been developed to benefit from their much higher heat capacity (e.g. water ran remove c4,200 times more heat than air) and provide data centres with a more efficient thermal-management strategy. One of the key benefits of liquid cooling solutions is that they allow the use of a more water efficient RHS and thereby significantly reduce a data centre's water consumption. Another benefit from liquid cooling is that they tend to be more energy efficient than air-based cooling. Reasons why liquid cooling has not been adopted more aggressively in the past include the fact that it is more complex and costly to install. Another relevant factor relates to the liquid that is used and the potential risk that it may come in contact with server electronics and damage it.

Each cooling solution involves different trade-offs between energy efficiency and water consumption (Figure 64). Their performance is also heavily influenced by local climate conditions, meaning that the optimal cooling strategy can vary significantly across locations. As a result, cooling technology choices are becoming an increasingly important determinant of both the power and water footprint of data centres.

Figure 63

Cooling system structure of a data centre



Source: Standard Chartered Bank

Figure 64

Key data centre cooling solutions

Non-exhaustive

Solution	Principle	Best climate conditions	Energy usage	Water usage
Air and evaporative cooling solutions				
Evaporative	Uses direct or indirect evaporative heat rejection where water evaporation removes heat from air or a water loop.	Hot, dry climates with low wet-bulb temperatures	Low	Very high
Adiabatic	Uses adiabatic pre-cooling (water mist or wetted media) to reduce intake air temperature before dry cooling.	Temperate to hot climates where humidity is moderate.	Below average	High
Computer room air handling - dry	Chilled-water air handling units connected to dry coolers or water-side economizers with no evaporative heat rejection.	Cold to cool-temperate climates.	High	Above average
Computer room air handling - wet	Chilled-water air handlers connected to cooling towers or evaporative condensers.	Warm climates with adequate water availability.	Average	High
Computer room air conditioning - dry	DX refrigeration units with dry fluid coolers; mechanical refrigeration provides primary cooling.	Cool-to-moderate climates.	High	Above average
Computer room air conditioning - wet	DX system with evaporative condenser or cooling tower support.	Warm climates where water-energy trade-off favours evaporation.	Average	High
Free air cooling solutions				
Direct free air cooling	Outdoor air directly enters white space after filtration and conditioning.	Cool, clean-air climates with moderate humidity.	Low	Low
Indirect free air cooling	Outdoor air cools indoor air through heat exchangers without direct mixing.	Cool to temperate climates; more tolerant of humid environments than direct-air systems.	Low	Average
Liquid cooling solutions				
Direct to chip liquid (closed loop)	Liquid circulated through cold plates on CPUs/GPUs.	Largely climate independent, although dry coolers perform best in cooler climates.	Below average	Low
Rear-door heat exchanger (RDHx)	Water-cooled heat exchanger mounted on rear of rack captures exhaust heat before entering room.	Climate independent	Average	Low
Immersion (1 phase)	Servers fully submerged in dielectric fluid that remains liquid throughout operation.	Climate independent.	Low	Very low
Immersion (2 phase)	Dielectric liquid boils on hot components and condenses via an internal heat exchanger.	Climate independent.	Very low	Very low
Hybrid cooling solutions				
Hybrid air-liquid	Combines direct liquid cooling for high-density equipment with air cooling for remaining loads.	Suitable for virtually all climates.	Below average	Above average
Hybrid dry/adiabatic	Dry cooling for most operating hours with adiabatic assist during peak temperatures.	Temperate climates with seasonal hot periods.	High	Very low

Source: Different desk research sources, Standard Chartered Bank

Data centre water efficiency: hyperscalers set the tone

Investors, regulators and data centre operators cannot effectively manage, benchmark or reduce water-related impacts without a clear understanding of how efficiently water is being used across different data centre facilities. However, measuring water performance is far from straightforward as differences in cooling technologies, climate conditions, power sourcing and reporting practices can create significant variations in water intensity, making comparisons between operators challenging. As a result, robust and transparent metrics are essential.



Measuring efficiency: Water Use Effectiveness (WUE)

Water Use Effectiveness (WUE) has emerged as one of the most commonly adopted metrics for assessing data centre water efficiency. It measures the amount of water consumed relative to the electricity used by a data centre's IT equipment. However, WUE is not a straightforward metric as various methodologies have been developed by different organisations. The most used methodologies are:

The International Organisation for Standardisation (ISO): As part of its series on best practices for data centres the ISO has defined three categories for the calculation of WUE (MOSL, 2026).

- **Category 1:** This measures only a data centre's direct water consumption relative to the electricity consumed by its IT equipment. This approach excludes indirect water consumption associated with electricity generation and does not account for water that is discharged or reused.
- **Category 2:** This measure refines the calculation by excluding water that is used for non-IT-related purposes, such as irrigation, landscaping or toilet flushing. The resulting ratio therefore focuses more specifically on the water consumption associated with supporting data centre operations and cooling activities.
- **Category 3:** This measure provides the most comprehensive assessment of water intensity. It incorporates both direct and indirect water consumption while also accounting for water that is recycled, reused or replenished. By capturing water impacts across both facility operations and electricity supply, this approach offers the broadest view of a data centre's overall water footprint.

The Green Grid, an organisation consisting of stakeholders associated with the global data centre industry, has defined two ways to calculate a data centre's WUE that broadly speaking coincide with category 1 and 3 of the WUE definitions listed earlier (The Green Grid, 2011).

- **WUE:** This measure assesses the total amount of direct water usage on-site for the operation of a data centre relative to the level of energy consumed by the data centre's IT equipment. This WUE measure is similar to the category 1 ratio as defined by ISO and used in Europe.
- **WUE – source:** This ratio includes the total amount of direct water usage and the water used off-site in the production of the energy used on-site by the data centre.

Assessing WUE requires a view on PUE too

The development of water intensity metrics for data centres is helpful, however, assessing a data centre's water efficiency requires an approach that incorporates the wider context of its operations. More specifically it should acknowledge the interplay between water and power consumption. A data centre's water and power consumption are closely connected suggesting that strategies that reduce direct water consumption may cause a deterioration in energy efficiency and indirect water consumptions and vice versa.

- For example, air-cooled systems can significantly lower a data centre's direct water use and improve its WUE (category 1), however, this may require more electricity which worsens its PUE and increases a data centre's indirect water consumption or WUE (category 3).
- Conversely, evaporative cooling systems can deliver superior energy efficiency and lower PUE and indirect water consumption by reducing cooling-related electricity demand, but these cooling systems typically cause direct water consumption to increase.

As a result, optimising one metric in isolation can unintentionally increase environmental impacts elsewhere. Effective water stewardship therefore requires a holistic approach that simultaneously considers both water and energy efficiency.

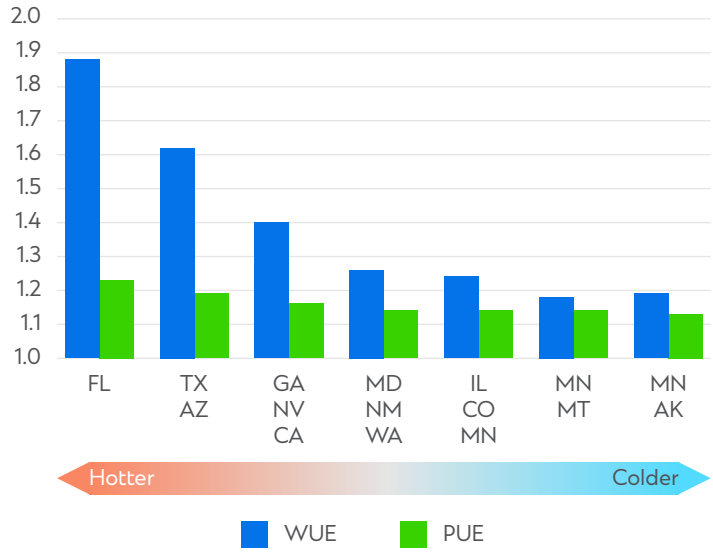
Climate conditions matter when judging a data centre's efficiency

Even if power and water characteristics of a data centre are included, a review of a data centre water performance is only just if its performance metrics are benchmarked against those of data centres located in areas with similar climate conditions. Data centres located in warmer or more humid climates generally require more cooling to maintain optimal operating temperatures. The nature of the traditional cooling solutions available for this suggests that this increases the water and power footprint of a data centre. Calculations from (Lei & Masanet) illustrate this as water and power intensity levels of data centres in warmer climates tend to be higher than those located in cooler areas (Figure 65). Their analysis also suggests that data centres located in humid areas tend to have higher water intensity (Figure 66).

Figure 65

WUE vs PUE by climate

L million / MWh for WUE and ratio for PUE

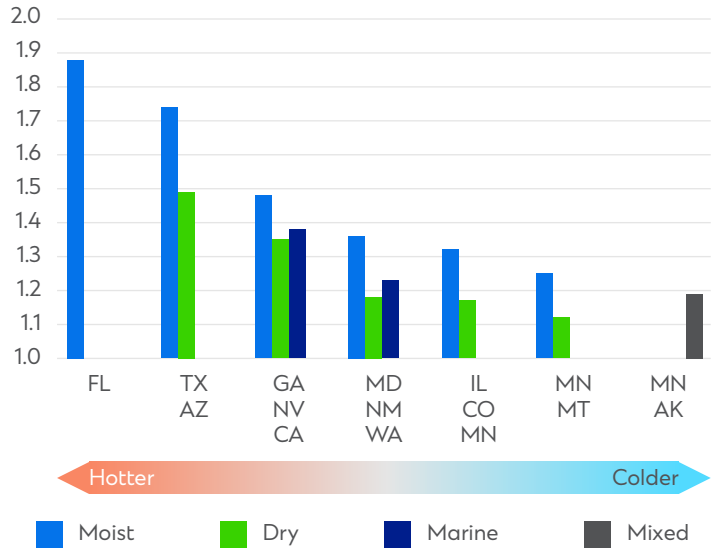


Source: Lei and Masanet, Standard Chartered Bank

Figure 66

WUE for different climate conditions

L million/MWh



Source: Lei and Masanet, Standard Chartered Bank

Hyperscalers lead the sector on water efficiency

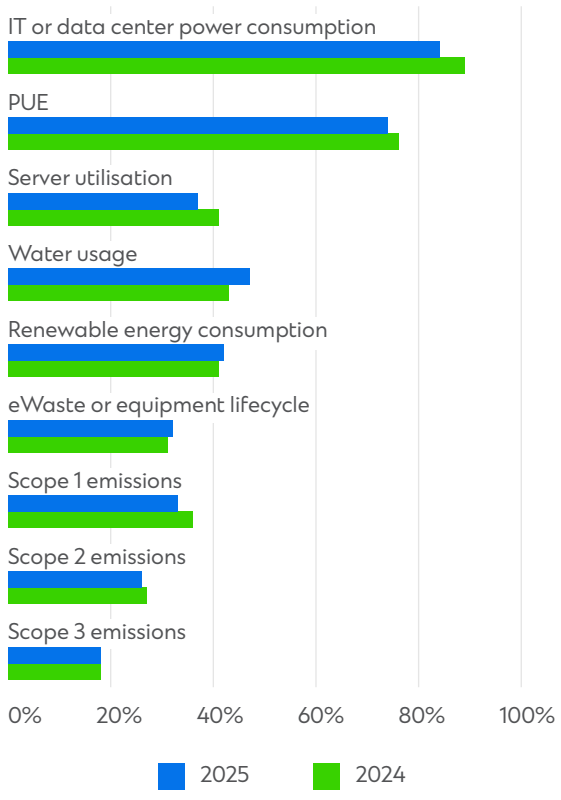
As the data centre sector expands, financiers and policymakers are likely to place greater emphasis on benchmarking different data centre operators on their power and water efficiency. However, as the previous section highlighted, most jurisdictions have yet to establish comprehensive disclosure frameworks suggesting that a data gap likely exists at present. Data from a recent Uptime survey highlights this point as it indicated that less than half of surveyed data centre operators actively collect water consumption data (Figure 67).

A review of power and water-related disclosure of more than 30 of the world's largest data centre operators shows that a data gap exists. While the largest data centre operators tend to provide more detailed information on their power and water performance, current disclosure levels on average are insufficient to identify with confidence which data centre operators perform best in terms of power and water intensity or who is best positioned to deliver future improvements on this (Figure 68).

Figure 67

Which of the following metric does your organisation collect

% of respondents



Source: Uptime, Standard Chartered Bank

Key observations on power and water efficiency related disclosure levels by key data centre operators include:

- Reported metrics are frequently aggregated at the corporate level rather than at individual facility level. For the companies reviewed more than 60% disclose their overall energy or water usage but less than 10% provide data on a facility level. The absence of site-specific data limits the ability for investors to make meaningful comparisons between data centre portfolios or analyse investment proposals for new individual data centres.
- Water disclosures often exclude the water footprint that is embedded in the electricity used to power a data centre operator's facilities. The only company of the reviewed list that publicly discloses its Scope 2 WUE ratio is CyrusOne.
- As highlighted earlier a data centre's power and water intensity depends largely on the type of cooling technology used, workload intensity, utilisation rates, hardware efficiency and local climate conditions. Operators do not disclose data on these factors. Competitive pressures are often highlighted as a reason for this.
- Although some operators provide data regarding the water efficiency of their operations, most do not provide a detailed description of the methodology used to calculate this. This makes it more difficult to compare one operator's calculations with those of another. Key uncertainties include which data centres are included in the WUE calculation. Some operators only focus on fully owned or leased centres while others focus on those that have been in operation for more than 1 or sometimes 2 years. Some operators exclude data centres that have utilisation rates of less than 20% or 30% while others might only include data for periods that a centre has a full load. Another area of uncertainty relates to how water use is defined by an operator. For example, is WUE based on water withdrawals or consumption and to what extent does it make corrections for certain types of water use such as recycled, reclaimed or non-potable water.

As a result of these challenges, current reporting frameworks provide only a partial view of operational sustainability performance and leave investors with limited ability to assess current efficiency performance of data centre operators or their potential to improve future efficiency, resilience and water-related risk exposure.

Figure 68

Disclosure of electricity and water efficiency levels for key data centre companies

	Electricity use reporting?		Water use reporting?		PUE reporting?		WUE reporting?			Specific water target disclosed?
Company	Company	By data center	Company	By data center	Company	By data center	Company	Scope 2	By data center	
AirTrunk	✓	✗	✓	✗	✓	✗	✓	✗	✗	✓
Alibaba	✓	✗	✓	✗	✓	✗	✓	✗	✗	✗
Aligned	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Amazon	✗	✗	✓	✓	✓	✗	✓	✗	✗	✓
Apple	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
Baidu	✓	✗	✓	✗	✓	✗	✓	✗	✗	✗
Beale Infrastructure	✗	✗	✗	✓	✗	✗	✗	✗	✗	✗
China Telecom	✓	✗	✓	✗	✗	✓	✗	✗	✗	✗
China Unicom	✓	✗	✓	✗	✓	✗	✗	✗	✗	✗
CloudHQ	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Compass Datacenters	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Crusoe	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
CyrusOne	✓	✗	✓	✗	✓	✗	✓	✓	✗	✓
DayOne	✗	✗	✗	✗	✓	✗	✗	✗	✗	✗
Digital Realty	✓	✗	✓	✗	✓	✗	✓	✗	✗	✗
EdgeConneX	✓	✗	✓	✗	✓	✗	✗	✗	✗	✗
Equinix	✓	✗	✓	✗	✓	✗	✓	✗	✗	✗
GDS Holdings	✓	✗	✓	✗	✓	✗	✓	✗	✗	✗
Google	✓	✗	✓	✓	✓	✓	✓	✗	✗	✓
Meta*	✓	✓	✓	✓	✗	✗	✓	✗	✗	✓
Microsoft	✓	✗	✓	✗	✓	✗	✓	✗	✗	✓
Nebius	✓	✗	✓	✗	✓	✗	✗	✗	✗	✗
NTT Global Data Centers	✓	✗	✗	✗	✓	✗	✓	✗	✗	✗
Oracle	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
QTS Data Centers*	✓	✗	✓	✗	✓	✗	✓	✗	✗	✓
Rowan Digital Infrastructure	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
SoftBank Group	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Stack Infrastructure*	✓	✗	✓	✗	✓	✗	✓	✗	✗	✗
Switch	✓	✗	✓	✗	✓	✗	✓	✗	✗	✓
T5 Data Centers	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
TECfusions	✗	✗	✗	✗	✗	✗	✗	✗	✗	✗
Tencent	✓	✗	✓	✗	✓	✗	✓	✗	✗	✓
Vantage Data Centers	✓	✗	✓	✗	✓	✗	✗	✗	✗	✗

Source: Companies mentioned, Standard Chartered Bank,
Based on most recent publicly disclosed information in annual and sustainability reports

Size seems to matter when it comes to efficiency

The limited water intensity data currently available reveals a significant performance gap between data centre operators. Reported WUE levels range from more than 1.5 million litres per MWh to just 50,000 litres per MWh, suggesting that some operators use over 30 times more water per unit of computing activity than others (Figure 69).

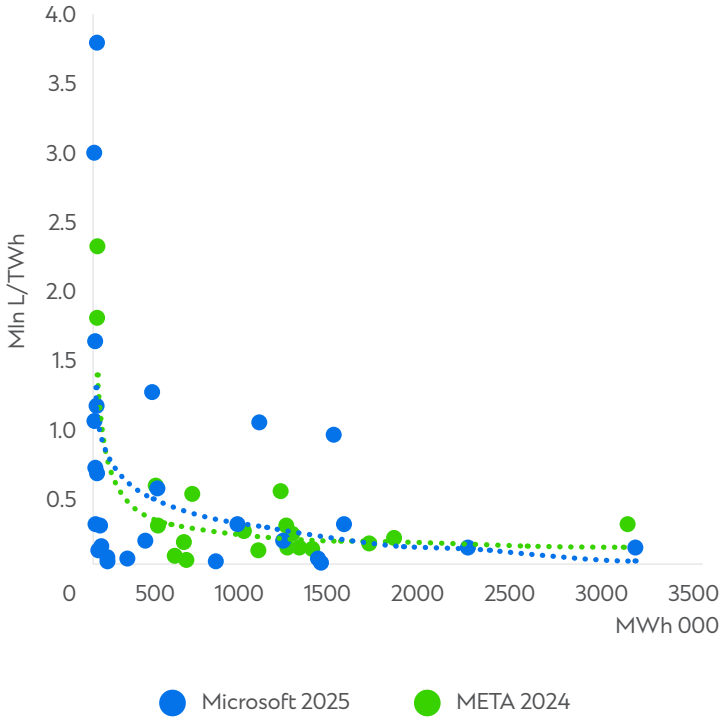
Reported data shows that of the larger data centre operators, Amazon, Meta and Microsoft rank among the most water efficient (Figure 69). The improvements they report in terms of water efficiency are noteworthy as average WUEs for Amazon, Meta and Microsoft have declined by 43-52% since 2020 (Figure 71).

META and Microsoft provide data for the individual facilities in their respective portfolios that imply that larger facilities tend to have superior water efficiency levels than smaller ones (Figure 70). One reason for this may be that larger data centres benefit from increased economies of scale which support investments in more efficient cooling systems, operational optimisation and water-management strategies which can deliver measurable reductions in water and power efficiency.

Figure 70

Water withdrawal intensity by data centre

Mln L/TWh for Microsoft (MSFT, 2025) and META (2024)

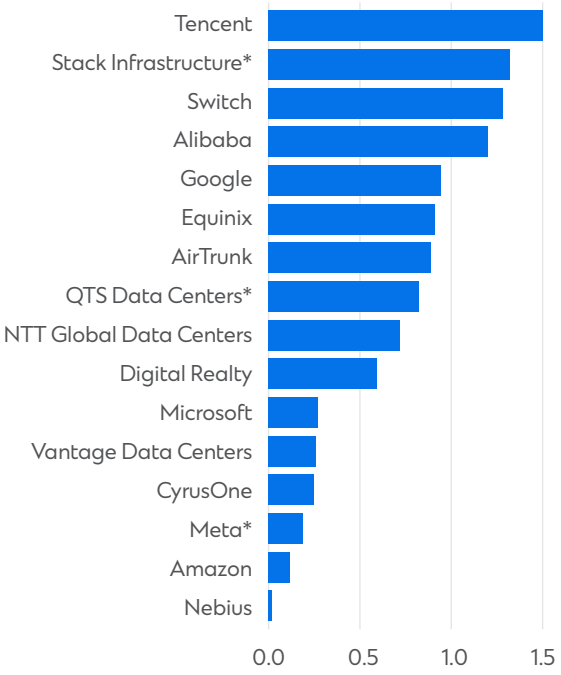


Source : Sustainability reports of companies mentioned, Standard Chartered Bank

Figure 69

WUE for key data centre operators

L million / MWh, 2025, * = 2024

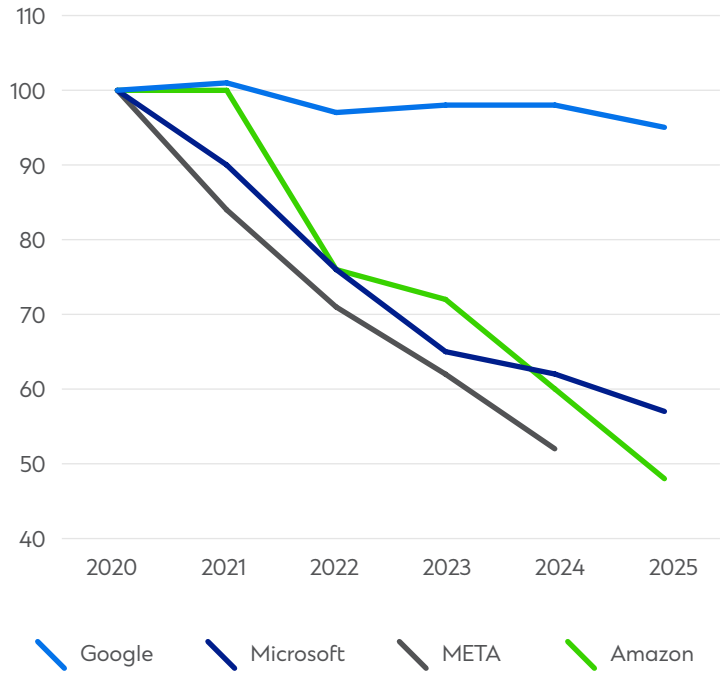


Source : Sustainability reports of companies mentioned, Standard Chartered Bank

Figure 71

Water efficiency of hyperscalers has strongly improved

Rebased to 2020 (=100)



Source: Sustainability reports of companies mentioned, Standard Chartered Bank

Solutions to reduce data centre water use further

As shown earlier in this report, estimates regularly suggest that water consumption by data centres is set to increase during the next few years as total capacity of the sector rises (e.g. (IEA, 2025), (GWi/Xylem, 2026), (Li & et al., 2025)). However, a range of available solutions and structural drivers suggest that further improvements in water intensity are likely.



Solution 1

Maximise use of dry and immersion cooling

While the largest data centre operators generally appear to achieve better water efficiency outcomes, Figure 67 shows that some smaller operators also rank among the sector's most water efficient. This suggests that scale is not the only determinant of water performance. A key differentiator is the ability to deploy cooling technologies with a low-water footprint.

For its sites, neocloud company Nebius uses direct free air cooling or closed-loop liquid cooling which results in very low-water intensity. The company has only been in operation a few years but for its owned and fully operational site in Finland, Nebius reported a 2025 WUE of 0.015 L/kWh (Nebius, 2025). CyrusOne provides another good example. Approximately 76% of its facilities use adiabatic or dry cooling with an associated WUE (category 1) of 0.00. This more than offsets the 21% of its data centre portfolio that uses evaporative or wet cooling with a WUE of 1.35 leaving the company with an overall WUE (category 1) of just 0.25.

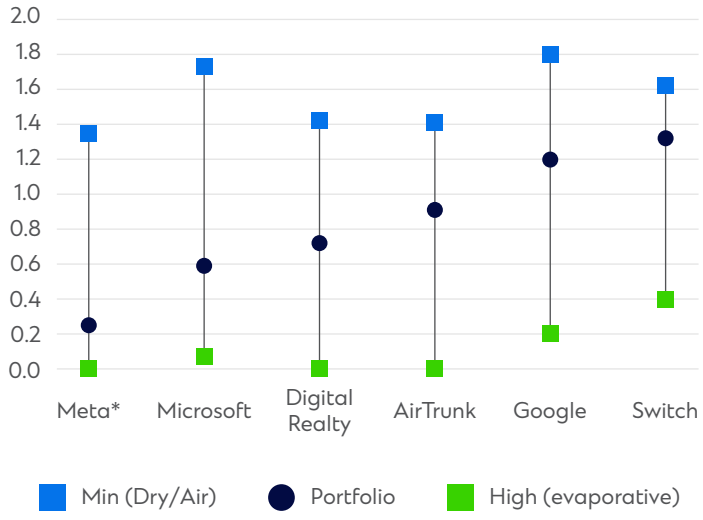
Figure 72 shows several other companies that report water efficiency ratios for different cooling technologies within their respective data centre portfolios as well as their overall WUE. It clearly supports the idea that data centres can materially improve their WUE if they are able to switch to dry or air-based cooling technologies all else being equal.

Liquid cooling technologies, including direct-to-chip and immersion cooling, offer even lower water intensity, with WUE levels that are close to zero (Figure 73). However, these technologies are still at an early stage of adoption across much of the industry. As a result, operators seeking to reduce water consumption today are often focused on increasing the use of adiabatic, dry cooling and free-air cooling solutions where climatic conditions allow. As liquid cooling solutions mature and become more widely available they have the potential to substantially reduce data centre water intensity and reduce the sector's overall water footprint.

Figure 72

The impact of cooling technologies

WUE of Air-based vs. evaporative cooling



Source: Sustainability reports of companies mentioned, Standard Chartered Bank

AI may accelerate the switch to low intensity water cooling solutions

The development of AI may well be a catalyst for the adoption of exactly those cooling solutions with the lowest water intensity. The reason for this is that traditional cooling technologies are not effective enough to cool AI data centres as these typically operate at much higher rack densities than traditional facilities (Figure 74). The processing units most often used in these facilities such as those supplied by NVIDIA require liquid cooling (Data Center Dynamics, 2026), (Al Kez & et al, 2025). As highlighted before in this report, these solutions have virtually zero water usage implying that a strong AI-driven expansion of the data centre sector is unlikely to show the proportional increase in water usage seen historically and may in some scenarios even coincide with a reduction of the sector's water footprint.

Retrofitting existing data centres may reduce water usage too

The review of key data centre operators shows that most have water-related strategies that focus on cooling solutions with low- or zero-water intensity. As existing data centres are upgraded to support AI workloads, operators may well replace existing cooling solutions with those that have low to zero water intensity. This can include better air management systems or upgrading mechanical systems with more efficient ones. Moves such as these would help contribute to a gradual decline in overall water intensity of the data centre sector.

The industry intends to move towards zero water use too

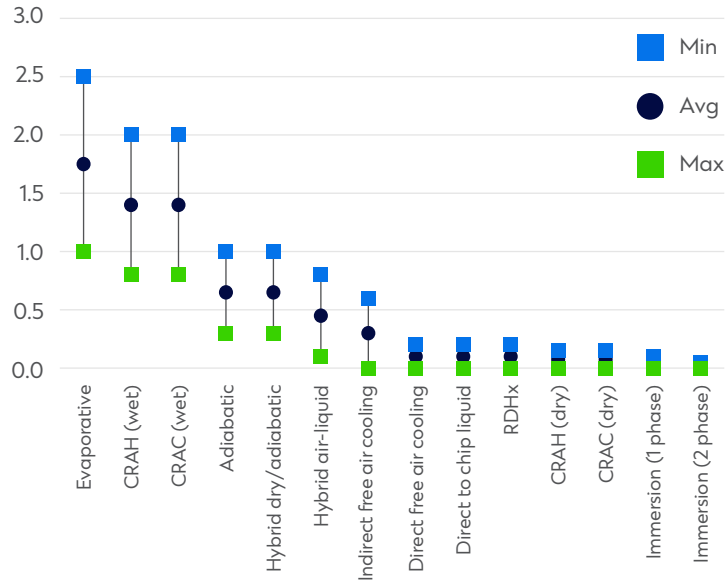
The above two points are strongly supported by various recent data centre operator surveys. Survey results from Data Center Dynamics among more than 100 data centre operators globally showed that 84% of operators either already use direct liquid cooling or are in some way planning or evaluating this. In addition, only 6% of operators suggested that they are currently not confident that liquid cooling would scale to support AI infrastructure during the next 3-5 years. (Data Center Dynamics, 2026). Survey work from Turner and Townsend meanwhile implies that 47% of surveyed data centre operators expect that high density liquid cooling will account for more than 50% of data centre capacity within the next two years (Figure 75).

A recent study from (Alissa & et al, 2025) showed what the impact of a shift from air cooling to liquid cooling could be on the environmental footprint of data centres. Using a life cycle assessment, they calculated that replacing air cooling with immersion cooling could reduce a data centre's water consumption by up to 48% when it is powered with grid electricity. They also found that water consumption could reduce more than 80% when combined with 100% renewable electricity rather than grid electricity. These findings suggest that data centre-related water consumption may well be substantially lower than is generally believed if, as operators believe, liquid cooling solutions are adopted at scale.

Figure 73

WUE ranges for different cooling technologies

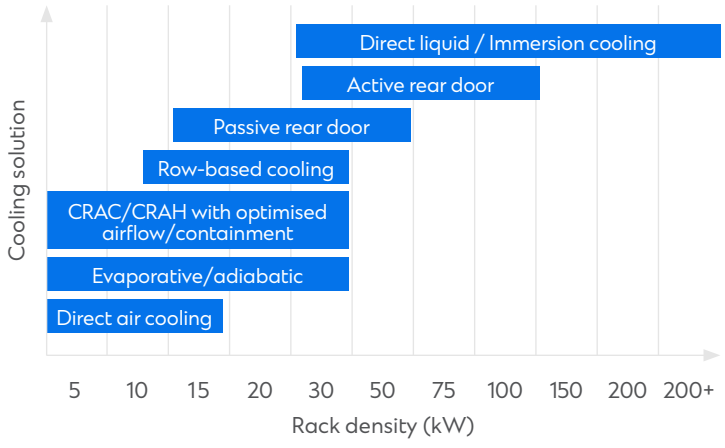
l/kWh



Source: Standard Chartered Bank

Figure 74

Applicability of DC cooling solutions for different power racks

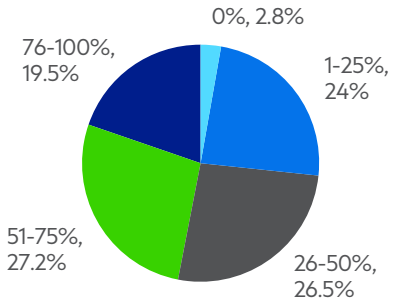


Source: Kez et al, Standard Chartered Bank

Figure 75

Share of high-density liquid cooling data centres

(% of total within 2 years)



Source: Turner and Townsend, Standard Chartered Bank

Improving water efficiency may be more profitable

Financial considerations may become a driver for the adoption of low-water cooling technologies. While traditional air cooling systems often have lower upfront installation costs, liquid cooling solutions can offer materially lower operating costs over time. These systems typically consume less electricity for cooling as reflected in their lower PUEs, while also requiring significantly less water as reflected in substantially lower WUEs. Additional benefits can include higher rack densities, which reduce space requirements and real estate costs and the useful life of equipment. Maintenance expenditures may also be lower, as conventional air cooling systems often require frequent filter replacements, whereas liquid cooling systems generally require only periodic coolant testing and infrequent pump maintenance (Introl, 2026).

The economic advantage of liquid cooling is also shown in Total Cost of Ownership (TCO) calculations for different cooling technologies. Work from Energy Solutions Intelligence suggests that the TCO of a single immersion system can be almost 30% lower than that of an air cooling system using rear door heating exchanges (Energy Solutions Intelligence, 2026). GRC suggests that a 39% reduction in TCO can be achieved by switching to immersion cooling (GRC cooling, 2026). The economic case in favour of liquid cooling appears to be strengthening with time as analysis from Schneider Electric suggested that liquid cooling only had a 14% CAPEX advantage in 2020 (Bunger & et al., 2020).

Finally, calculations from Ampcom and Energy Solutions Intelligence suggest that liquid cooling systems have a pay-back period relative to traditional air-cooled systems of around 3 years without incorporating any revenues that might be enabled by AI workloads that air cooling cannot support (Ampcom, 2026).

Taken together these calculations suggest that the power and water intensity of the data centre sector may well decline organically as operators recognise the financial benefits provided by cooling solutions with lower power and water intensity over traditional cooling sources.

Solution 2

Allow data centres to run at higher temperatures

Cooling-related electricity consumption has historically accounted for up to 40% of a data centre's total power demand (McKinsey, 2023). Some of this can possibly be linked to the impact that industry guidelines such as those developed by engineers at ASHRAE, the international society of heating, refrigerating and air-conditioning.

ASHRAE's TC 9.9 guideline outlines data centres' cooling, temperature and humidity handling (ASHRAE, 2026). Historically these guidelines were largely developed based on traditional air-cooled data centre environments and included a recommended temperature range for data centres of 18-27°C. Criticism from the data centre industry regarding ASHRAE's guidelines has been growing not least following the more recent development of liquid cooling solutions that can operate on much higher temperatures than the recommended range as suggested by ASHRAE (Green Calculus, 2026).

Case study: NVIDIA

Running AI data centres at 45°C has significant benefits

The impact that running a data centre on higher temperatures can have on its power and water consumption are significant as was shown by NVIDIA when it launched its DSX design platform for AI factories (NVIDIA, 2026c) and its Rubin liquid cooled AI platform (NVIDIA, 2026d).

Using the DSX reference design for AI factories NVIDIA's S45 architecture uses 100% liquid cooling based on a closed loop system. One of the key aspects of NVIDIA's design approach is that it allows its AI infrastructure to run at 45°C or substantially higher than the ASHRAE recommended range of 18-27°C. This according to the company eliminates the need for evaporative cooling water in facility heat rejection as the need for evaporative cooling water reduces (NVIDIA, 2026e). As a result, AI factories using this architecture can be deployed in water-stressed regions too without challenging with fresh-water availability. The design also significantly reduces power usage as the reduction of cooling overhead frees existing electrical capacity for direct AI use. Delaying cooling to much higher temperatures is also financially relevant as each degree of temperature increase can cut cooling-related energy costs by 4-5% (Energy Star, 2026). NVIDIA's calculations suggest that letting data centres run warmer reduces the need for mechanical or power and/or water intensive cooling significantly across most parts of the US (Figure 76).

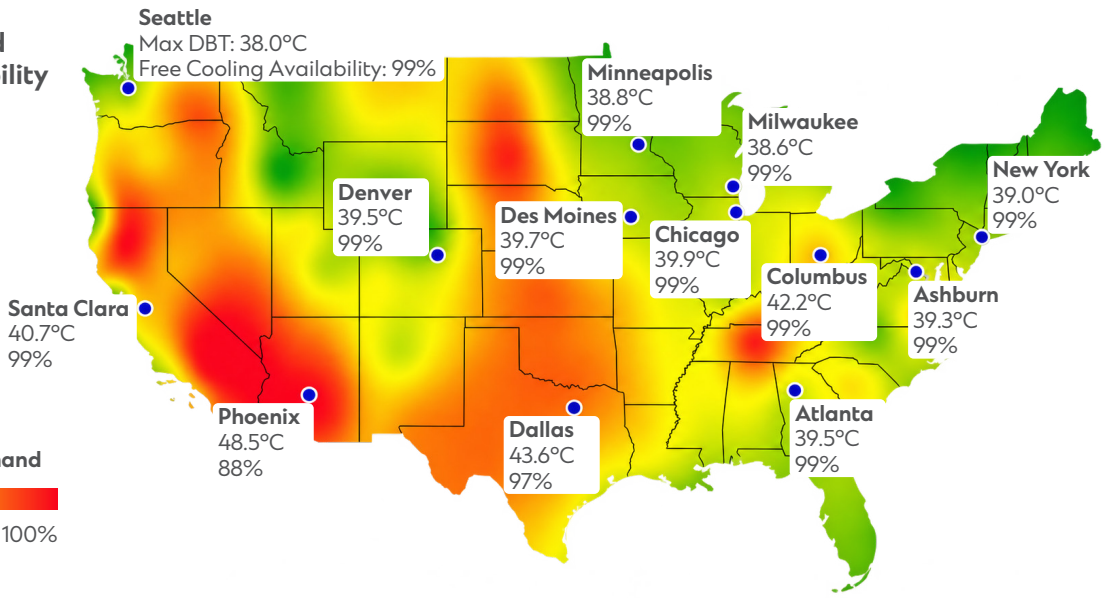
The benefits of running data centres at higher temperatures not only include a reduction in power and water usage but also lower noise reduction as cooling fans, responsible for much of a data centre's noise production, are used far less to not at all.

Another benefit of running AI factories at higher temperatures is the greater potential for waste heat recovery. This waste heat may help lower energy requirements across local communities if it is used for district heating or commercial and industrial facilities. Furthermore, integrating this waste heat into a data centre's on-site power generating facility could create a closed loop heating and power system with potentially additional environmental benefits especially if small modular reactors or fuel cells are part of this too.

Operating data centres at higher temperatures does come with some potential risks such as a reduced thermal safety margin and the potential for accelerated wear on certain components. While higher operating temperatures can improve energy efficiency and reduce cooling demand, operators must balance these benefits against potential impacts on reliability, resilience and long-term infrastructure performance.

Figure 76

NVIDIA's 45-degree solution and implied free cooling availability across the US



Source: NVIDIA

Note: Percentages show the share of hours in a typical meteorological year when dry coolers alone can meet heat rejection (45°C rack supply, 4°C CDU approach, 5°C dry-cooler approach, no adiabatic assist). The map shows the mechanical cooling capacity required at the design dry-bulb condition (20-year extreme maximum, ASHRAE 2021), as a share of the liquid-cooled load. Liquid loop only.

Solution 3

Integrate circularity solutions

Data centre operators can improve their direct environmental footprint not only through the adoption of new, more efficient technologies or services but also by incorporating circular economy solutions. These help to reduce resource consumption, lower waste, cut lifecycle emissions and can ease pressure on local ecosystems which can help strengthen operators' engagement with local communities and regulators. The circular solutions that data centre operators can adopt are broadly part of two categories:

- **Develop alternative water sources:** Data centres that do not use low or near-zero water cooling solutions typically rely on potable water. This may create supply shortages and tensions with local communities. Various circular water solutions exist that help data centres switch part of their potable water use to using recycled water. Sources for this can be internal through recycling existing cooling water or external using treated municipal or industrial wastewater or the use of rainwater or effluent. Figure 77 shows some circular water solutions and companies that have a strategy for these. It shows that the data centre sector is starting to get engaged with these solutions, however, in most cases these solutions are used at a relatively small scale suggesting that significant upside potential to reduce direct water consumption remains.
- **Recycle materials and equipment:** Limited data is disclosed by the level of embedded water consumption from data centres in relation to purchased IT equipment or building and construction materials. This may not generally impact local water supply conditions but can be a factor on a more national level or in regions elsewhere where this equipment or materials are produced. Re-use and refurbishment of IT equipment and recycling of building and construction materials centres can adopt to reduce not only embodied carbon or emissions but also their 'Scope-3' water consumption.

Figure 77

Alternative circular water solutions for data centre operators

Source	Application	Advantages	Key constraints	Companies exposed
Circular water-related solutions				
Treated municipal wastewater	Used for cooling-tower or evaporative-cooling system	Most scalable alternative to potable water; AWS plans expansion to 120+ U.S. locations, preserving an estimated 530mn gallons of drinking water annually	Requires dedicated "purple-pipe" infrastructure, pretreatment and reliable municipal supply	Aitrunk, Amazon, Beale, CyrusOne, Digital Realty, Equinix, Google, Microsoft, QTS, Stack, Switch, Vantage
Industrial process wastewater	General cooling or internal use after treatment	Useful for facilities located in industrial parks; converts a neighbouring facility's waste stream into a resource	Variable chemistry, contamination and supply; potentially substantial treatment requirements	Airtrunk
Recycled cooling-tower blowdown	Internal reuse as make-up water	Reduces both withdrawals and wastewater discharge; enables higher internal water circularity	Dissolved solids become concentrated, often requiring filtration, reverse osmosis or other treatment	Aitrunk, Aligned, Meta, Microsoft, NTT, Oracle, Switch, Tencent
Harvested rainwater and stormwater	Supplemental cooling, landscaping and non-potable building uses	Lowers demand on municipal water and helps manage site runoff	Seasonal and intermittent; requires storage, filtration and sufficient roof/ catchment area. Amazon explicitly includes harvested rainwater in its reuse strategy	Amazon, CyrusOne, Digital Realty, Meta, Microsoft, QTS, Stack, Tencent

Source	Application	Advantages	Key constraints	Companies exposed
Circular water-related solutions				
Greywater and onsite treated sewage	Cooling make-up, toilet flushing and irrigation	Creates a local circular supply; suitable for campuses with offices or other occupied buildings.	Data centres alone produce relatively little greywater; treatment and health controls are required.	Amazon
Seawater or brackish water	Once-through or closed-loop heat exchange	Avoids freshwater withdrawal and can reduce chiller electricity. Google uses Bay of Finland seawater at its Hamina site.	Limited to coastal sites; corrosion, marine impacts, intake permitting and discharge-temperature controls.	Equinix, Google
HVAC condensate	Supplemental cooling-tower make-up	Typically low in dissolved minerals and generated during humid conditions when cooling demand is high.	Volumes may be small or seasonal; collection plumbing and biological controls are needed.	
Mine water, quarry lakes or deep surface-water reservoirs	Heat sink through closed-loop heat exchangers	Can provide stable low-temperature cooling without consuming the source water.	Early-stage and location-specific. Research estimates one Pennsylvania mine pool could support at least 9.5MW of conventional cooling or 30.6MW with liquid cooling.	
Circular equipment, material and heat solutions				
Re-use and refurbishment of equipment	Extend the useful life of servers, racks, drives, network equipment and components through repair, refurbishment, redeployment, resale and parts harvesting.	Reduces e-waste, virgin-material demand, embodied carbon and replacement CAPEX while improving supply-chain resilience.	Requires secure data sanitisation, component testing, traceability, compatibility with newer systems and reliable secondary/recycling markets has yet to be developed.	Aligned, Amazon, Apple, Google, Meta, Microsoft, Oracle
Use circular construction materials	Use recycled steel and metals, lower-carbon concrete/cement substitutes, reused structures, recycled plastics, modular components and materials designed for recovery.	Reduces embodied carbon and extraction of virgin materials; can lower construction waste and preserve value through future reuse or recycling.	Availability and certification vary by region; performance, warranty, fire/safety standards, cost premiums, quality and supplier traceability must be managed.	Amazon, Equinix, Meta, Microsoft, NTT Global Data Centers, STACK Infrastructure
Waste heat recovery / heat reuse	Capture server and cooling-system heat using heat exchangers; raise its temperature with heat pumps where needed and supply the heat to district-heating networks, nearby buildings, greenhouses or other thermal users.	Turns rejected heat into a useful energy stream, reduces fossil-fuel heating demand and community emissions, improves overall energy productivity and can strengthen local acceptance of data-centre development.	Requires a nearby, dependable year-round heat customer and distribution network; low-grade heat often needs heat pumps. Economics depend on distance, heat temperature, utilisation, ownership and long-term contracts.	Amazon, CyrusOne, Digital Realty, Equinix, Google, Microsoft, NTT Global Data Centers, Vantage Datacenters

Source: Companies mentioned, Standard Chartered Bank

Solution 4

Maximise the use of renewable power

The previous chapter discussed how clean, or low emission, power can be a strategy that eases power-related challenges created by the expansion of data centre capacity. The same strategy also materially reduces a data centre's water footprint. Figures 55 and 56 provided more context to this argument. Figure 55 showed that the majority of a data centre's water footprint is associated with the water intensity of its power consumption. Figure 56 showed that to generate electricity through solar panels or wind turbines requires virtually no water at all whereas other power sources such as coal are much more water intensive.

Using the data from Figures 55 and 56 a data centre operator could theoretically reduce its overall water footprint by close to 50% if it were able to switch its power source fully to renewable electricity. A full reduction like this is unlikely to occur in the short term as most data centres remain connected to the local grid and therefore depend on the speed at which their local utility companies reduce the water intensity of their power mix. The largest data centre operators may have the scale to more aggressively expand their own renewable power capabilities and help reduce their indirect water footprint. This would also require investment in energy storage capacity given the intermittency that solar and wind power have.

Solution 5

Develop water replenishment programmes close to data centre locations

The solutions thus far have focused on improving a data centre's direct water use, through the use of low-water consuming cooling solutions, or indirect water use by switching the power mix towards renewable source. However, to minimise the overall water and ecosystem impact of data centre growth, helping to secure support from local communities and reduce permitting and project-delay risks likely require a broad stewardship approach that focuses on active engagement with the communities, utilities and local authorities in the regions where data centres are built and operated.



Water replenishment can meaningfully reduce water intensity

Recognising this challenge, data centre operators, particularly the larger hyperscalers, are increasingly investing in water initiatives to reduce their net impact on local freshwater availability. These efforts not only include expanding the use of recycled and non-potable water but also focus on developing water-replenishment projects and supporting ecosystem restoration initiatives in local communities.

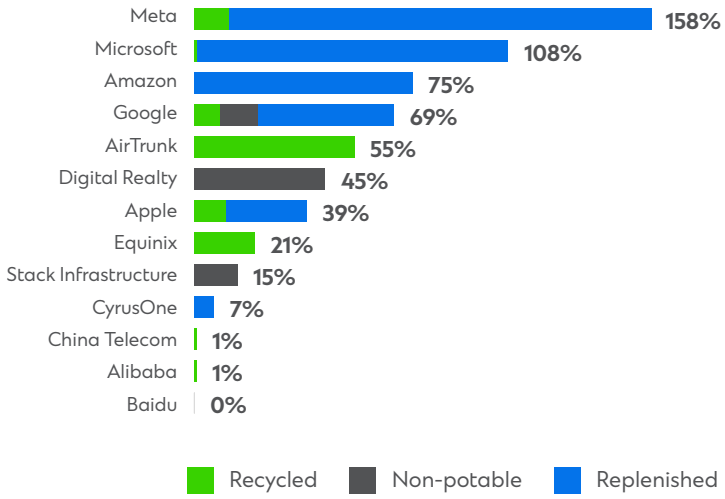
The latest available data for a range of data centre operators suggests that, for several leading operators, recycled water, replenished water and non-potable water sources already account for a meaningful share of total water withdrawals (Figure 78). Some of the largest data centre operators now replenish more water than their data centres use (Figure 79).

The scope of the replenishment programmes established by some of the main data centre operators is outlined in Figure 80. It suggests that the annual level of water replenishment when these programmes are fully developed is expected to be well above recent total water consumption levels. Water conservation therefore can be an important component of data centre sustainability strategies with the potential to materially reduce the sector's long-term impact on freshwater resources.

Figure 78

Data centre water conservation

% of withdrawals for most recent financial reporting year

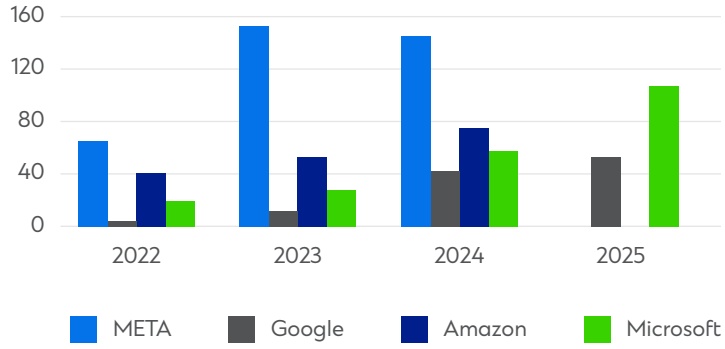


Source: Sustainability reports of Companies mentioned, Standard Chartered Bank

Figure 79

Replenished water volumes

As % of annual water withdrawal



Source: Sustainability reports of companies mentioned, Standard Chartered Bank

Figure 80

Replenishment details for selected data centre operators

Operator	Water consumption (bn litres)	Projects	Countries	Annual water replenished (bn litres)	Target annual replenishment (bn litres)
Amazon	9.5	58	16	9.4	18
Apple	6.8	9	1	3.0	6.8
Google	40	165	21	29	75
Meta	3.0	84	3	6.8	12
Microsoft	8.2	111	20	14.3	133*

Source: Companies mentioned, Standard Chartered Bank
* total replenishment over the lifetime of the projects rather than an annual maximum level of replenishment

Case study: Microsoft

Microsoft's approach to water management is partly embedded within its Datacenter Community Pledge (Microsoft, 2024), launched in 2024 and its US Community-First AI Infrastructure plan launched in 2026 (Microsoft, 2026a).

The Datacenter Community Pledge is Microsoft's long-standing commitment focused on general cloud and data centre operations. It includes design and operational aspects for data centres aimed at supporting society's climate goals. Earlier this year Microsoft expanded its Datacenter Community Pledge with the introduction of its Community-First AI Infrastructure initiative. This initiative consists of 5 key commitments or promised to each of the communities where the company builds, owns and operates data centres (Figure 81). One of the five commitments relates to water specifically.

In relation to water Microsoft's targets include the reduction of water usage per data centre. Where possible and, dependent on local conditions and specifics such as climate or operational needs, the company is using closed-loop liquid cooling, chip-level cooling through the use of microfluidics and zonal cooling as some of the solutions that can help reduce the need for potable water for cooling.

Figure 81

Microsoft's Community-First AI pledges

Electricity	**We'll pay our way to ensure our datacenters do not increase your electricity prices**
	Pay utility rates that are high enough to cover our electricity costs Collaborate with local utilities on plans to add to the electricity we will need Innovate to make our data centres more efficient Advocate for public policies needed for affordable, reliable and sustainable power
Water	**We'll minimise our water use and replenish more of your water than we use**
	Reduce the amount of water our datacentres use Replenish more water than we use Provide greater local transparency Advocate for public policy that helps minimise water use
Jobs	**We'll create jobs for your residents**
	Invest in partnerships to train local construction workers Expand our Datacenter Academy program to train more individuals for ongoing operational roles Encourage policymakers to support new job opportunities
Taxes	**We'll add to the tax base that funds local hospitals, schools, parks and libraries**
	We won't ask municipalities to reduce their local property tax rates for datacenters We'll support policies to invest the added taxes we pay in the vital services the local community cares about
Training/support	**We'll strengthen your community by investing in local AI training and non-profits**
	Partner with schools, community colleges and universities to provide AI training Support adults with AI tools and skills through AI learning hubs in local libraries Support AI skills training for businesses Invest in local non-profits

Source: Microsoft Building Community-First AI Infrastructure, Standard Chartered Bank

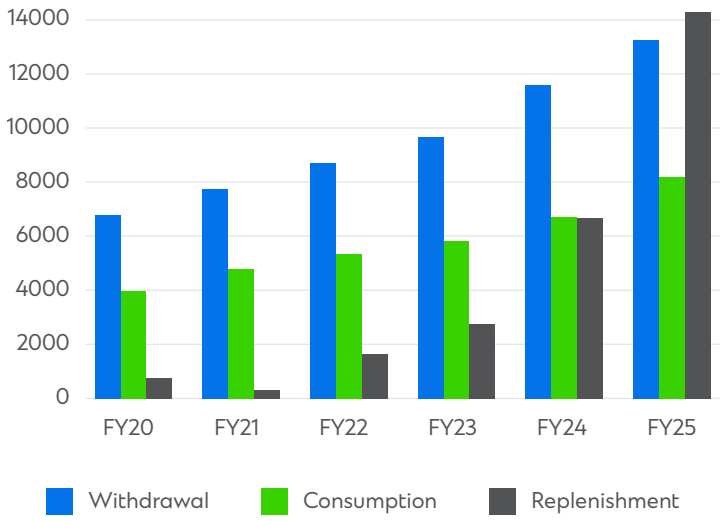
To further reduce water constraints Microsoft collaborates with local utility companies to find solutions to avoid adding further pressure on water resources. One example of this is the Quincy Water Reuse Utility that the company constructed in Quincy Washington that treats and recirculates data centre cooling water which in turn limits the use of potable or drinking water for cooling purposes. Efforts such as these has allowed Microsoft to reduce the use of potable water across some of its largest data centres significantly. In Quincy, San Antonio and Singapore only 26%, 24% and 1% of total water withdrawal in 2025 was potable water. For its new data centre in the Amsterdam region the company has installed a full-scale rainwater harvesting system which is expected to collect and store more than three times the water annually required for cooling the data centre (Microsoft, 2026b). To ensure the resilience of local water systems, Microsoft fully funds the infrastructure buildout that is needed to support its data centres, it invests in community water infrastructure and helps modernise local water systems to expand access and water reliability.

Water replenishment forms another key part of Microsoft commitments to local communities and the company's goal is to replenish more water than it withdraws by 2030. The company has a replenishment project portfolio consisting of more than 100 initiatives with a total replenishment of more than 133 million m3 of water over their lifetime (Microsoft, 2026c). Several projects already replenish more water per year than the data centres located in those areas use (Figure 83). Microsoft's progress in terms of replenishment is clear as shown by the annual level of replenishment globally since its fiscal year of 2020 (Figure 82).

Figure 82

Microsoft total replenishment per year

M³ 1000

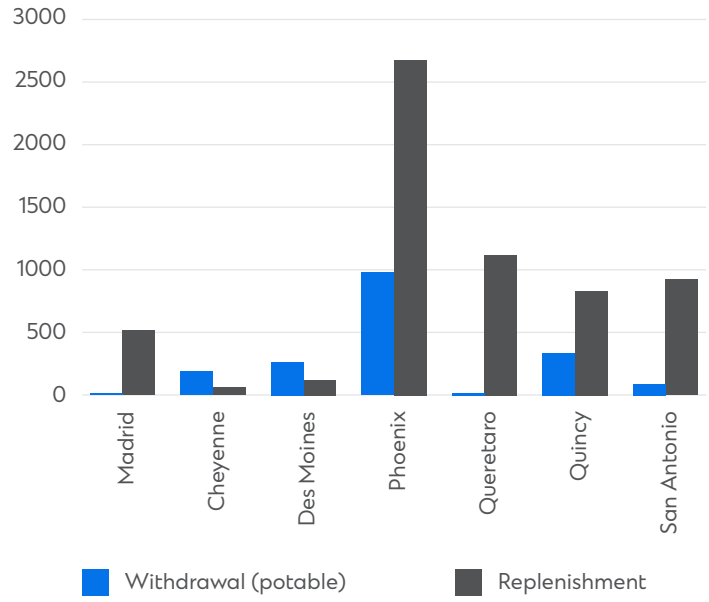


Source: 2025 Sustainability report Microsoft, Standard Chartered Bank

Figure 83

Replenishment by data centre location

ML (Mega litres)



Source: 2025 Sustainability report Microsoft, Standard Chartered Bank

AI as a tool to address water leakage

Water leakage is one of the largest and most overlooked inefficiencies in global freshwater systems. An estimated 30% of treated water is lost before it reaches customers, primarily due to ageing infrastructure, undetected leaks, and poor network management. In some poorly maintained water networks, losses can exceed 50-60%, highlighting the scale of the challenge facing utilities worldwide (World Water Tech, 2025). The problem is not confined to developing economies. In the United States, water utilities lose approximately 20% of treated drinking water on average, while leakage rates in parts of Europe remain significantly above best-practice levels despite decades of investment in water infrastructure.



Artificial intelligence is emerging as a powerful tool to address this challenge. By analysing data from pressure sensors, flow meters and acoustic monitoring devices, AI systems can identify hidden leaks, estimate their size and pinpoint their location more quickly and accurately than traditional inspection methods. FIDO AI, for example, uses acoustic signatures collected from water networks to detect underground leaks that may never become visible at the surface. Microsoft has partnered with FIDO Tech as part of its water replenishment programme, supporting deployments across water networks in London, Phoenix, Las Vegas and Querétaro (Mexico) to help utilities reduce water losses and improve water security (Microsoft, 2025).

Several leading technology companies are now investing in AI-driven leakage reduction as part of their water stewardship strategies. In 2025, Meta partnered with FIDO Tech to fund the deployment of AI leak detection across 300 km of water infrastructure in Farmington, New Mexico, within the water-stressed Colorado River Basin (FIDO, 2025). The initiative aims to reduce non-revenue water, strengthen local water resilience and contribute to Meta's commitment to become water positive by 2030. Microsoft is also a founding partner of the broader Water United collaboration, which seeks to scale AI-enabled leakage reduction across the Colorado River Basin and other water-stressed regions.

These initiatives demonstrate how AI can create environmental benefits beyond the data centre itself and help utilities recover millions of litres of treated water that would otherwise be lost through leakage. In regions facing increasing water stress, AI-enabled leak detection could become one of the most cost-effective and rapidly deployable tools for improving water efficiency and strengthening long-term water security.

Solution 6

Regulation: a potential external driver?

The rapid expansion of data centres has triggered growing levels of political and public scrutiny as communities seek greater insight into their broader impact. Greater regulatory oversight may be needed to establish robust, transparent and comparable metrics for assessing the environmental and social impacts of the rapidly expanding data centre sector and help reduce its water footprint further. A more developed regulatory framework could also help align the interests of key stakeholders, including governments, investors, financiers, operators and local communities. The current policy landscape for data centres appears fragmented, reflecting the relatively early stage of the sector's development. Recent policy initiatives in the European Union may act as a catalyst for the development of more comprehensive data centre regulation globally.

Current policy landscape is a work in progress

Clear disclosure requirements, consistent efficiency metrics and transparent thresholds for power and water performance give governments, communities, investors and operators a strong basis for assessing the impact of data centre development. Well-designed regulation can also reduce permitting uncertainty by clarifying what is expected from operators, including requirements around grid access, renewable power procurement, water efficiency, reuse and local resource impacts. Such regulation could therefore help improve public confidence and create support for approval of new projects and encourage investment.

The current policy landscape around data centres, however, does not appear to be clear or consistent and is best described as a 'work in progress' as governments seek to build their knowledge regarding data centres and their potential impact on societies. The most recent example of this is the consultation by the Australian government on data centres which closed on 1 September [2026] (Australian Government Department of the Prime Minister and Cabinet, 2026). Policies that have been established appear to be more reactive such as setting moratoriums on new data centre development as has happened in different countries including Singapore, Ireland, Thailand and several states in the US. Limited requirements have so far been put in place by most governments while thresholds or targets that have been set tend to vary between countries (Figure 84).

A review of current policies across some of the main data centre-related countries shows that:

- The European Union's Energy Efficiency Directive requires any data centre with installed IT power demand of more than 500kW to annually report 24 data points across energy, sustainability metrics, ICT capacity and data traffic (European Commission, 2026). The EU then calculates PUE, WUE, Energy Reuse Factor and the Renewable Energy Factor. Mandatory reporting requirements elsewhere are very limited compared to those implemented in the EU.
- Where efficiency policies have been established, governments have tended to focus solely on PUE-related targets (e.g. China and Singapore).
- Except for China, none of the main data centre-related jurisdictions have incorporated water performance metrics or targets. The EU is currently in the process of establishing a rating scheme for data centres which is expected to include targets for water efficiency too. The government of Singapore has discussed setting water-related targets for data centres but that has not yet been implemented as a universal requirement.
- Several Asian countries that may experience some of the highest data centre growth rates during the next few years have not yet developed significant data centre-related policies and requirements.
- While the US has very limited data centre related policies on a federal level there are various state level laws and regulations that are relevant. In 2025, over 40 states considered 267 related bills. Legislative activity has continued into 2026, with energy, water, zoning, and tax issues all active across multiple states. Energy concerns remain the dominant issue, with states weighing how to provide sufficient power to data centres while protecting residential ratepayers. South Carolina and Maryland enacted rate negotiation laws in 2025 with similar proposals moving in additional states in 2026. Water usage legislation is gaining momentum

in 2026, with states increasingly requiring data centres to disclose and manage water consumption. Minnesota established a separate water permitting requirement for data centres in 2025, and more states are following suit. Local regulation is emerging as a significant battleground. State-level moratorium bills have been introduced in 11 states in 2026 but face resistance; meanwhile, dozens of municipalities have moved ahead with local construction pauses, shifting the centre of gravity towards local government action. Tax incentives for data centres are under renewed scrutiny. States that previously offered generous incentives are reassessing their value as construction costs soar and community opposition grows.

Can the EU's rating scheme for data centres become the standard?

One of the most significant policy developments aimed at addressing the disclosure challenges outlined earlier is the European Union's Data Centre Sustainability Rating Scheme. The initiative builds on the revised Energy Efficiency Directive (EED), which introduced mandatory reporting requirements for data centres with an installed IT power demand of more than 500kW (EU, 2026). This scheme sets out key performance indicators, the methodology to measure these and requirements in terms of how often these are made publicly available. Following further publications and initial assessments in 2024 (EU, 2024) and 2025, a supplement directive was issued in March 2026 with the aim to improve some of the reporting aspects for data centres (European Commission, 2026).

The proposed rating scheme introduces a unified and transparent approach to disclosure for data centres operators in Europe. They will report standardised information on metrics such as PUE and WUE, enabling more consistent comparisons between facilities. Much of the reported information is expected to be made available through a centralised public database, including details on whether the reported WUE values reflect only direct water consumption or also incorporate the water footprint associated with electricity generation. The scheme is also intended to support the development of sustainability ratings that assess the relative energy and water performance of data centres.

Adoption of the proposed legislation by the Commission is planned for the third quarter of 2026 although the Climate Neutral Data Centre Pact (CNDCP) has raised concerns about aspects of the methodology (CNDCP, 2026). The CNDCP is a voluntary initiative that has over 100 data centre members plus representative bodies that represent 85% of the European data centre market. One of their key concerns relate to the fact that the EU's proposal does not take climate conditions of data centres into consideration. As highlighted in the previous section such an adjustment would acknowledge that current cooling and power usage requirements for data centres differ depending on local climate conditions.

Aside from the concerns raised by the CNDCP, adoption of the EU's data centre rating scheme creates a transparent, harmonised and publicly available dataset. For policymakers, investors and financiers such a dataset would improve their ability to access the potential impact of data centres.



Figure 84

Summary of data centre-related policies and approaches for key markets (non-exhaustive)

Category	China	European Union	United States	UK
Overall approach	Direct performance regulation with explicit efficiency targets	Mandatory disclosure and transparency regime	Limited federal regulation; state permitting and emerging water scrutiny	Critical National Infrastructure (CNI)-led approach combining planning reform, energy-system integration, sustainability oversight and resilience requirements.
Key policy instruments	Green Development Action Plan for Data Centres East Data West Computing programme Green Data Centre standards	Energy Efficiency Directive (EU 2023/1791) Delegated Regulation 2024/1364 EU rating scheme	DCOI/DCCOI (federal facilities) Energy Act 2020 state permitting processes	Critical National Infrastructure designation National Planning Policy Framework Nationally Significant Infrastructure Projects (NSIP) regime Forthcoming Cyber Resilience Bill Environmental and water policy reviews
National mandatory PUE target	Yes	Operators must track PUE but no EU-wide threshold yet	No federal PUE requirement for private-sector facilities	No
Typical PUE requirement for new facilities	< 1.5 nationwide for all <1.25 for large and mega DCs <1.2 for national hubs <1.3 government procurement	None at EU level	None at federal level	No
Mandatory WUE requirement	Limited but emerging; government procurement standard includes water intensity metric	No EU-wide WUE threshold yet	No federal WUE requirement	No
Explicit water performance metric	Water intensity <2.5L/kWh in government procurement standard	WUE reporting required but a threshold has not yet been set	No	No
Water-use reporting requirement	Limited currently, though water efficiency is increasingly incorporated into approvals	Mandatory reporting of water use and WUE for facilities ≥500kW IT load	No federal reporting requirement; some states are considering water reporting legislation	No
Renewable energy requirements	Renewable energy utilisation in data centres to increase. Government targets 50% by '27, 80% by '30	Renewable Energy Factor (REF) reported and benchmarked	No federal renewable-energy requirement specific to data centres	No
Mandatory sustainability disclosure	Limited	Yes globally through EU Data Centre Database	Generally no nationwide requirement specific to data centres	No
Use of sustainability rating system	Green Data Centre certification programme	Common EU sustainability rating scheme under development	Primarily voluntary certification schemes (e.g. ENERGY STAR and industry standards)	Primarily voluntary. Climate Neutral Data Centre Pact, European Code of Conduct for Data Centres, LEED, BREEAM commonly referenced.

Category	Singapore	India	Indonesia	Thailand	UAE
Overall approach	Capacity-allocation and certification-led, moving toward mandatory regulation	Incentive-driven, limited direct performance regulation	Investment-led green growth strategy	Renewable-energy and investment-led approach	Sustainability-framework-led approach
Key policy instruments	Green Data Centre Roadmap; DC-CFA programme BCA-IMDA Green Mark proposed Digital Infrastructure Act	Renewable Purchase Obligations (RPOs) Green Energy Open Access Rules state DC policies	Making Indonesia 4.0 PPP Green Data Centre initiatives PLN renewable programmes	BOI incentives Clean Energy Policy Direct Renewable PPA framework	Sustainable Digital Transformation Framework Net Zero 2050 Strategy
National mandatory PUE target	Proposed mandatory PUE requirements under the forthcoming Digital Infrastructure Act. Specific thresholds not yet enacted.	No	No	No	No
Typical PUE requirement for new facilities	PUE ≤1.25 for DC-CFA allocated new projects. Not a universal requirement	Some procurement programmes have used PUE <1.35 criteria	Market-driven rather than regulation-driven	BOI incentives favour high-efficiency facilities but no national threshold identified	Market-driven, sustainability-focused
Mandatory WUE requirement	No	No	No	No	No
Explicit water performance metric	Government objective discussed of WUE ≤2.0 m^3/MWh, but not identified as mandatory	No	No	No	No
Water-use reporting requirement	Large water users must submit Water Efficiency Management Plans	No	No	No	No
Renewable energy requirements	Strong preference for green power and low-carbon energy under DC-CFA process	Renewable Power Obligation rises to 43.33% by FY2029-30; Open Access Rules support direct renewable procurement	Policy support for renewable powered data centres but no specific mandatory renewable-energy requirement	Direct PPA programme allows renewable procurement; some programmes require 100% renewable sourcing	Green data centres explicitly encouraged under sustainability framework
Mandatory sustainability disclosure	Limited	Limited	No	No	Limited
Use of sustainability rating system	Green Mark for Data Centres certification	Emerging IGBC Green Data Centre certification framework	Emerging focus on sustainable and green data-centre development	BOI incentive framework rewards efficient facilities	Sustainability embedded in digital-transformation policy

Source: Standard Chartered Bank

Building-labelling schemes can provide additional insights

Where regulation or policies are insufficient building-related labelling schemes can perhaps be an alternative driver for the adoption of energy and water efficient technologies by data centres. By incorporating PUE, WUE, renewable energy sources, water reuse and replenishment strategies these schemes can help create benchmarks for operators, investors and regulators if regulatory requirements fail to achieve this.

Some of the building-related labelling schemes that are in operation are listed in Figure 85. Most of these use metrics that are relevant for data centres including PUE, renewable energy use or cooling system efficiency. As a water intensity measure, WUE is also included in several of these schemes including those used in Singapore, Germany and China. However, the more prominent schemes such as the LEED and Energy Star have yet to include more specific water efficiency metrics such as WUE.

Figure 85

Labelling policies that impact data centres for key markets

Country/ jurisdiction	Scheme name	Nature	Primary metrics	WUE included?
Australia	NABERS Data Centre	Voluntary (mandatory for government contracts)	PUE, GHG emissions	No
Austria	Ecolabel UZ80	Voluntary	PUE, ERG, Energy efficiency, Renewable energy use	Yes
China	Data Center Green Rating Assessment	Government sponsored performance award	PUE, Energy efficiency, green management including water utilisation	Yes
EU	EU Code of Conduct for Data Centre Energy Efficiency	Voluntary	PUE, ITE energy efficiency, IEE energy	No
Germany	Blue Angel / PEER-DC	Voluntary (mandatory for federal IT contracts)	PUE, ERG, Energy efficiency, Renewable energy use	Yes
Hong Kong	BEAM Plus Data Centres	Voluntary	Energy use, sustainability of site, materials and waste, energy use	Yes
Malaysia	Green Building Index Data Centre Tool	Voluntary	PUE	No
Singapore	BCA-IMDA Green Mark for Data Centres	Voluntary	PUE, Cooling system efficiency, Energy management	Yes
South Korea	Green Data Certification	Voluntary	PUE, Green practices, Technology excellence	No
Switzerland	SDEA Label	Voluntary	PUE, Carbon footprint, Heat recovery	No
United States / Global	LEED BD+C: Data Centers; LEED O+M	Voluntary green-building certification	Energy efficiency; building and advanced energy metering; cooling/ system efficiency; renewable energy procurement. PUE and WUE not universal LEED threshold metrics	No
USA	Energy Star	Voluntary	PUE	No

Source: Standard Chartered Bank

Strong operator targets and disclosures as catalyst for improvement

Notwithstanding the relatively underdeveloped regulatory landscape and the absence of comprehensive standardised disclosure requirements, most of the leading data centre operators have established targets associated with power and water intensity. Figure 86 provides an overview of some of the water-related targets and plans adopted by different data centre operators. Almost all are pursuing lower water intensity through more efficient technologies and operating practices although not all have publicly committed to reducing net water consumption to zero or disclosed a timeline for doing so. While data centre operators are starting to develop strategies focused on reducing the power and water intensity of their facilities it is worth pointing out that other areas of concerns including noise and air pollution and land take may remain unaddressed.

As public scrutiny, resource constraints and requirements for permitting increase, data centre operators have an opportunity to move the sector onto the front foot by proactively adopting transparent disclosure frameworks and measurable sustainability commitments. Doing so can help build trust with regulators, investors and local communities before concerns over power, water and environmental impacts become barriers to growth. Figure 85 shows examples of disclosure elements that could support this objective and strengthen operators' long-term license to operate.

Data centre operators may find some of the disclosure suggestions included in Figure 87 difficult to accept. This is likely to be the case with those related to individual facilities as these may trigger security issues or concerns around customer confidentiality or operational sensitivity. One potential solution could be to develop public reporting requirements and reporting requirements for regulators or appropriate authorities.

The industry is reacting: AI Infrastructure Coalition

The level of public scrutiny related to the topic of data centres has not gone unnoticed within the industry. On September 30, the AI Infrastructure Coalition (AICC) which represents companies across the AI supply chain including semiconductor manufacturers, energy providers, hyperscalers, data centre operators and AI model developers announced five Community Principles that it said will guide data centre development for its member across the country (AI Infrastructure Coalition, 2026). These principles are:

- **Protect ratepayers: cover our full energy costs:** Data centres will bear the full cost of the energy they use and will work with grid operators to contribute to a more reliable grid.
- **Benefit the community:** Data centre operators will help fund local priorities including schools, roads and hospitals to help support local communities and economies.
- **Minimize water usage and act as responsible environmental stewards:** Data centres will minimise water use and protect local water supplies by using advanced cooling solutions, investing in local water infrastructure and developing water replenishment programmes where feasible. Responsible construction practices will be prioritised as will solutions to protect communities from noise pollution.
- **Hire local and invest in American workers:** Data centres will prioritise hiring residents from the communities and support local technical-education and workforce training programs.
- **Earn the buy-in of each community they call home:** Data centres will engage with local leaders and residents early, listen and be transparent. Infrastructure projects will respect community concerns and quality of life.

Figure 86

Summary of data centre-related policies and approaches for key markets (non-exhaustive)

Company	Water-use-specific targets or strategies	Target year / timeframe
AirTrunk	Minimise water withdrawal, pursue alternative water sources in water-stressed regions, and optimise water productivity and WUE.	Not publicly disclosed
Aligned	Uses waterless heat rejection and closed-loop cooling; states near-zero outside water usage and zero potable water in its closed-loop design.	Not publicly disclosed
Amazon	AWS to be water positive, returning more water to communities and the environment than it uses in data-centre operations. Also plans to expand recycled-water use to more than 120 US locations by 2030.	2030
Apple	Supplier Clean Water Program drives participating suppliers toward an average water reuse rate of 50%. Apple also says recent data centres have been designed with cooling systems that do not consume water.	2030 for supplier reuse target
Beale Infrastructure	Modern closed-loop data centres have near-zero day-to-day consumption and air-cooled systems skip water almost entirely. Beale prioritises recycled and non-potable water sources.	Not publicly disclosed
CloudHQ	Minimise water use, avoid irrigation, and only use water for heat rejection under specific and sustainable circumstances.	Not publicly disclosed
Crusoe	Designs AI data centres with closed-loop, non-evaporative cooling; aims for near-zero Cooling Water Usage Effectiveness.	Not publicly disclosed
CyrusOne	Target to make all facilities in high-water-stress regions net positive water facilities. Net positive is defined as restoring 20% more water than the facility withdraws from local watersheds.	Not publicly disclosed
DayOne	Integrates energy, water and waste efficiency into every new build; public targets include design PUE by 2030 and 100% renewable energy by 2030, but no specific WUE or water-positive target found.	Not publicly disclosed
Digital Realty	Tracks water use monthly and sets site-specific targets. WUE to be less than 0.4 by 2030.	2030
EdgeConneX	Committed to water neutrality across global operations.	2030
Equinix	Reports portfolio WUE performance, including 0.95 average annual WUE for the global portfolio.	Not publicly disclosed
Google	Replenish more freshwater than it consumes, on average, across data centres and offices.	2030
Meta	Become water positive; restore 200% of consumption in high-water-stress regions and 100% in medium-water-stress regions.	2030
Microsoft	Become water positive by replenishing more water than it withdraws. Also commits to a 40% improvement in datacentre water-use intensity.	2030
NTT Global Data Centers	Global standard design includes dry cooling systems by default for new projects.	Not publicly disclosed
Oracle	Halve water use in its most at-risk regions; use closed-loop, non-evaporative liquid cooling systems.	2035
QTS Data Centers	Freedom Design delivers WUE of 0 because its water-free cooling system does not consume water.	Annual 5% WUE reduction from 2022
SoftBank Group	FY2030 target to keep water consumption stable at current levels despite business expansion.	FY2025 and FY2030
Stack Infrastructure	Conserves potable water through efficiency, onsite reuse and connection to reclaimed water systems. Commits to efficient water practices and reclaimed-water/rainwater strategies.	Not publicly disclosed
Switch	Net positive water strategy intended to replenish and restore up to two times more water than it uses operationally.	"During the coming decade"
T5 Data Centers	Zero Water Loss Cooling; modular air-cooled chillers with WUE = 0.	Not publicly disclosed
TECfusions	Uses adaptive reuse, microgrid deployment and sustainability-focused design.	Not publicly disclosed
Tencent	Environmental policy states each office and data-centre campus should develop its own energy- and water-saving plan and actively promote water reuse and recycling.	20% WUE reduction by 2030 from 2025 levels
Vantage Datacenters	Air-cooled model with closed-loop chillers that need no ongoing water source; WUE described as virtually zero. For its Wisconsin Lighthouse campus, Vantage says it will pursue water positivity by restoring more water than the campus consumes.	Not publicly disclosed

Source: Company announcements and sustainability reports, Standard Chartered Bank

Figure 87

Examples of data centre disclosure areas

Disclosure area	Key statistics
Facility identity and capacity	Location, operating model, commissioning year, installed IT power, gross and live capacity, average and peak IT load, rack count and density.
Energy use and PUE	Total facility and IT electricity consumption, annual, monthly and peak demand, PUE.
IT efficiency and utilisation	Server, CPU, GPU, storage and network utilisation, compute per watt, tokens or operations per joule where applicable, idle-energy share, virtualisation and workload utilisation.
Water use and WUE	Water withdrawal, consumption and discharge, direct WUE, source-based WUE.
Water source and stewardship	Potable, surface, groundwater, recycled, reclaimed, harvested, brackish and other non-potable water, replenishment volume and location, percentage reused.
Cooling and thermal management	Cooling technology, heat-rejection system, liquid-cooled IT share, cooling efficiency ratio.
Heat recovery	ERF, recoverable heat, useful heat exported, delivery temperature.
Power sourcing	Grid electricity, onsite generation, PPAs, virtual PPAs, utility green tariffs, EACs, RECs or GOs, storage, nuclear and other low-carbon sources.
Carbon and air emissions	Location- and market-based Scope 1 and 2 emissions, relevant Scope 3 emissions, CUE, backup-generator fuel.
Circularity and lifecycle impacts	Embodied carbon, equipment lifetime, reused or refurbished hardware, material recovery, e-waste, construction waste, landfill diversion.
Resilience and grid interaction	UPS and battery capacity, backup generation, microgrid assets, demand-response participation, curtailed or shifted load, grid services, outage and uptime statistics.
Local environmental context	Basin water stress, drought exposure, grid carbon and water intensity, ambient climate, land take, biodiversity sensitivity, noise and community complaints.
Targets, methodology and assurance	Baselines, facility-level targets, progress, calculation standards, data coverage, estimates, exclusions, restatements, assurance level.

Source: Standard Chartered Bank

In summary

This chapter shows that the data centre water story is more nuanced than is often suggested. A data centre's total water footprint is shaped not only by direct cooling water use, but also by local climate, site selection, cooling architecture, power sourcing and the water intensity of the electricity system. This means that Water Usage Effectiveness alone can be an incomplete metric unless it is assessed alongside power efficiency, energy mix and local water conditions. The chapter also highlights that disclosed water intensity already varies significantly across operators, with some hyperscalers and smaller specialist providers demonstrating that much lower water intensity is achievable when scale, design, technology and operating discipline are aligned.

Looking ahead, several indicators point to future data centre water use being lower than headline capacity-growth forecasts imply. Liquid cooling, dry cooling, higher-temperature server operation, AI-enabled optimisation, recycled or non-potable water use, renewable power procurement and stronger operator targets have the potential to reduce direct or indirect water intensity. AI itself may accelerate this shift because high-density workloads are increasingly difficult to cool with conventional evaporative systems and may therefore push new facilities towards lower-water cooling architectures. However, low-water data centres are unlikely to become the norm automatically. Achieving that outcome will depend on careful location choices, transparent and comparable disclosure, continued technology adoption, stronger regulation where needed and active engagement with local communities to ensure that data centre growth does not come at the expense of local ecosystems. The most recent announcement by the AIIC in the US seems to indicate that the industry is moving towards adopting most of these solutions and strategies suggesting that tomorrow's data centres are likely to have a much lower environmental footprint than is the case today.



What does good look like?

This report has examined how data centres intersect with rising power and water demand and highlighted some of the solutions available that can help reduce a data centre's environmental footprint.

As AI drives a new wave of infrastructure investment, the challenge is not simply how quickly and successfully operators can adopt these solutions to minimise environmental impacts. It is also how financiers, policymakers and local communities distinguish between projects that are compatible with long-term sustainability objectives and those that may create additional environmental, resource or infrastructure pressures. The following section sets out what good looks like for three key stakeholder groups.

1. Data centre operators

Maximise the use of liquid and dry cooling technologies

This report shows that data centres can make meaningful progress in reducing both their power and water consumption profile by adopting the latest generation of cooling solutions. Adopting liquid cooling or direct free air cooling where climate conditions permit has the potential to reduce power intensity levels and reduce direct water consumption close to zero.

Shift to power sources with a low-carbon and water footprint

Developing alternative low-carbon sources of power allows data centre operators to reduce power supply-demand challenges that may arise due to the expansion of data centre capacity. These alternatives also help to reduce the indirect water footprint of data centres and reduce emission intensity of data centre operations. Options available include solar and wind, nuclear power including small modular reactors but also the use of geothermal, biofuels and hydrogen or fuel cell technology.

Make maximum use of circular solutions

Water and material-related circular economy solutions help reduce resource consumption, lower waste, cut lifecycle emissions and ease pressure on local ecosystems. Data centre operators can recycle municipal wastewater, cooling-tower water, rainwater or wastewater from industrial processes for cooling purposes. Embodied carbon or water use can be reduced through the reuse or refurbishing of IT equipment and recycling of building and construction materials.

Optimise data centre design

Optimising a data centre's design can reduce power and water use beyond the solutions listed so far. Examples of this include locating data centres in areas with favourable climate, water and power conditions and with access to low-carbon electricity. Other strategies include the adoption of AI to optimise facility performance, deploy the latest building and IT infrastructure solutions that allow servers to operate at higher temperatures, install higher voltage, direct current, power systems to reduce power loss and build battery energy storage capabilities (BESS).

Develop deep engagement strategies

To help create support from local communities and policymakers data centre operators need to implement deep strategies that directly address community concerns. These include developing water replenishment programmes and invest in local power and water infrastructure. Transparent disclosure policies help to alleviate concerns over the net impact of a data centre on local ecosystems. Using local resources such as labour, suppliers and services can also help ensure that local communities benefit economically from the construction and operation of a data centre.

Incorporate comprehensive disclosure policies

Current disclosure levels by data centre operators are often insufficient for external stakeholders including policymakers, financiers and local communities to assess the impact of an individual data centre. More transparent and standardised reporting would enable these stakeholders to make better-informed decisions.

A comprehensive disclosure policy is one where consistent, independently assured annual data is published for individual facilities. This data should cover resource inputs, operational efficiency and environmental outcomes. Some of the key data items that should be included are listed in Figure 85. The results should use standardised boundaries and methodologies, show absolute values as well as intensity ratios, state exclusions and estimation methods, provide monthly or seasonal ranges alongside annual averages where material, and be contextualised by facility location, climate, grid carbon and water intensity, and local water stress.

An integrated systems-level approach is needed

A systems-level approach to data centre management that integrates all these actions is essential because there is no universally "best" cooling or infrastructure solution. Lowering direct water use can increase electricity demand where dry or mechanical cooling replaces evaporative systems, while improving energy efficiency can increase water consumption if more water-intensive cooling is used. The optimal combination should assess energy, carbon and water together across IT efficiency, cooling and heat rejection, power sourcing, heat reuse, circularity, storage and grid interaction. It should also reflect the local climate, carbon and

water intensity of the grid, and prevailing water stress. Assessing a data centre's absolute impact alongside a balanced set of complementary metrics would help operators, policymakers and financiers avoid relying on or optimising a single measure, such as PUE or WUE and in doing so shift environmental pressures elsewhere. Instead, it would support facility designs and operating models that minimise total environmental impact in each location.

2. Financiers

Direct capital towards the highest-performing projects

Not all data centres have the same environmental footprint. Future capital allocation should increasingly distinguish between facilities based on their energy efficiency, water performance, carbon intensity, resilience and approach to community engagement.

Rather than relying on a single metric such as PUE, financiers should consider a broader range of indicators covering energy, water, emissions, circularity and local environmental conditions.

Reward transparency and disclosure

Limited facility-level disclosure makes it difficult to compare projects and assess risk. Greater transparency on operational performance, water consumption, power sourcing, emissions and local environmental impacts would improve decision-making, support more efficient capital allocation and help financiers identify projects that demonstrate leading practice.

Assess infrastructure and resource constraints alongside financial returns

Project viability increasingly depends on access to power, water and supporting infrastructure. Assessments should therefore consider factors such as local water stress, grid constraints, availability of low-carbon electricity, permitting risk and the potential for community opposition alongside traditional financial metrics.

Support solutions that reduce system-wide impacts

The transition to more sustainable AI infrastructure will require significant investment in areas including renewable energy, storage, grid infrastructure, advanced cooling technologies, heat reuse systems and water recycling solutions.

Financiers can play a role in accelerating deployment of these solutions by directing capital towards projects that strengthen the resilience and sustainability of the wider system.

3. Policymakers

Move beyond single sustainability metrics

No single indicator can fully capture a data centre's environmental performance. Policy frameworks should encourage assessment across energy, water, emissions, resilience, grid impacts and local environmental conditions to avoid creating incentives that improve one metric while worsening another.

Encourage greater transparency

More standardised, facility-level disclosure would improve market transparency and enable policymakers, investors and communities to make more informed decisions. The development of common reporting frameworks would also improve comparability across jurisdictions and operators.

Consider local conditions when defining policy priorities

The most appropriate solutions differ significantly between locations. Policy frameworks should reflect local factors including climate conditions, water availability, power-system characteristics, existing infrastructure and patterns of resource stress.

Prioritise long-term system resilience

As AI infrastructure expands, planning decisions will increasingly influence the resilience of electricity, water and digital systems. Policies that encourage efficient siting, infrastructure investment, grid flexibility and resource-efficient technologies can help support economic growth while reducing environmental pressures.

Summary

The growth of AI is turning data centres into one of the most important infrastructure themes of the coming decade. While this could create legitimate environmental pressures, the tools, technologies and operating practices available to address them are more advanced and scalable than recent headline concerns suggest.

Data centre capacity, electricity demand and local grid pressure are likely to rise materially as AI workloads become larger, denser and more power intensive. However, the sector has multiple levers available to it with the potential to reduce its power footprint substantially. These include more efficient IT hardware, better hardware utilisation, AI-enabled optimisation, advanced cooling, waste-heat reuse, higher-voltage power distribution and the large-scale procurement of low-carbon energy. The key question is not whether data centres will grow, but whether that growth will be supported by the right combination of efficient design, clean power, grid investment and transparent reporting.

Water use by data centres follows the same logic. Direct water use for cooling is important, but the broader water footprint of data centres is shaped by climate, location, power sourcing and the water intensity of the electricity system. New cooling technologies such as liquid cooling, dry cooling and free-air cooling have the potential to materially reduce water intensity and, in some cases, bring direct water consumption close to zero.

However, technical efficiency alone is unlikely to be sufficient. Best-in-class operators will also need to choose locations carefully, maximise low-carbon and low-water power supply, embed circularity, develop water recycling and replenishment programmes, disclose performance transparently and engage actively with local communities. In that sense, the future licence to grow will increasingly belong to those operators that can show they are not only building more capacity, but that do so with lower power intensity, lower water intensity and a smaller impact on local ecosystems.

Ultimately, the environmental footprint of AI infrastructure is not fixed. It is shaped by the choices made by operators, financiers and policymakers around technology, infrastructure and resource use. A broader assessment of energy, water, emissions, resilience and local conditions can help direct capital, policy and innovation towards the solutions and projects most capable of supporting responsible and sustainable AI-driven growth.

Appendix

Acknowledgement

This report benefited from discussions and consultation with Microsoft, Trane and NVIDIA, who provided background information and data. The conclusions and views expressed in this report are those of the author and do not necessarily represent those of Microsoft, Trane or NVIDIA.

About the author

Eugène Klerk is a Managing Director at Standard Chartered and as the Head of Sustainability Insights is responsible for generating and coordinating the firm's sustainability-related content. Eugène has published extensively on Nature, Blue Economy, Carbon, Circular Economy and Energy Transition-related challenges and solutions.

Prior to joining Standard Chartered Eugène worked for Credit Suisse, where for more than 10 years he was responsible for sustainable thematic and ESG research and managed the Global ESG research team. In addition, Eugène was the head of the Sustainability pillar of the Credit Suisse Research Institute.

Eugène started his career in the mid-90s as an emerging markets fixed income and equity analyst. During this time, he was responsible for the EMEA Research team at Credit Suisse First Boston and achieved multiple top 3 rankings in major surveys. Eugène holds a master's degree in applied mathematics.

Eugène Klerk
Managing Director
Head of Sustainability



References

ABI Research (a). (2026). Neocloud Data Centers. Retrieved from https://www.abiresearch.com/news-resources/chart-data/number-of-neocloud-data-centers.

ABI Research (b). (2026). Forecasting data center capacity increases in the AI era. Retrieved from https://www.abiresearch.com/blog/data-center-capacity-growth-forecast.

ABI Research. (2026). Forecasting data center capacity increases in the AI era. Retrieved from https://www.abiresearch.com/blog/data-center-capacity-growth-forecast.

AI Infrastructure Coalition. (2026). AIIC announced community principles, setting a gold standard for AI infrastructure development.

Al Kez, D., & et al. (2025). AI-driven cooling technologies for high-performance data centres: state-of-the-art review and future directions. Sustainable Energy Technologies and Assessments.

Ali, F. (2026). AI-driven cooling optimisation for data centers: Architectures, algorithms, real-world performance and a roadmap for sustainable computing infrastructure.

Alissa, H., & et al. (2025). Using life cycle assessment to drive innovation for sustainable cool clouds. Nature.

Ampcom. (2026). Data center cooling: Types, benefits and best practices.

Australian Government Department of the Prime Minister and Cabinet. (2026). Getting it right: Building AI infrastructure that works for Australia. Retrieved from https://www.pmc.gov.au/resources/getting-it-right-building-ai-infrastructure-works-australia

Bloom Energy. (2025). Data center survey (Nov. 2025).

Bloom Energy. (2026). 2026 Data Center Power Report.

Bloomberg. (2026). Microsoft reconsiders clean-energy pledge made harder by AI.

BloombergNEF. (2025). Global grid investment could top USD470 billion for the first time in 2025.

BloombergNEF. (2025). New Energy Outlook 2025: Grids.

Bunger, R., & et al. (2020). Capital cost analysis of immersive liquid-cooled vs air-cooled large data centers.

CBRE. (2025). UK Real Estate Market Outlook Mid-Year Review 2025.

Center for sustainable systems (a). (2026). Photovoltaic Energy Factsheet.

Center for sustainable systems. (2026). Wind energy factsheet.

CNDCP. (2026). Comments on the draft EU Data Centre Sustainability Rating Scheme.

Cushman & Wakefield. (2026). Asia Pacific data centre construction cost guide.

Data Center Dynamics. (2026). Survey: Cooling, as the data center industry evolves so do its cooling solutions.

Datacenters.com. (2026). AI workloads are increasing data center rack density by 5x. Retrieved from www.datacenters.com.

de Vries-Gao, A. (2025). The carbon and water footprints of data centers and what this could mean for artificial intelligence. Patterns.

Deloitte. (2024). Powering Artificial Intelligence.

Deloitte. (2025). AI for energy systems.

Deloitte. (2025). Can US infrastructure keep up with the AI economy?

Dos Santos Junior, A. (2025). AI-driven predictive maintenance and energy optimisation in intelligent manufacturing. Iconic Research and Engineering Journals.

EDCA. (2026). State of European Data Centres.

EIA. (2024). Carbon Dioxide Emissions Coefficients.

EIA. (2025). Electric Power Annual 2025 - Table 8.2 Average tested heat rates by prime mover and energy source.

EIA. (2026). US Energy-related Carbon Dioxide Emissions 2025.

Ember. (2026). Yearly electricity data. Retrieved from ember-energy.org/data.

Energy Solutions Intelligence. (2026). Data center cooling 2026: Liquid immersion vs air: TCO & performance.

Energy Star. (accessed 2026). Virtualise servers. Retrieved from www.energystar.gov.

EU. (2024). Commission delegated regulation (EU) 2024/1364, on the first phase of the establishment of a common Union rating scheme for data centres.

EU. (2026). Minimum performance standards for EU data centres. Retrieved from https://energy.ec.europa.eu/resources/preparatory-studies/minimum-performance-standards-eu-data-centres_en

European Commission. (2026). Energy performance of data centres. Retrieved from https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-directive/energy-performance-data-centres_en

European Commission. (2026). Supplementing Directive (EU) 2023/1791.

Financial Times. (2026). 'Bold but sane': what we learnt from Anthropic's 240,000-word IPO filing. Retrieved from https://www.ft.com/content/a6c43402-719e-4295-ae5f-541edb0aeb29?syn-25a6b1a6=1

Forbes. (2026). Retrieved from https://www.forbes.com/sites/giacomotognini/2026/10/16/why-debt-laden-neoclouds-could-become-a-trillion-dollar-wework-20/

Gartner. (2025). Retrieved from https://www.gartner.com/en/newsroom/press-releases/2025-11-17-gartner-says-electricity-demand-for-data-centers-to-grow-16-percent-in-2025-and-double-by-2030.

GEM. (2026). US gas power proposals tied to data centres nearly double in six months.

Global Electricity. (2026). Electricity Consumption by Type: Global Sector Breakdown and Trends.

Google. (2016). DeepMind AI reduced energy used for cooling Google data centers by 40%.

Google. (2026). Environmental report 2026.

GRC cooling. (2026). GRC TCO calculator.

GWi/Xylem. (2026). Watering the new economy - Managing the impacts of the AI revolution.

Herrera, M., & et al. (2025). Sustainable AI infrastructure: A scenario-based forecast of water footprint under uncertainty. Journal of cleaner production.

Ho, A., & et al. (2024). Algorithmic progress in language models.

IEA. (2025). Energy and AI.

IEA. (2026). Global Electric Vehicle outlook 2026.

IEA. (2026). Key questions on energy and AI.

IEA. (2026). World Energy Outlook 2025.

IMF. (2025). Power hungry: how AI will drive energy demand.

Introl. (2026). Liquid cooling vs air cooling for AI data centers: 2025 analysis.

Ipsos. (2026). The Ipsos AI monitor 2026.

Jegham, N., & et al. (2025). How hungry is AI? Benchmarking energy, water and carbon footprint of LLM inference.

JLL. (2026). 2026 Global Data Center Outlook.

Lei, N., & et al. (2025). The water use of data center workloads: A review and assessment of key determinants. Resources, conservation and recycling.

Lei, N., & Masanet, E. (2022). Climate- and technology-specific PUE and WUE estimations for I.S. data centers using a hybrid statistical and thermodynamics-based approach. Resources, conservation and recycling.

Li, P., & et al. (2025). Making AI less thirsty. Communications of the ACM.

Masanet, E., & et al. (2024). To better understand AI's growing energy use, analysts need a data revolution. Joule.

Maslej, N., & et al. (2025). The AI Index 2025 Annual Report.

Mazza, M., & et al. (2025). The evolution of neoclouds and their next moves. Retrieved from https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-evolution-of-neoclouds-and-their-next-moves.

McKinsey (a). (2025). The cost of compute a 7 trillion dollar race to scale data centers.

McKinsey (b). (2025). Scaling bigger, faster, cheaper data centers with smarter designs.

McKinsey (c). (2025). Beyond compute: infrastructure that power and cools AI data centres.

McKinsey. (2023). Investing in the rising data center economy.

Moduledge. (2026). Solar powered data centres: Economics, design and why the math now works.

MOSL. (2026). Water efficient data centres.

Mytton, D., & Ashtine, M. (2022). Sources for data centre energy estimates: A comprehensive review. Joule.

OECD.AI. (2026). Retrieved from https://oecd.ai/en/data?selectedArea=investments-in-ai-and-data&selectedVisualization=worldwide-vc-investments-in-ai.

Ofgem. (2025). Ofgem sets out major reform package in next step to accelerate grid connections.

PWC. (2026). Could net-zero AI become a reality. Value in motion.

Ranghino, F. (2024). Data center computing efficiency. Retrieved from www.ambientasgr.com.

S&P Global. (2025). Grid congestion remains key issue as data center load growth stresses system. Retrieved from spglobal.com.

Sarioguz, O. (2025). AI for sustainable logistics: reducing carbon emissions and fuel consumption through route optimisation. International Journal of Science and Research Archive.

Schneider Electric. (2023). Microgrids and batteries: A partnership for a more resilient and sustainable future.

Schneider Electric. (2024). Industrial AI: Optimising energy efficiency with predictive AI.

Schneider Electric. (2024). Understanding BESS: Battery energy storage systems for data centres.

Shehabi, A., & et al. (2024). 2024 United States Data Center Energy Usage Report.

Smith, S., Hubbard, A., & et al. (2026). United States Data Center Energy Usage Report: 2025 Update. Berkeley Lab.

Supplychain briefing. (2025). AI route optimisation: case studies in energy savings.

Synergy Research Group. (2025). The world's total data center capacity is shifting rapidly to hyperscale operators. Retrieved from https://www.srgresearch.com/articles/the-worlds-total-data-center-capacity-is-shifting-rapidly-to-hyperscale-operators.

The Green Grid. (2011). Water Use Effectiveness: A green grid data centre sustainability metric.

Transition Pathway Initiative. (2026). https://www.transitionpathwayinitiative.org/.

Turner and Townsend (b). (2025). Data centre construction cost survey.

Turner and Townsend. (2025). Data centre cost index. Retrieved from https://reports.turnerandtownsend.com/data-centre-construction-cost-index-2025/data-centre-cost-trends.

Uptime Institute. (2023). Server energy efficiency: Five key insights.

Vertiv. (2024). Strengthening mission-critical microgrids with a battery energy storage system.

Vertiv. (2025). Charting the financial roadmap and economics of BESS adoption for enterprises and data centres.

Water Europe. (2024). Data management: A factsheet on the sector's water-intensive practices.

Wood Mackenzie. (2026). Orbital data centres cost three times more than terrestrial alternatives.

World Nuclear Association. (2026). Carbon Dioxide Emissions From Electricity.

World Water-Tech. (2025). Plugging the leak: Innovative solutions for reducing water loss and its economic impact.

Disclaimer

This document is for information and discussion purposes only and does not constitute an invitation, recommendation or offer to subscribe for or purchase any of the products or services mentioned, or to enter into any transaction. The information contained herein is not intended to be used as a general guide to investing and does not constitute investment advice or a source of any specific investment recommendations, as it has not been prepared with regard to the specific investment objectives, financial situation or particular needs of any particular person.

The reader should be aware that this document and the information contained within it have been prepared on the following basis:

i. all material, positions and statements set out in this document are subject to change without notice;

ii. the material included in this document does not constitute investment, accounting, legal, regulatory, sustainability or tax advice, nor an invitation or recommendation to enter into any transaction;

iii. this document may have been prepared using models, methodologies and data which are subject to certain limitations. These limitations include the limited availability of reliable data, data gaps, and the nascent nature of the methodologies and technologies underpinning the data; the limited standardisation of data (given, amongst other things, limited international coordination on data and methodology standards); and future uncertainty (due, amongst other things, to changing projections relating to technological development and global and regional laws, regulations and policies, and the current inability to make use of strong historical data);

iv. some of the models, external data and methodologies used in this document are or could be subject to adjustment which is beyond our control;

v. any opinions and estimates should be regarded as indicative, preliminary and for illustrative purposes only;

vi. some of the information appearing in this document may have been obtained from public and other sources and has not been independently verified and no representation or warranty is made as to its quality, completeness, accuracy, fitness for a particular purpose or non-infringement of such information;

vii. for the purposes of this document, a number of key judgements and assumptions have been made. It is possible that the assumptions drawn, and the judgement exercised may subsequently turn out to be inaccurate. The judgements and data presented in this document are not a substitute for judgements and analysis made independently by the reader.

viii. any opinions or views of third parties expressed in this document are those of the third parties identified. Any incorporation of or reference to opinions and views of third parties in this document, is not, in any way, endorsing or supporting such opinions or views;

ix. the data contained in this document reflects available information and estimates at the relevant time;

x. where this document uses any methodology developed by a third party, the application of the methodology (or consequences of its application) shall not be interpreted as conflicting with any legal or contractual obligations and such legal or contractual obligations shall take precedence over the application of the methodology;

xi. where this document uses any underlying data provided or sourced by a third party, the use of the data shall not be interpreted as conflicting with any legal or contractual obligations and such legal or contractual obligations shall take precedence over the use of the data;

xii. further development of reporting, standards or other principles could impact this document or any metrics, data and targets included in this document (it being noted that Environmental, Social and Governance reporting and standards are subject to rapid change and development);

xiii. no representation or warranty is made as to whether sustainability labels or statements are appropriate or meet specific classification requirements or standards. Actual outcomes and results may differ materially from those expressed in, or implied by, any forward-looking statements relating to, without limitation, sustainability targets, strategy or objectives; and

xiv. while all reasonable care has been taken in preparing the information included in this document, no representation or warranty is made as to its quality, accuracy or completeness, and no responsibility or liability is accepted for the contents of this information, including any errors of fact, omission or opinion expressed.

You are advised to exercise your own independent judgement (with the advice of your professional advisers as necessary) with respect to the risks and consequences of any matter contained in this document.

The Standard Chartered Group, its affiliates, directors, officers, employees or agents expressly disclaim any liability and responsibility for any decisions or actions which you may take and for any damage or losses you may suffer from your use of or reliance on the information contained in this document.

Copyright in all materials, text, articles and information contained in this document (other than third party materials, text, articles and information) is the property of, and may only be reproduced with permission of an authorised signatory of, the Group.

© Standard Chartered Bank 2026. All rights reserved. Copyright in materials, text, articles and information created by third parties and the rights under copyright of such parties, are hereby acknowledged. Copyright in all other materials not belonging to third parties and copyright in these materials as a compilation vests and shall remain at all times copyright of Standard Chartered Bank and should not be reproduced, distributed or otherwise used without the prior written consent of Standard Chartered Bank.

We welcome feedback on our report.

Please contact us at
sustainability.feedback@sc.com