Expert Commentary:

Electricity demand growth for data centres and AI and implications for natural gas-fired power generation

Prepared by:

Dr. Alexander Ermakov, Energy Econometrician

October 2024

GECF

Contents

1. Introduction

The digital transformation of the global economy has sparked an unprecedented surge in data generation and processing, with cloud computing, big data analytics, and online services becoming indispensable to modern business operations. At the core of this transformation is the rapid advancement of artificial intelligence (AI), particularly generative AI, whose data-intensive machine learning models require extraordinary computational power. This AI-driven growth, combined with the increasing use of mobile data networks, video streaming, Internet of Things (IoT) devices, and energy-intensive applications like cryptocurrency mining, is placing immense pressure on data centres to expand their capacity. By 2030, the global data centre market is projected to grow at a compound annual growth rate exceeding 10%, with hyperscale facilities expected to dominate, pushing total data centre floor space to over 300 million square feet. This explosive expansion highlights the urgent need for scalable infrastructure capable of meeting the skyrocketing demand for data processing and storage.

As data centres expand to accommodate growing demand, one of the most critical challenges they face is the sharp rise in electricity consumption. AI, in particular, is an important driver of this growth. AI workloads, especially those involving deep learning and generative models, require massive computational resources, relying heavily on energy-intensive hardware like graphics processing units (GPUs) and tensor processing units (TPUs), which consume far more electricity than traditional processors. Training large AI models can take days or weeks of continuous processing, putting immense strain on the power infrastructure of data centres. Moreover, real-time AI applications – ranging from autonomous systems to large-scale predictive analytics – require a continuous supply of energy for inference and decision-making processes. As AI continues to dominate data centre workloads, the demand for stable, scalable, and reliable energy solutions becomes increasingly urgent.

This surge in electricity demand also presents broader challenges, particularly as it intersects with global energy transition strategies. Electrification is emerging as a cornerstone of mainstream energy policies, with electrons replacing molecules in sectors like transport, where electric vehicles (EVs) are gaining traction, and residential heating, where heat pumps are becoming more common. As these sectors increasingly electrify, competition for electricity resources is intensifying. AI-driven data centres, with their growing power needs, are competing with these newly electrified sectors for a finite supply of energy. Addressing the electricity needs of both AI-driven infrastructure and electrified industries will require significant investments in grid modernisation, energy infrastructure, and renewable energy integration to ensure the energy supply remains affordable and reliable.

While our latest GECF Global Gas Outlook 2050, published in March 2024, reflects some upside power demand from data centres and AI, uncertainties remain regarding the exact scale of this growth. Consequently, more analysis is needed to assess the real potential. In this expert commentary, we aim to discuss current trends and identify the major drivers that shape future electricity demand in this sector. Additionally, we have developed three scenarios to project electricity demand growth from data centres and AI by 2030, evaluating their implications for future power generation. Special attention is given to the role of natural gas, which is poised to emerge as a pivotal source in addressing surging electricity needs while balancing sustainability and grid stability.

2. Global electricity demand from data centres: background

The global push toward "smartification" – the widespread integration of advanced, interconnected technologies across homes, cities, industries, and infrastructure – is driven by the need for enhanced efficiency, automation, and data-driven decision-making. At the heart of this transformation are smart devices powered by the IoT and AI, which continuously collect, analyse, and act on vast amounts of data. By 2025, global data generation is expected to reach 175 zettabytes annually, a monumental rise from the 64 zettabytes generated in 2020, reflecting the exponential growth in data needs (1). This surge in data generation demands significant processing power, much of which will be handled by hyperscale data centres. The continuous operation of smart grids, autonomous vehicles, and intelligent appliances depends on these facilities, driving a strong link between data growth and electricity demand. As smart technologies become more deeply integrated into everyday life, data centres will need to scale rapidly to manage this rising volume of information, and the power required to support these systems will increase accordingly. This inextricable connection between data and electricity underscores the critical need for reliable, scalable energy infrastructure to support the global shift toward smart, interconnected ecosystems.

A key contributor to this rising energy demand is AI, which requires even more power due to the complexity and intensity of its operations. Training AI models – especially deep learning and generative models – requires processing big datasets over extended periods, with calculations running continuously across thousands of high-performance processors like GPUs and TPUs. These specialised chips consume significantly more electricity than traditional processors, amplifying the power needs of AI. Additionally, AI applications such as real-time analytics, autonomous systems, and natural language processing require constant power to handle inference tasks – the real-time decisions AI systems make. As AI scales, its computational and power demands grow exponentially, making it one of the most energy-intensive technologies in use today. With AI becoming integral to sectors such as healthcare, finance, and manufacturing, its expanding role will further strain electricity resources, highlighting the urgent need for scalable, reliable power solutions.

Between 2015 and 2019, the data centre sector managed to maintain remarkable stability in electricity consumption despite workloads more than doubling during that period. Electricity usage remained flat at around 200 TWh annually, primarily due to substantial energy efficiency improvements (Figure 1). This reflects the sector's ability to decouple data processing growth from energy consumption through continued efficiency gains. However, since 2020, data centres' global electricity demand has been on the rise, growing at an average annual rate of 20% and surpassing 400 TWh in 2023 (excluding contribution from cryptocurrency[1]). This now accounts for approximately 1.6% of global electricity consumption – up from 0.5% in 2000 – placing data centre electricity usage on par with countries like France, the 10th-biggest consumer of electricity globally. The COVID-19 pandemic accelerated this upward trend by increasing reliance on remote work, online education, e-commerce, and digital entertainment, all of which significantly boosted data centre workloads. Moreover, the continued demand for cloud services and AI technologies has continued to fuel this surge, while the diminishing potential for further energy efficiency gains signals a potential plateau in the sector's ability to balance energy use with workload expansion.

---

[1] Global cryptomining is estimated to consume around 110 TWh in 2022 and is excluded from analysis to maintain focus on data center operations and AI-driven computations.

Figure 1. Global electricity demand from data centres (TWh) and workloads (mn compute instances), 2015-2023

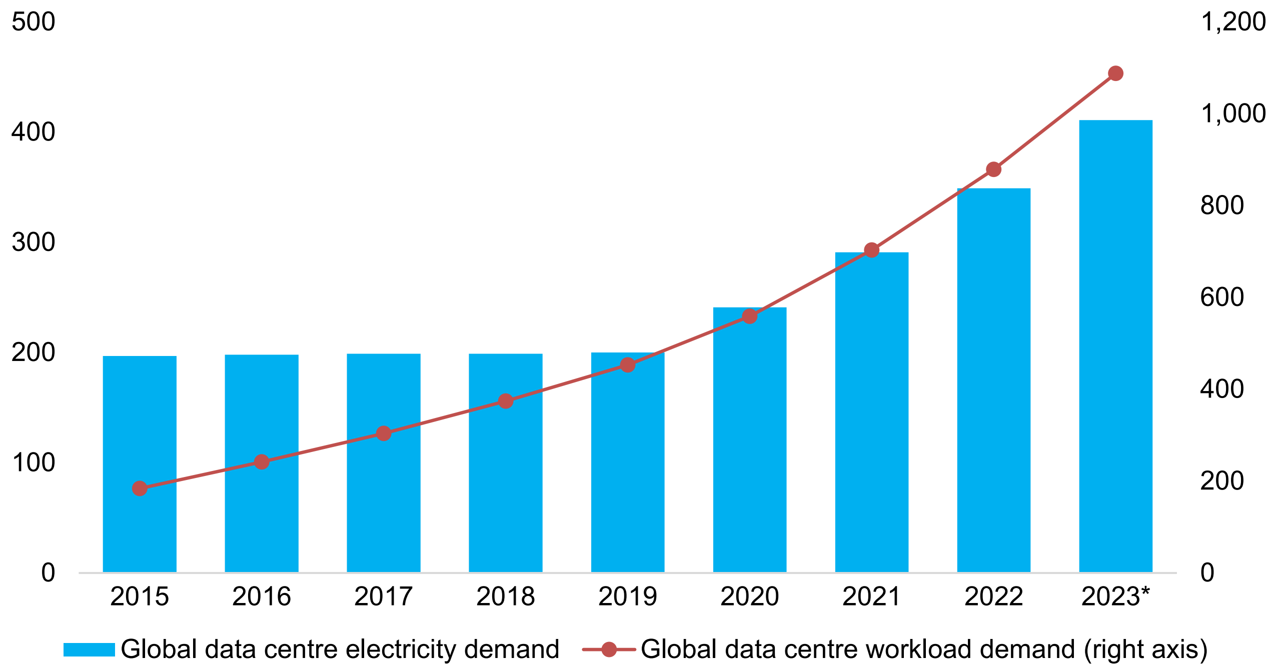

Source: GECF Secretariat based on data from Cisco, IEA, Goldman Sachs Global Investment Research. Data with the asterisk (*) is estimated.

Governments worldwide are grappling with the energy challenges posed by the rapid expansion of data centres. In 2019, Singapore temporarily halted the construction of new data centres, citing concerns about its ability to meet net-zero emissions targets by 2050. While the moratorium was lifted in 2023, with new permits granted – including one to ByteDance, the parent company of TikTok – this pause highlighted the difficulty of balancing infrastructure growth with sustainability goals. Similarly, in Ireland, data centres now account for 18% of the country's electricity consumption. To safeguard grid capacity, EirGrid, the national utility, has halted new data centre connections until 2028. Some companies, such as Microsoft, are exploring alternatives, such as using stand-alone gas-fired plants to power their data centres, as seen with a proposed project in Dublin.

The energy infrastructure in northern Virginia, known as "data centre alley," is also feeling the strain. In 2022, Dominion Energy, the regional utility provider, paused new power connections to data centres to ensure the grid could support the growing demand. With over three-square kilometres of data centres concentrated in Loudoun County – an area that is projected to see an 85% increase in data centre energy demand over the next 15 years – managing electricity supply in this high-density area presents significant logistical challenges (2). Although regulatory measures like these have not halted data centre growth, they have shifted it to less constrained markets like Maryland and Malaysia. However, as data centre activity grows in these secondary regions, governments may soon face similar challenges in managing energy resources and protecting their grids from being overwhelmed.

Despite this rapid expansion and the increasing global reliance on digital services, the internet's overall electricity consumption has remained remarkably efficient. Additionally, generative AI models have not yet been launched. According to the IEA, between 2015 and 2022, the number of internet users increased by 78%, global internet traffic soared by 6 folds, and data centre workloads surged by 340%. Yet, electricity consumption in data centres only rose by 20-70% during the same period. This efficiency was largely driven by innovations in cooling technologies, power management, and server optimization. The transition from traditional data centres to cloud and hyperscale facilities, which operate with significantly lower electricity consumption intensity, also played a key role. By 2020, over 90% of workload compute instances were processed by cloud and hyperscale data centres, which operate at nearly 80% better energy efficiency than smaller facilities (3).

One key factor driving these improvements is Koomey's Law, which states that the energy required to perform a given amount of computation halves every two-and-a-half year. This trend has allowed data centres to manage increasingly complex workloads with far less energy. Moreover, as data centres grow in size – particularly with the rise of hyperscale facilities – their energy efficiency has improved, with a greater share of power dedicated to computation rather than auxiliary functions like cooling and infrastructure. These hyperscale facilities have led innovations in cooling, AI-powered energy optimisation, and other technologies that reduce power consumption relative to workloads.

A key metric reflecting this energy efficiency improvement is Power Usage Effectiveness (PUE), which compares the total power used in a facility with the power consumed by IT equipment. As Figure 2 illustrates, global PUE values have fallen significantly since 2007, but progress has stagnated in recent years, with the global industry average hovering between 1.55 and 1.59 since 2020 (4). This stagnation is largely due to the growing number of data centres in regions like Asia, the Middle East, Africa, and Latin America, where facilities are typically smaller and located in warmer climates, requiring more energy for cooling. Even hyperscale data centres with advanced designs have struggled to make substantial gains. For instance, Google reported an average PUE of 1.1 in its latest environmental report, but improvements have been incremental over the past five years (5).

Figure 2. Global average annual PUE in data centres, 2007-2023

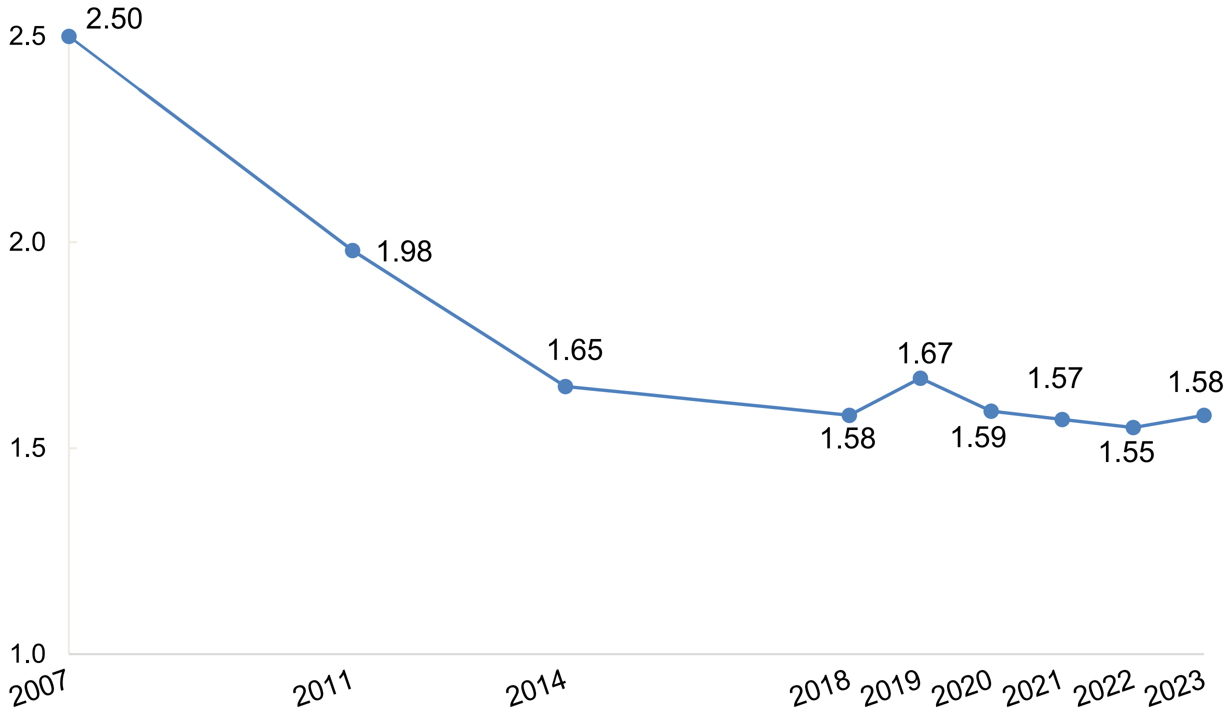

Source: Uptime Institute Intelligence

This highlights a critical challenge: the energy efficiency gains that have driven the stability of electricity consumption in data centres are starting to fade as many of the low-hanging fruits have already been achieved. Future improvements will depend on breakthroughs in both technological and algorithmic methods. Advanced cooling techniques, such as liquid cooling, and AI-powered energy optimisation, as well as smarter data routing, will be crucial to keep pace with the increasing workloads brought by AI and other energy-intensive technologies.

While these efficiency gains have allowed the sector to manage exponential growth without a proportional surge in electricity consumption, maintaining this trajectory will become more challenging. The rise of AI, machine learning, and real-time data processing will require continued innovation in computational efficiency and sustainable energy sourcing. As the world continues its digital transformation, the ability to balance increasing workloads with sustainable power usage will be critical for the long-term growth of the internet and data infrastructure.

3. Global electricity demand from data centres: prospects

Looking ahead, the world is poised to experience an unprecedented surge in data generation and usage. Projections indicate that global data volumes could reach a staggering 335 zettabytes by 2030 – up from 64 zettabytes in 2020 (6). This exponential growth is being driven by an increasingly digital global economy, where businesses, governments, and industries rely more heavily on data-driven insights, automation, and advanced analytics. As a result, the demand for expanded data centre capacity is set to soar, with these facilities becoming the backbone for processing, storing, and managing this tidal wave of information. This surge in data will not only require a vast increase in computational power but also fuel a significant rise in electricity demand as data centres expand to handle the growing workload.

A major contributor to the rising electricity demand is the rapid adoption of AI as a general-purpose technology. AI is revolutionizing industries such as healthcare, finance, manufacturing, and autonomous systems, leading to a sharp increase in the computational workloads handled by data centres. These AI-driven workloads are projected to significantly boost data centre electricity consumption by 2030, further intensifying the sector’s energy needs. As AI continues to evolve and become more integrated into daily operations across various industries, the pressure on data centres to expand will grow, establishing a clear link between AI adoption and the increasing demand for electrical power.

As indicated above, while advancements in energy efficiency, such as improved cooling technologies and server optimisation, have historically helped offset rising power demands, the potential for future efficiency gains is diminishing. Hence, as data centres grow in size and complexity to meet the surging demands of both AI and data generation, cooling systems, power distribution networks, and server infrastructure may struggle to keep pace. The result will likely be diminishing returns in energy efficiency, meaning that the more data and AI-driven workloads are generated, the more energy will be required, without proportional gains in efficiency.

The convergence of these trends – explosive data growth, the rise of AI, and limited future efficiency improvements – points to a future where electricity demand from data centres could rise significantly. However, uncertainties surrounding anticipated workload, efficiency improvements and technological advancements create challenges in accurate projecting future electricity consumption for data centres. Particularly, this is valid for AI-related demand, which remains highly unpredictable as the adoption of AI technologies, especially generative AI, is just emerging. The release of OpenAI’s ChatGPT in 2022 sparked significant public and corporate interest, raising questions about the potential impact of these tools on electricity needs.

Calculations hinge on multifaceted factors, encompassing details about the chips, cooling systems, data centre design, software, and workload. Estimates exist, but offer partial and approximate numbers, providing a general overview of AI's electricity consumption. One important factor that can be identified is the difference between initial model training phase and deploying it to users, a process known as inference phase.

Training an AI model, the process by which it learns patterns from huge datasets, requires using GPUs, which are power-hungry hardware. As an example, the GPUs that trained a large language model Generative Pre-trained Transformer 3 (GPT-3, the precursor to ChatGPT) consumed 1,287 MWh. The electricity consumption for training is set to increase as models get more advanced and have been steadily trending upward in size. Particularly, it is estimated that GPT-4 consumed 62,319 MWh for training, over 40 times higher than what its predecessor, GPT-3, required (7). The scale of the program's training is evident also from the number of parameters involved: GPT-4 features 1.7 trillion machine learning parameters, a substantial increase compared to GPT-3's 175 billion.

Training is only a part of the picture, given that this process incurs a one-time cost. The ongoing use of an AI model consumes electricity over time, and it stands to reason that most of the electricity used for AI will eventually stem from inference. For example, a study from Schneider Electric estimates that 80% of the AI workload in data centres in 2023 is from inference and 20% is for training (8).

Within just two months of its launch, ChatGPT attracted 100 million global users, prompting major tech companies like Microsoft, Alphabet, Meta, and Bing to introduce their own large language model (LLM) chatbots. In terms of electricity consumption, these LLMs represent a new challenge, with their overall impact largely depending on how widely the 5.3 billion internet users embrace these new AI-driven features. The extent of this adoption will significantly influence future data centre electricity demand.

For example, AI queries are estimated to consume up to 10 times more electricity than traditional Google searches, which typically use around 0.3 Wh per query. Semi Analysis projects that if LLMs were integrated into every Google search, daily electricity consumption could reach 80 GWh, totalling 29.2 TWh annually (9). Similarly, New Street Research estimates that approximately 400,000 servers would be required, consuming 62.4 GWh per day, or 22.8 TWh per year. Given that Google currently handles up to 9 billion searches daily, these projections suggest that each AI-powered search could use between 6.9 and 8.9 Wh of electricity. This aligns with data from Hugging Face’s BLOOM model, which averaged 3.96 Wh per request (10). Figure 3 provides various estimates.

A more pragmatic forecast of AI electricity demand has been proposed by Alex De Vries, based on the projected sales volume of AI servers by NVIDIA (which is accounted for around 95% of global market share for AI accelerators) and their rated power. Estimations indicate a substantial increase, from 100,000 servers in 2023 to 1.5 million units by 2027, assuming some industry parameters remain constant, such as the growth rate of the AI market and the availability of AI chips. The AI industry is expected to have grown exponentially to consume at least ten times its today demand, ranging 85.4-134 TWh by 2027.

Figure 3. Electricity consumption per query/search (Wh)

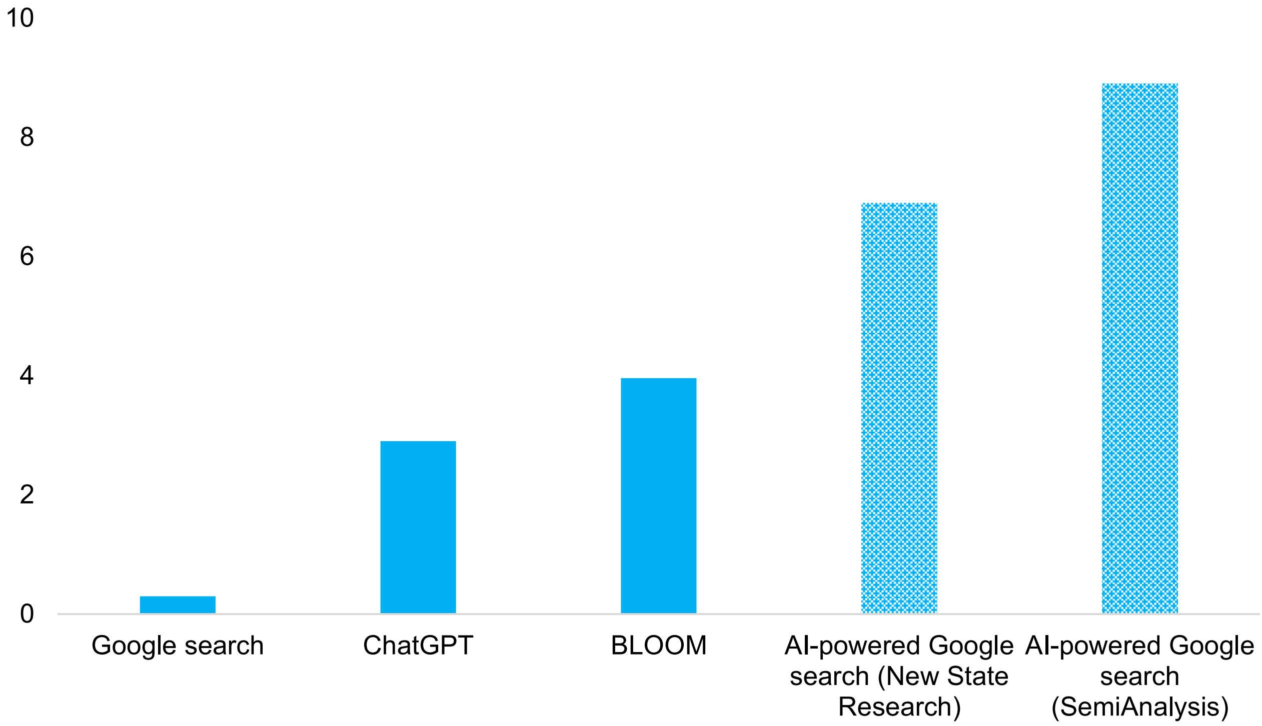

Source: De Vries, A. (2023). The Growing Energy Footprint of Artificial Intelligence.

It is important to note that efficiency improvements in chip hardware could have a big impact in electricity usage. For example, NVIDIA recently introduced a new line of GPUs with 25 times lower energy consumption than its previous model HGX H100. However, as computational power and efficiency rise, there is the potential for rebound effects, where the reduced cost of computation leads to increased demand. This phenomenon aligns with Jevons' Paradox, formulated in 1865, which suggests that improvements in efficiency can result in a net increase in resource use. This paradox has been observed in the history of technological change and automation and can be particularly relevant in the context of AI and data centres.

In addition to uncertainties surrounding workload, AI adoption, and advances in energy efficiency and data centre expansion, factors such as AI model complexity and the exponential growth of data further contribute to the unpredictability of energy demand in this sector. The evolving regulatory landscape surrounding emissions and sustainability adds another layer of complexity. Market fluctuations in demand for cloud services and data storage only amplify the challenges in forecasting long-term electricity consumption.

Given these significant uncertainties and based on publicly available information about existing data centres and industry growth estimates, the GECF has developed three scenarios to project future electricity consumption for data centres and AI by 2030. This analysis draws on historical trends in the industry, compute workloads, and internet traffic, while also accounting for the increasing computational intensity and widespread adoption of AI models. Each scenario reflects different trajectories for data centre growth, AI adoption and energy efficiency improvements.

1. Base Case Scenario

Assumptions: The base case represents the most likely scenario. AI adoption grows steadily, with moderate increases in data centre capacity to meet the expanding demand for cloud services, storage, and AI applications.

Energy Efficiency: Energy efficiency improvements, such as advanced cooling systems and more efficient hardware, are implemented at a moderate pace, helping to offset some of the increased demand. We assume that power efficiency gains, which slowed to around 2% from 2020 to 2023, will remain modest throughout this decade.

Outcome: The base case projects electricity demand to grow from an estimated 414 TWh in 2023 to 1,135 TWh by 2030, representing an annual growth rate of 15.5%. This reflects a balanced progression of AI workload expansion and data centre efficiency improvements. The forecast also suggests that data centres sector to account for 3.6% of global electricity consumption by 2030, up from the current 1.6%.

2. Low Case Scenario

Assumptions: In this scenario, AI adoption is slower than anticipated, and growth in data generation is less pronounced. The demand for cloud services and data centre expansion is more limited, reducing the overall need for additional data centre capacity.

Energy Efficiency: Significant advancements in energy-efficient hardware, cooling technologies lower electricity consumption per unit of data processed. Furthermore, the rapid integration of renewable energy drives investments in complementary technologies like energy storage or AI-driven load optimisation, indirectly enhancing energy efficiency. We assume a re-acceleration the efficiency assumption, compared to a base case, to about 5-7% during 2024-2030.

Outcome: The low case projects electricity demand to grow by an average 9.3% per annum, reaching 770 TWh by 2030, with sector's contribution rising to 2.5% of the total electricity demand. In addition to the early adoption of energy-efficient technologies and renewables, this scenario anticipates reduced data centre expansion due to grid constraints.

3. High Case Scenario

Assumptions: This scenario assumes rapid growth in AI adoption, exponential data generation, and the widespread use of computationally intensive AI models. There is also strong demand for cloud services and hyperscale data centres, leading to a significant increase in data centre capacity.

Energy Efficiency: In this scenario, advancements in energy-efficient technologies and cooling systems are slower than in Base Case, resulting in higher electricity consumption per unit of data processed. This scenario assumes a near-zero annual efficiency gains.

Outcome: The high case projects a steep rise in electricity demand, to about 1,560 TWh by 2030, corresponding to a 20.9% annual growth rate, as data centres expand rapidly to meet surging AI workloads and data processing needs, with almost negligible progress in efficiency improvements. As a result, data centres sector is set to account for almost 5% of global electricity consumption by 2030.

Figure 4. Forecast for global data centre and AI electricity consumption (TWh)

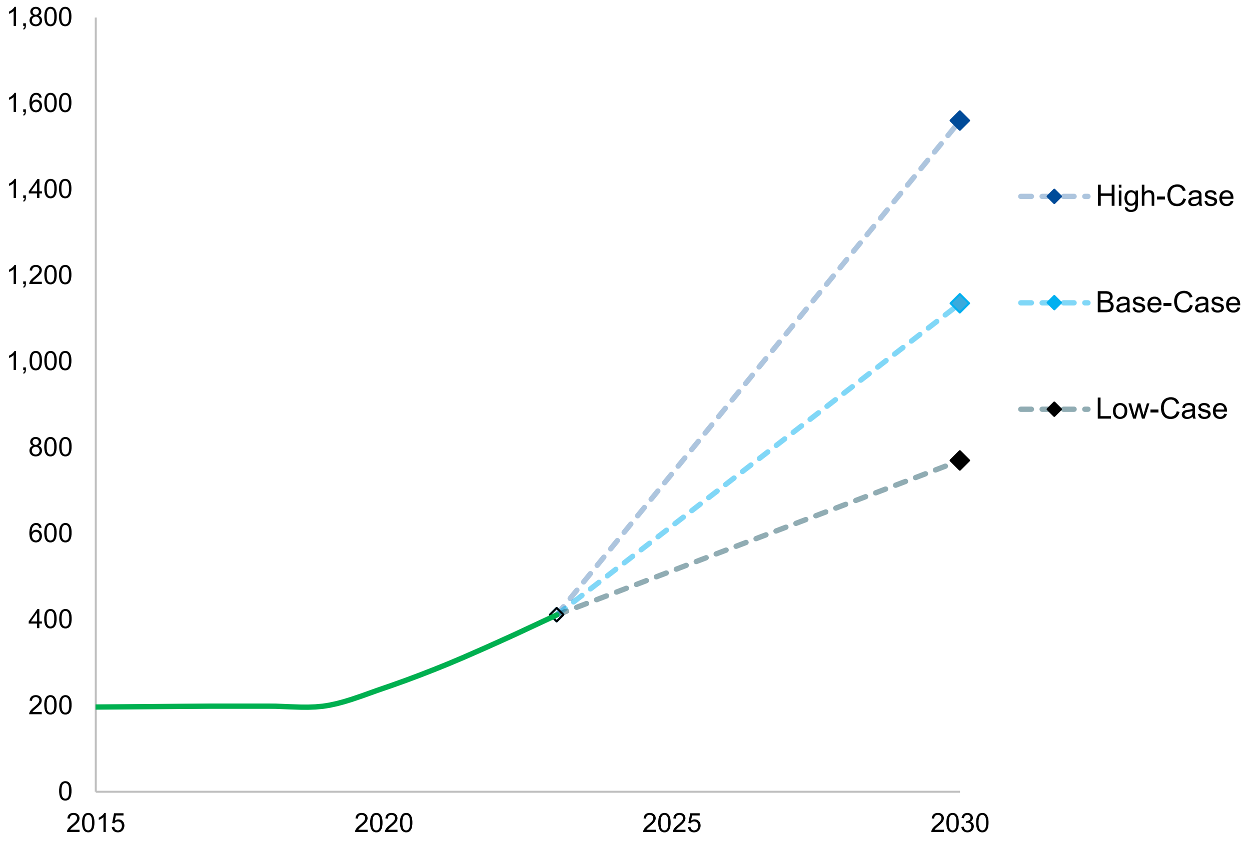

Source: GECF secretariat based on data from the GECF GGM

The regional distribution of electricity demand growth by 2030 is expected to be predominantly driven by the Asia Pacific region, North America, and Europe across all scenarios. Asia is expected to exhibit the fastest growth due to rapid digital transformation, rising data consumption, expansion of 5G networks and IoT, and government initiatives promoting AI adoption. China's significant investments in digital infrastructure, combined with India's growing tech services sector, will significantly contribute to rising electricity consumption.

North America, with its leadership in hyperscale data centres and cloud service providers, is forecast to be a major contributor as well. The United States continues to see strong adoption of cloud-based services as well as AI across sectors like finance, healthcare, and tech, and its well-developed grid infrastructure, coupled with access to abundant natural gas resources, supports this trend. Furthermore, competitive state tax incentives along with ongoing investment in AI research and development by large tech companies ensure sustained growth in electricity demand. Similarly, Europe is projected to see strong growth of electricity demand from data centres up to 2030, with significant number of additional data centres planned. Although the region's ambitious energy efficiency targets as well as prioritising the integration of renewable energy will balance electricity consumption and environmental goals.

4. Global electricity generation from data centres: implications for natural gas-fired generation

The demand for reliable and efficient electricity sources is increasingly critical for data centres' operation. The most attractive power generation options are those that can provide continuous, uninterrupted power, low operational costs, and the flexibility to scale in line with the growing energy needs of digital infrastructure.

In theory, renewables would be an ideal option due to their low environmental impact and alignment with corporate sustainability goals. Companies like Google and Microsoft are making significant investments in renewable energy to power their data centres, proactively pursuing energy efficiency technologies and signing power purchase agreements (PPAs) to fund renewable energy development. However, the intermittent nature of wind and solar power poses a significant challenge, as data centres require 24/7 electricity. While battery storage technologies, particularly lithium-ion batteries, have made notable advancements, they are still evolving and not yet widely deployable at the scale needed for grid-level reliability.

Nuclear power provides another attractive option, offering carbon-free baseload power. With a high-capacity factor and low operational costs, it is well-suited for facilities requiring constant electricity. However, high upfront capital costs, lengthy development timelines, and low public acceptance, owing to concerns over safety and waste management, limit nuclear power's ability to quickly respond to the rising electricity demand, particularly in the coming decade. Although there is growing interest in emerging Small Modular Reactor (SMR) technologies, their development remains in the early stages. Given the time required for regulatory approvals, pilot projects, and initial deployments, SMRs are unlikely to play a significant role before the mid-2030s.

While both renewables and nuclear are set to contribute to data centre power requirements, these options will not be sufficient to cover the growing needs as well as to maintain grid stability, especially during peak load periods. This highlights the ongoing need for natural gas and, to some extent, coal to fill the gap and ensure stable electricity supply.

Incorporating the abovementioned scenarios for data centres (Base, Low, and High Cases), global power generation is forecast to reach 36,007 TWh, 35,601 TWh and 36,483 TWh, respectively, by 2030 (Figure 5). This translates into an increase of 803 TWh, 397 TWh and 1,275 TWh above the projection for the Reference Case Scenario presented in the 8th Edition of the GECF Global Gas Outlook (GGO-RCS). Natural gas-fired generation is anticipated to account in the range of 41-45% (depending on scenario) of the additional electricity needs from data centres, making it a crucial energy source for meeting this surge. Additionally, gas-fired plants are favoured for their ability to provide reliable, on-demand power, ensuring operational continuity in data centres, particularly when renewable sources face intermittency.

Figure 5. Global power generation outlook, 2023-2030 (TWh)

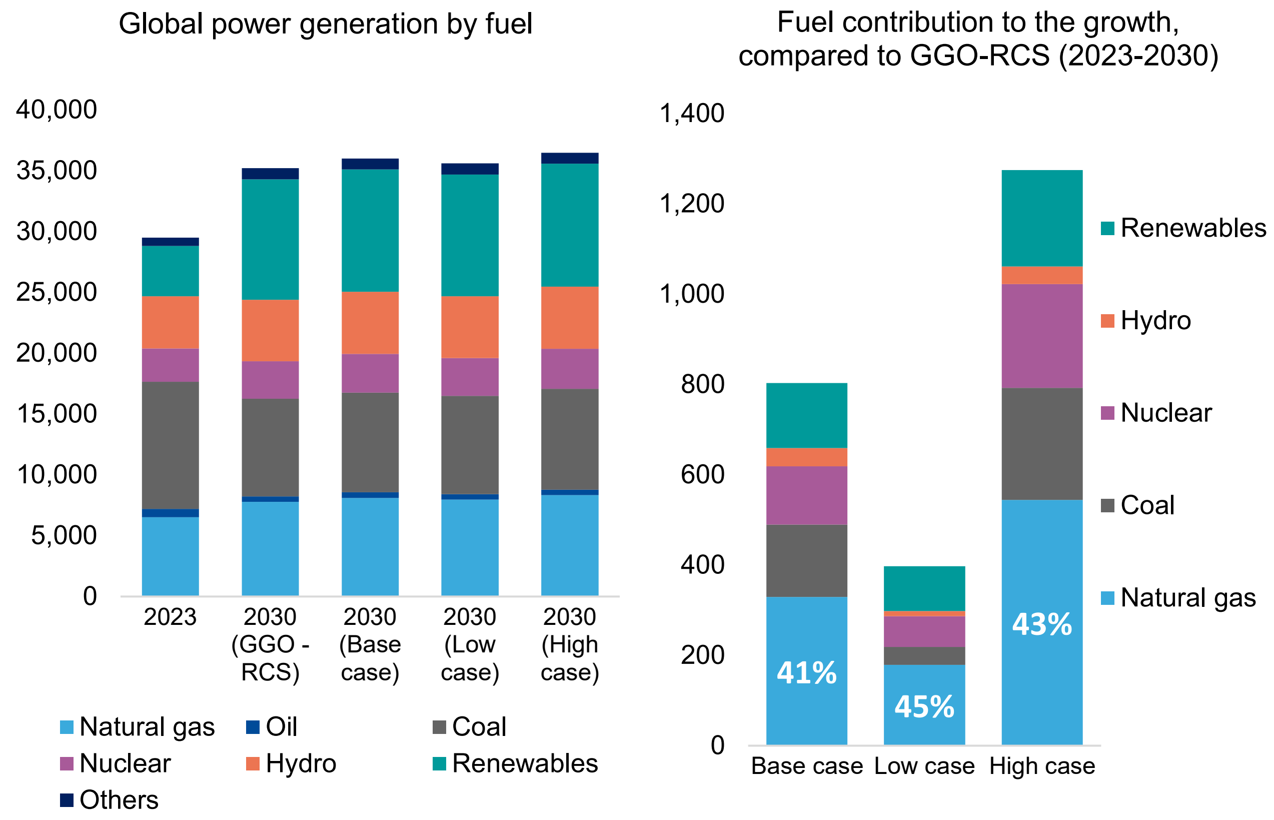

Source: GECF Secretariat based on data from the GECF GGM

For a more detailed view, we focus on the Base Case scenario, considered the most realistic and pragmatic trajectory. (Figure 6). The regional distribution of electricity generation growth varies by fuel type, with Asia Pacific, North America, and Europe leading the expansion. Together, these regions are forecast to account for 83% of the incremental global power generation growth.

In the Asia Pacific region, additional coal-fired power cannot be ruled out, as coal maintains cost advantages and remains a significant part of the power generation mix. Meanwhile, gas-fired generation is expected to contribute to 36% and will play a critical role in balancing supply, while helping to minimise CO_2 emissions.

In North America and Europe, natural gas and renewables are set to be central for electricity generation growth. The United States, a leader in hyperscale data centres, will rely heavily on gas-fired power plants for their flexibility and reliability. Companies have already announced plans to build more new gas-fired power capacity across the country to meet a surge in demand from power-hungry AI data centres. Europe, in contrast, is expected to adopt a more balanced mix, focusing on integrating renewable energy into data centre networks while still relying on natural gas for grid stability. Renewables are anticipated to make a substantial impact in both regions, driven by corporate sustainability commitments and government policies supporting clean energy investments.

Meanwhile, natural gas emerges as a pivotal resource in addressing the electricity needs of data centres and AI. Offering unmatched flexibility, scalability, and reliability, natural gas is uniquely suited to meet the high and often fluctuating power demand of modern data centres. Unlike intermittent renewable energy sources, natural gas-fired plants can provide stable, around-the-clock power, enabling data centres to operate continuously without disruption. Moreover, natural gas-fired plants are well-suited to respond quickly to demand spikes, which are typical of AI workloads that often require sudden bursts of computational power. In regions where full renewable energy deployment is not yet feasible, natural gas serves as a cleaner alternative to coal, offering a lower carbon footprint while ensuring a stable energy supply.

Figure 6. Base case: fuel contribution to the regional electricity generation growth from data centre and AI between 2023 and 2030 (TWh)

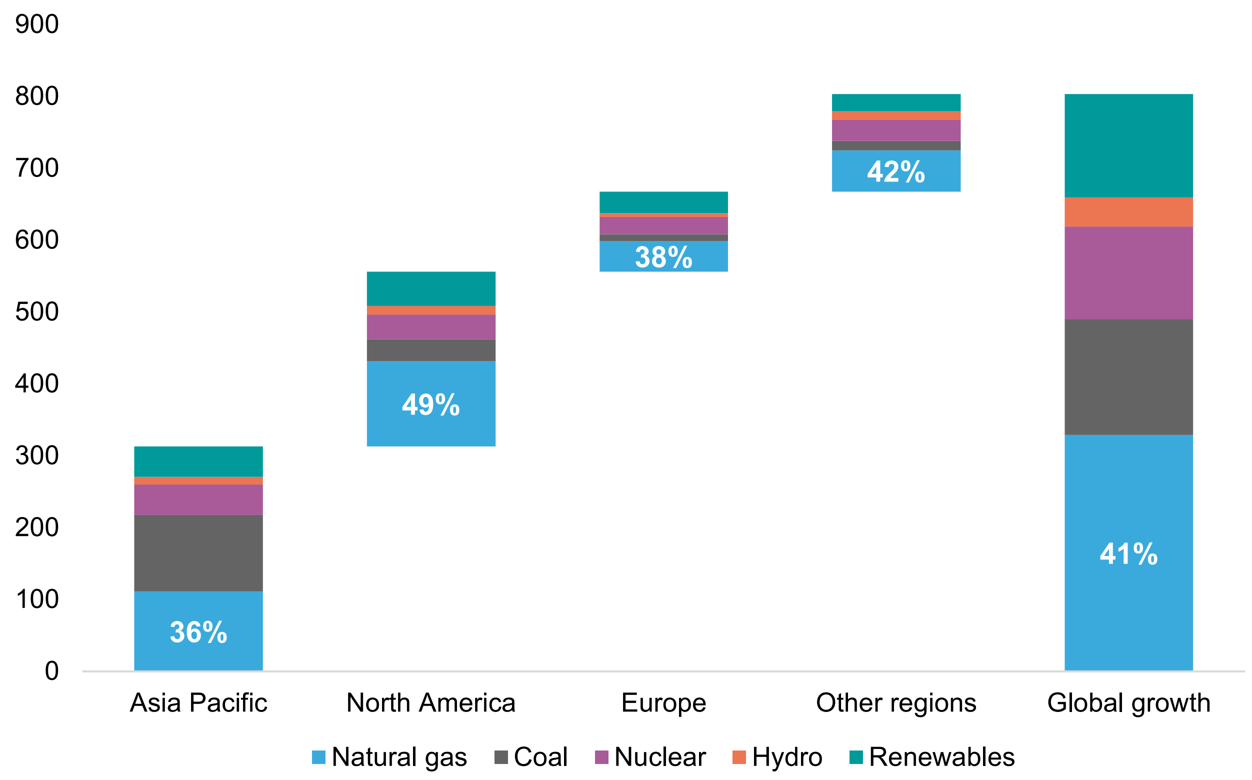

Source: GECF secretariat based on data from the GECF GGM

5. Conclusion

The expansion of data centres and AI technologies is set to drive a substantial rise in global electricity demand by 2030. While advancements in energy efficiency have historically helped manage this growth, the increasing scale and complexity of future workloads will require far more power, outpacing expected efficiency improvements. As projected in the scenarios, electricity demand from data centres and AI is forecast to grow significantly, from 414 TWh in 2023 to between 770 TWh and 1,560 TWh by 2030. The Base Case, regarded as the most realistic scenario, suggests a 174% increase to 1,135 TWh by 2030. This surge presents a significant challenge for power systems. The demand for reliable, continuous, and cost-efficient electricity sources is essential for ensuring the uninterrupted operation of data centres.

Moreover, addressing these rising energy needs will require power sources that can scale rapidly while also aligning with sustainability goals. Renewables and nuclear power are expected to play an important role, but each has limitations. Renewables struggle with intermittency, and nuclear requires high upfront capital costs and long development timelines, making it difficult for both to fully meet the sector’s electricity demand growth, specifically within the timeframe to 2030.

In this context, natural gas-fired generation, with its flexibility, scalability, and lower carbon footprint compared to coal, is positioned to remain crucial to support rapidly rising electricity needs. In the Base Case, natural gas is forecast to account for 41% of the additional global electricity generation growth, coming from the data centre sector. This is particularly relevant for regions such as North America and Europe, where natural gas will complement renewable energy integration and provide grid stability.

In conclusion, the growth of data centres and AI presents a dual challenge: how to meet the surging electricity demand while maintaining a commitment to sustainability and emissions reduction. This requires a multifaceted approach – one that integrates renewables, leverages the flexibility of natural gas, and supports ongoing innovations in energy efficiency. As the digital and energy landscapes continue to evolve, strategic investments in power infrastructure will be essential to ensure that the global economy's digital transformation is supported by a stable, scalable, and sustainable energy.

In this commentary, we have focused on the additional electricity demand arising from the growth of AI data centres and their expanding infrastructure. However, it is important to note that the broader economic impact of AI, including the increased energy consumption driven by productivity gains across various sectors, will also drive electricity consumption higher. This important impact will be addressed separately in an upcoming analysis. Understanding both these aspects is crucial for a comprehensive view of AI's influence on energy demand.

6. References

1. International Data Corportaion (IDC), Global DataSphere Forecast.
2. S&PGlobal (2024, April 1), Utilities face challenges, opportunities from AI-driven data center power demand growth.
Retrieved from https://www.spglobal.com/commodityinsights/en/market-insights/latest-news/electric-power/040124-utilities-face-challenges-opportunities-from-ai-driven-data-center-power-demand-growth-report
3. Goldman Sachs (2024, April 28). AI/data centres' global power surge and the sustainability impact.
Retrieved from https://www.goldmansachs.com/images/migrated/insights/pages/gs-research/gs-sustain-generational-growth-ai-data-centers-global-power-surge-and-the-sustainability-impact/sustain-data-center-redaction.pdf
4. Uptime Institute (2024, February 7). Large data centres are mostly efficient.
Retrieved from https://journal.uptimeinstitute.com/large-data-centers-are-mostly-more-efficient-analysis-confirms/
5. Google (2024, July). 2024 Environmental Report.
Retrieved from https://www.gstatic.com/gumdrop/sustainability/google-2024-environmental-report.pdf
6. International Data Corporation (IDC), Data Age 2025.
7. Extreme networks (2024, August 15). Artificial Intelligence, Real Consequences: Confronting AI's Growing Energy Appetite.
Retrieved from https://www.extremenetworks.com/resources/blogs/confronting-ai-growing-energy-appetite-part-1
8. Schneider Electric (2023, December 6). The AI Disruption: Challenges and Guidance for Data Centre Design.
Retrieved from https://www.se.com/ww/en/download/document/SPD_WP110_EN/
9. SemiAnalysis (2023, February 9). The Inference Cost Of Search Disruption – Large Language Model Cost Analysis.
Retrieved from https://www.semianalysis.com/p/the-inference-cost-of-search-disruption
10. De Vries, A. (2023, October 18). The Growing Energy Footprint of Artificial Intelligence. Joule, 7 (10), 2191-2194. Retrieved from https://www.sciencedirect.com/science/article/abs/pii/S2542435123003653?dgcid=author