BANCA D'ITALIA
EUROSISTEMA

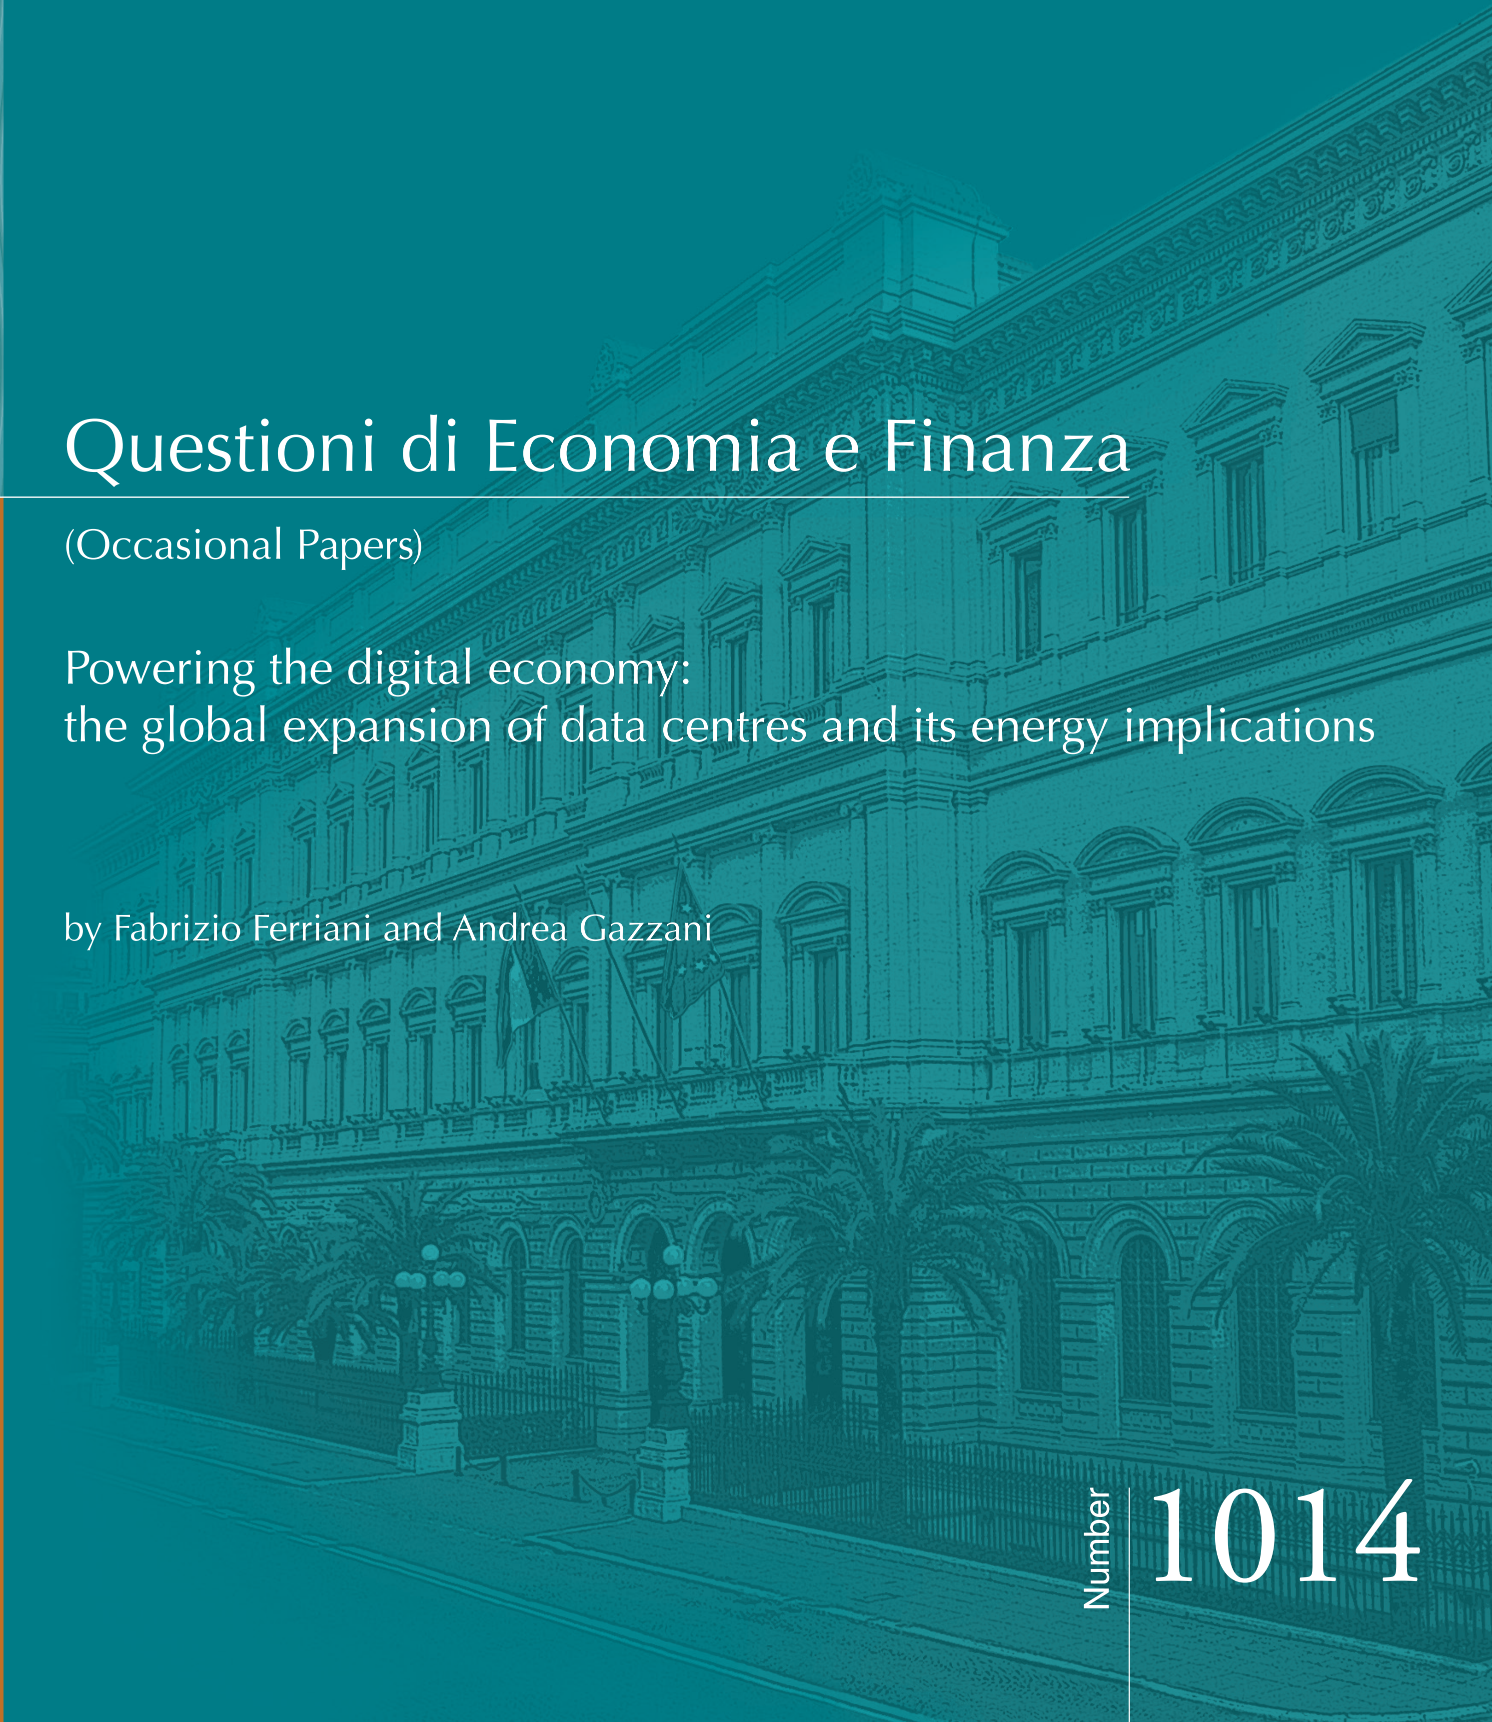

Questioni di Economia e Finanza

(Occasional Papers)

Powering the digital economy: the global expansion of data centres and its energy implications

by Fabrizio Ferriani and Andrea Gazzani

June 2026

Number 1014

BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Powering the digital economy: the global expansion of data centres and its energy implications

by Fabrizio Ferriani and Andrea Gazzani

Number 1014 – June 2026

ISSN 1972-6643 (online)

Designed by the Printing and Publishing Division of Banca d'Italia

POWERING THE DIGITAL ECONOMY: THE GLOBAL EXPANSION OF DATA CENTRES AND ITS ENERGY IMPLICATIONS

by Fabrizio Ferriani* and Andrea Gazzani**

Abstract

The expansion of artificial intelligence and cloud computing is turning data centres into a strategic component of the global economy. This paper provides an empirical assessment of the global expansion of data centres and its energy implications, focusing on the United States and Europe. We combine facility-level data with electricity consumption data and document a sharp rise in global data-centre capacity over the past fifteen years, with the United States leading the market in terms of both the number of facilities and computing capacity. While data-centre electricity consumption remains relatively limited at the global level, it is growing rapidly and already has notable implications in some areas for local electricity demand, grid congestion, and electricity prices, especially in the United States. These dynamics are generating political tensions and have spurred a broad policy debate on grid access, cost allocation, energy security, and strategic autonomy.

JEL Classification: L86, Q41, Q48, O33.

Keywords: data centres, electricity demand, digital infrastructure, energy policy, grid congestion, hyperscalers, technological sovereignty.

DOI: 10.32057/0.QEF.2026.1014

---

* Bank of Italy. Email: fabrizio.ferriani@bancaditalia.it.

** Bank of Italy. Email: andreagiovanni.gazzani@bancaditalia.it.

1 Introduction[1]

The extraordinary rise of artificial intelligence (AI) has fuelled an extensive policy and academic debate on how the digital transition is reshaping the global economy (Aghion et al., 2017; Acemoglu, 2025; Barhoumi et al., 2026). At the heart of this transformation sit data centres, the physical backbone of the digital economy. Yet data centres are not only an economic asset but also an increasing source of pressure on energy consumption and electricity costs. The rapid expansion of data centres is thus emerging as a major challenge for energy policy: as their scale and energy intensity grow, the interaction between digital infrastructure expansion and electricity systems has become a first-order concern for regulators, grid operators, and policymakers.

The interaction between data centre growth and electricity systems operates through both demand and supply channels. On the demand side, the growth of IT capacity translates into a rapid increase in electricity consumption, often concentrated in a limited number of locations, which is therefore capable of driving up wholesale electricity prices and generating sizable local load impacts. On the supply side, the adjustment of generation, transmission and grid interconnection capacity is extremely expensive and inherently slow, owing to long planning and permitting timelines. As a result, electricity demand associated with new data centres may grow faster than the ability of the network to accommodate it, a mismatch that is particularly relevant in areas where data-centre activity is geographically concentrated.

The economic and policy salience of this issue is illustrated by recent regulatory developments in the United States. In January 2026, the Trump Administration urged PJM Interconnection - the largest electricity grid operator in the US, serving around 70 million people across 13 states - to hold an emergency auction requiring technology firms to fund new generation capacity via 15-year power purchase contracts. Under the proposal, tech firms would pay for the new electricity whether they use it or not, ensuring long-term revenue for generators and likely accelerating the funding to develop new natural gas and nuclear capacity. In March 2026, the US administration convened executives from leading technology companies to sign the Ratepayer Protection Pledge, committing them to securing new generation capacity for their data centres and bearing the full cost of the

---

[1] We thank Riccardo Cristadoro, Marco Taboga, and seminar participants at the Bank of Italy for useful comments. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Italy. All remaining errors are our own responsibility.

infrastructure upgrades. Although neither initiative is legally binding, both appear to reflect a political response to mounting pressure over energy affordability ahead of 2026 mid-term elections: US household electricity prices are 32% above their 2019 levels, with even larger increases in data centre-intensive regions (IEA, 2026a).

The expansion of data centres deserves policy attention for at least three reasons. First, it lies at the intersection of two major structural transformations: the energy and the digital transition. The former is associated with the increasing electrification and decarbonisation of economic activity, so that bottlenecks in generation, transmission, and grid connection may significantly shape the pace of digital investment. Second, even when data-centre electricity consumption remains limited in aggregate terms, its strong geographic concentration can generate sizable local effects on electricity prices and network congestion, with distributional impacts on affected communities and the potential to give rise to political and social tensions. Third, because computing capacity is highly concentrated among a small number of operators and jurisdictions (primarily the United States), the geography of data centres has implications that extend beyond energy markets, encompassing industrial policy and strategic autonomy considerations, particularly in light of the reliance on foreign-controlled digital infrastructure.

This paper contributes to that debate by providing a systematic empirical assessment of the global expansion of data centres and of its electricity implications, with a particular focus on the United States and Europe, the regions where most of the data centre IT capacity is located. For this purpose, we combine two complementary sources: a dataset containing detailed information on the location, capacity, ownership, and category of data centres, and national data on electricity consumption at country or state level.

We find that over the last 15 years, the number of data centres worldwide has risen sharply, while global IT capacity has increased even faster, indicating a marked shift toward larger and more power-intensive facilities. The United States is by far the leading market in both the number of facilities and installed IT capacity. Moreover, the IT capacity is increasingly concentrated in larger facilities especially in hyperscalers data centres operated by major cloud providers, designed to efficiently support massive computing and storage, characterised by higher energy efficiency and the ability to scale rapidly in response to demand. We show that, while at the global level electricity consumption by data centres still accounts for a relatively small share of total demand, it is growing rapidly and is expected to increase substantially over the next few years; see Shehabi et al. (2024), EPRI (2024), and De Roucy-Rochegonde and Buffard (2025) among many others. From a policy perspective, however, what stands out is the pronounced spatial heterogeneity in this demand, with a strong territorial concentration of pressures on electricity systems. For instance, in the United States, six states already exhibit data centre electricity demand exceeding 10% of total consumption, with Virginia alone accounting for more than a quarter. In Europe, Ireland represents a clear outlier, with data centres accounting for close to 25% of national electricity consumption, whereas the largest euro-area economies remain below 2%.

Related Literature. Our study complements strands of the economic literature focusing on how the digital transition is reshaping economic growth (Comunale and Manera, 2024; Acemoglu, 2025; Carpinelli et al., 2026), the labour market (Giuntella et al., 2023; Lane et al., 2023; Brynjolfsson et al., 2025), and the green transition (Stern et al., 2025; Bonfiglioli et al., 2025), among many other dimensions. We provide a detailed characterization of how the data centre industry has evolved over the last 15 years, documenting not only its rapid expansion, but also its increasing concentration and growing energy footprint. Consistent with the bottom-up approach commonly adopted in the literature on data centre energy use (Shehabi et al., 2016; Masanet et al., 2020; Shehabi et al., 2024), we rely on granular information at facility level rahter than on aggregate estimates. This allows us to identify where data-centre electricity demand is most concentrated, where pressures on power systems are likely to be stronger, and where the mismatch between digital investment and energy infrastructure may become a binding constraint. These dynamics are likely to be further reinforced by the growing importance of AI-related workloads and by the training of increasingly complex models, which require more powerful and energy-intensive computing infrastructure. Without making causal claims, we document that electricity demand associated with highly concentrated data-centre clusters is already correlated with upward pressure on local electricity prices, an effect that is likely to be understated in more aggregate national or global statistics. Understanding the detailed geographic concentration of data-centre electricity demand is therefore key to assessing risks to energy systems and informing appropriate policy responses. Moreover, our analysis highlights how the ongoing consolidation of both facilities and market power among a limited number of operators may give rise to broader industrial-policy and strategic-autonomy concerns. This issue is particularly salient for Europe (Mariniello, 2026), given the relatively limited availability of domestic capacity for training frontier AI models and the corresponding global dominance of a small number of US-based digital infrastructure firms. In this context, increasing reliance on foreign-controlled computing infrastructure may raise concerns related to technological dependence, resilience, and economic security.

The remainder of the paper is organized as follows. Section 2 documents the expansion of global data-centre capacity and discusses the increasing concentration of the market by size, operator and facility type. Section 3 turns to electricity demand from data centres and presents estimates of their energy footprint for the US and Europe. Finally, Section 4 concludes and discusses the main policy implications.

2 The anatomy of a booming market

Over the last 15 years, the number of global data centres has increased by more than 140%, reaching over 6,500 facilities in Q2 2025, see the upper panel of Figure 1.[2] At the country level, the US hosts the largest number of data centres, rising from 745 sites in Q1 2010 to 1,605 in Q2 2025, accounting for almost 25% of the global number of facilities (Table 1). Europe as a region had 1,920 facilities in Q2 2025, mainly concentrated in Germany, France, and the United Kingdom; also China reached a considerable number of 329 data centres at the end of our sample. An even more informative metric for data centres comes from the amount of running IT capacity.[3] This variable is usually measured in megawatts or gigawatts and refers to the electrical power available to run the computing equipment inside a data centre, which can be used for various purposes, from file storage to AI model training or cloud computing. Since 2010, global IT capacity surged approximately 900%, reaching about 55 GW in mid-2025, see the lower panel of Figure 1. From this perspective, the US dominance is even more striking: in Q1 2010, US data centres offered 2.5 GW of IT capacity, which increased tenfold to around 27.4 GW in Q2 2025, accounting for roughly 50% of global IT capacity; Europe as a region grew from 1.2 GW in Q1 2010 to 10 GW by the end of the period, around 18% of global IT capacity (Table 1).[4]

---

[2] All data-centre statistics are based on Bloomberg–DC Byte data, which cover facility-level, investable, and market-relevant data-centre sites. As a result, the figures reported in this study may differ from alternative sources that also include smaller, non-market-relevant facilities (e.g., micro, or in-building sites). For each reference period, the reported figures include only data centres classified as operational at that time.

[3] All figures in this study are based on live IT capacity in operational data centres at a specific point in time. This capacity may differ from headline capacity, i.e. the declared maximum capacity achievable by a facility.

[4] Further graphs on the evolution in the number of data centre and IT capacity in the US and Europe are available in the Appendix; details on specific countries are available from the authors upon request.

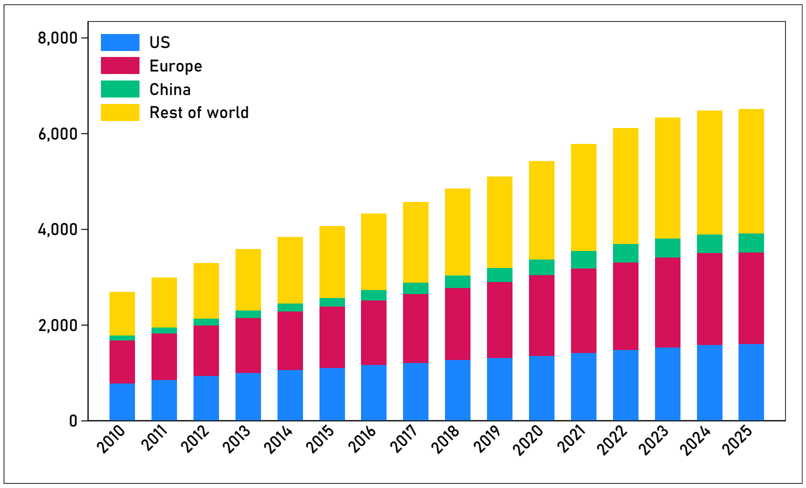

(a) *Number of data centres by region*

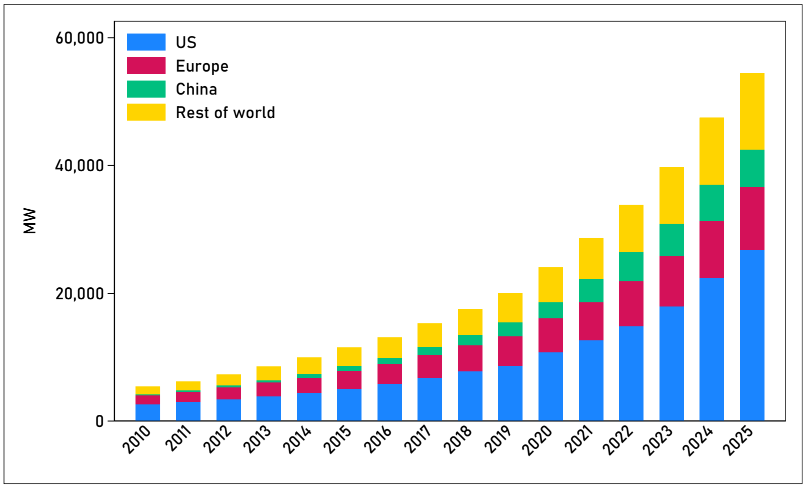

(b) *IT capacity (MW) by region*

FIGURE 1: *Number and IT capacity of data centres; latest data available is 2025 Q2. Europe includes countries in the European Union, but also UK, Norway, and Switzerland.*

Over the last 15 years, the growth of global IT capacity has increasingly been characterized by the concentration in larger, more powerful facilities. Large-scale data centres account for a rising share of worldwide computing capacity, as access to substantial computational power has become the prerequisite for running the most resource-intensive data centre workloads, most notably large-scale AI model training. We consider large data centres as facilities with more than 10 MW of IT capacity, corresponding to the 90th percentile in our sample. Despite accounting for only about one fifth of the global data centre count, these facilities have become increasingly dominant in terms of computing

Number of data centres			IT capacity of data centres (GW)		
Country	2010 Q1	2025 Q2	Country	2010 Q1	2025 Q2
US	745	1605	US	2.53	27.40
China	71	329	China	0.07	5.37
Germany	100	275	Japan	0.36	1.64
France	117	262	UK	0.41	1.59
UK	147	260	Canada	0.09	1.45
Japan	114	224	India	0.10	1.39
Canada	101	214	Australia	0.08	1.31
India	64	202	Germany	0.17	1.26
Australia	55	165	Ireland	0.04	1.23
Indonesia	41	138	Netherlands	0.12	1.11
Italy	76	135	Singapore	0.06	1.11
Netherlands	53	128	France	0.15	0.87
Spain	48	127	South Korea	0.13	0.76
Russia	36	116	Malaysia	0.02	0.71
Brazil	22	115	Sweden	0.03	0.65

Table 1: Number and IT capacity of data centres. Our elaborations on Bloomberg data. Countries are ranked with respect to their relative position on 2025 Q2 for respectively the number and the total IT capacity of data centres.

capacity. Their combined live capacity expanded from approximately 2 GW in Q1 2010 to over 45 GW in Q2 2025, representing more than 80% of total IT capacity by the end of the period (Figure 2, upper panel). The dominance of the United States is even more pronounced in this segment, as it hosts 55% of the global IT capacity deployed in large-scale facilities, while Europe accounts for approximately 16%. A similar pattern emerges for "hyperscaler" data centres, a subset of even larger and more computationally advanced facilities. While the precise definition of "hyperscalers" is not uniform, for simplicity and for the purposes of this analysis we define hyperscale facilities as large data centres, i.e. those with more than 10 MW of IT capacity, operated by one of the following firms: Amazon Web Services, Microsoft, Google, Meta, Oracle, IBM, and CoreWeave. These facilities are generally characterized by substantially higher average IT capacity - around 60 MW on average in our sample, but often exceeding 100 MW - as well as by a strong ability to scale efficiently to support increasingly computationally intensive workloads.[5] This dis-

---

[5] As of 2025 Q2, at the global level, there were 16 facilities with maximum theoretical headline IT capacity above 1 GW, almost all located in the US. The largest active project (Meta data centre currently under construction in Louisiana) displays capacities of up to 2 GW, which could be scaled up to 5 GW when

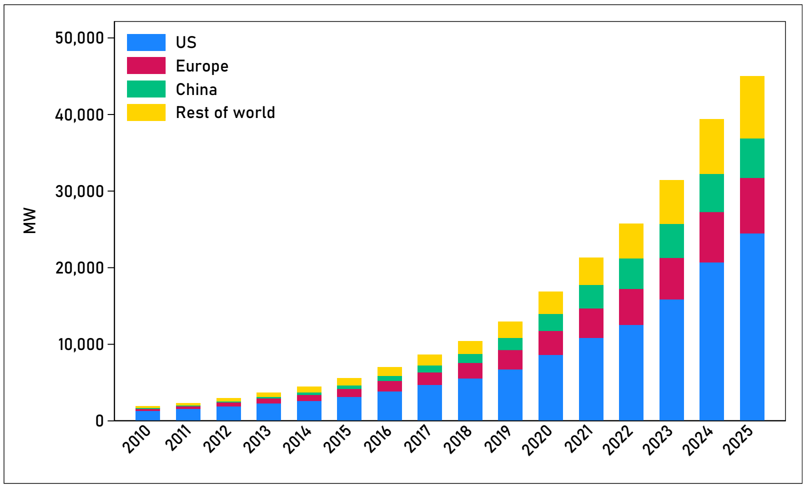

(a) *IT capacity of large data centres (>10 MW) by region*

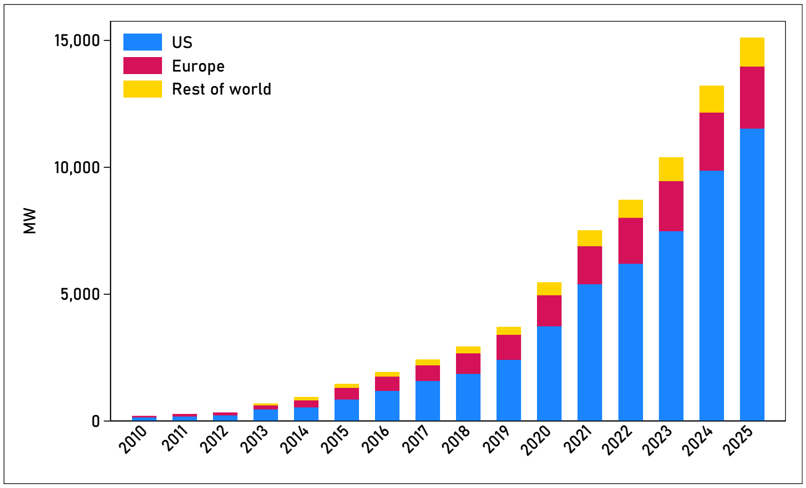

(b) *IT capacity of hyperscalers by region*

FIGURE 2: *IT capacity of large data centres and hyperscalers; latest data available is 2025 Q2. Large data centres are those with IT capacity larger than 10 MW, 90th percentile of the IT capacity in our sample. We define hyperscale facilities as large data centres (more than 10 MW of IT capacity) operated by one of the following firms: Amazon Web Services, Microsoft, Google, Meta, Oracle, IBM, and CoreWeave. Europe includes countries in the European Union, but also UK, Norway, and Switzerland.*

tinction between large data centres and hyperscalers is important, as the latter have been responsible for an increasingly large share of aggregate IT capacity in recent years. In Q2 2025, global IT capacity operated by hyperscalers reached more than 15.2 GW, equivalent to almost 28% of total IT capacity (Figure 2, lower panel). The dominance of the United States in this segment is striking: 75% of the facilities in this category are located in the

---

finished.

US, compared with 16% in Europe.[6]

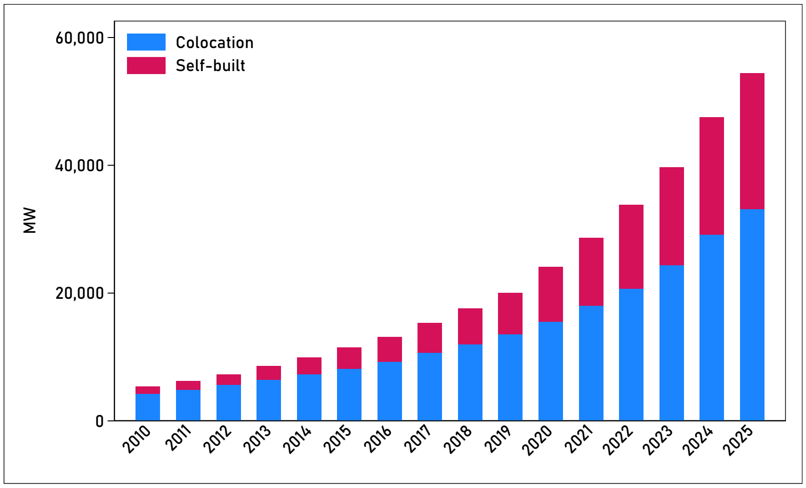

(a) *IT capacity of data centres by facility type*

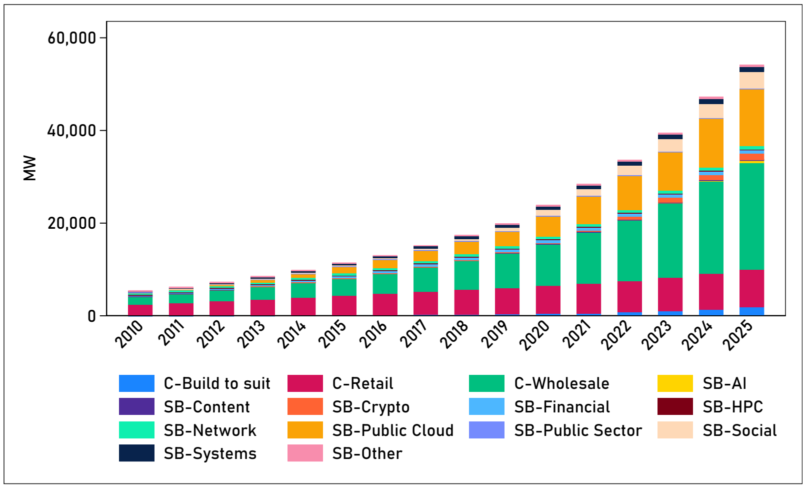

(b) *IT capacity of data centres by facility subcategory*

FIGURE 3: *IT capacity of data centres by facility type and subcategory. Our elaborations on Bloomberg data; latest data available is 2025 Q2. "Colocation" represents data centres operated by one party but where the services are leased by another; "self-built" represents data centres constructed and operated by a company for its own use. In the lower panel, C- stands for colocation, SB- for self-built; colocation subcategories refer to the nature of the customer, while self-built subcategories refer to the business of the operator. Exact definitions of each subcategory are available upon request.*

An important distinction within the data centre ecosystem is between colocation and self-built data centres. Colocation facilities are typically data centres where multiple

---

[6] The graph does not report data for China because, based on our definition of hyperscaler facilities, there are no facilities of this type in China.

clients (service providers) rent space, power, and connectivity from a third-party operator that owns and manages the facility. Conversely, self-built data centres are those in which companies own and operate their infrastructure directly for their own use. Colocation data centres have historically dominated global IT capacity, reflecting most organizations' preference for outsourcing their data centre needs (Figure 3, upper panel). The IT capacity of colocation facilities grew from 4 GW in 2010 to almost 34 GW in Q2 2025, accounting for approximately 61% of global IT capacity at the end of the period. However, the self-built segment has expanded rapidly in recent years, driven primarily by the large-scale investment programs of major technology companies such as Amazon Web Services (AWS), Microsoft, Google and Meta: from around 1 GW of IT capacity in 2020, self-built facilities reached almost 22 GW in Q2 2025, recording an even faster expansion over time. Within the colocation segment, wholesale data centres, where operators sell large volumes of IT capacity to multiple customers, have been the main driver sustaining the growth of IT capacity in recent years and represent by far the largest share of global IT capacity (around 43% in Q2 2025); retail data centres, where operators sell smaller amounts of computing capacity per contract, rank second at almost 15% , see Figure 3. Within the self-built segment, growth has been mainly driven by public cloud operators, which includes companies that own and operate their own data centre infrastructure and then rent out computing resources to external customers via cloud, on a pay-as-you-go basis; these facilities accounted for almost 23% of global IT capacity in mid 2025. Sizable capacity in the self-built segment is also concentrated in social media-oriented facilities (around 6% of total IT capacity, largely attributable to Meta), while, more recently, AI-dedicated facilities have emerged as a rapidly growing subcategory, reflecting the surging demand for computational resources required for large-scale AI model training (1% share of total IT capacity in Q2 2025).

At the global level, there is evidence that a small number of companies, namely AWS, Microsoft, Google, and Meta, account for a share between 60 to 70% of IT computing capacity (Shehabi et al., 2024; De Roucy-Rochegonde and Buffard, 2025; OMDIA, 2025). However, data in this context are difficult to obtain, and only broad sectoral estimates are available. In particular, for colocation facilities, the identity of the company renting computing capacity is generally not disclosed for security and competition reasons and is not reported in our dataset. By contrast, this limitation does not apply to the self-built data centre segment, where the operating company can be identified, although this segment represents only a fraction of the total IT market. Within the self-built segment, the

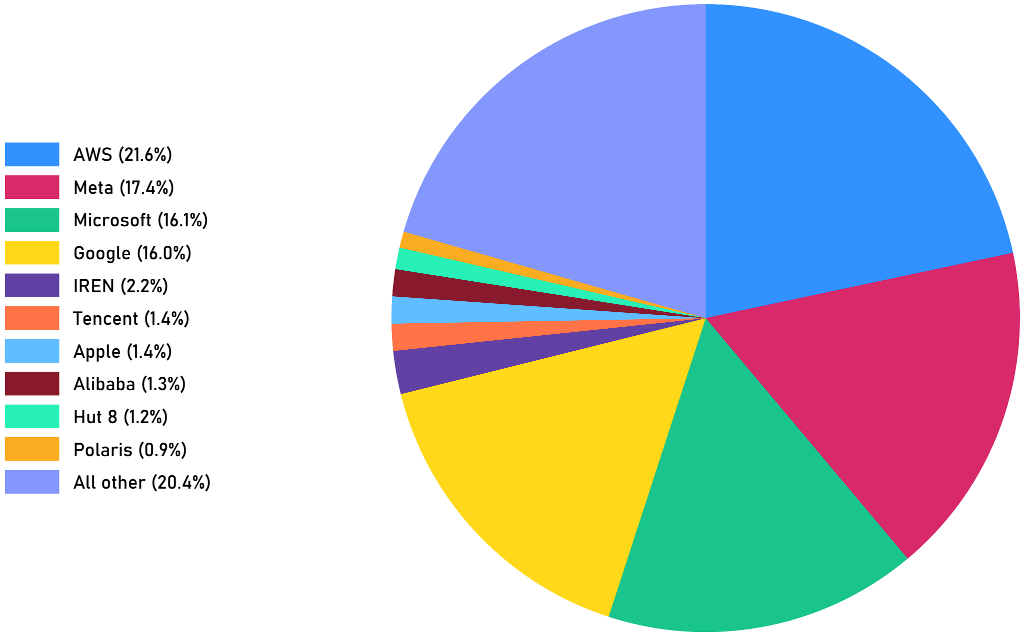

FIGURE 4: *IT capacity of top 10 companies (only self-built data centres). Our elaborations on Bloomberg data as of 2025 Q2. Data refers to IT capacity of self-built data centres only.*

top 10 companies account for around 80% of IT capacity, with the market being even more concentrated among a small number of major technology firms. The top four alone, typically hyperscale operators, account for more than 70% of global self-built capacity: AWS leads with 21.6% (4.8 GW), followed by Meta (3.8 GW), Microsoft (3.5 GW), and Google (3.5 GW). Except for Meta, the data centre capacity operated by these hyperscalers is not used exclusively for internal purposes but is largely devoted to providing public cloud services, whereby computing resources are supplied to customers through shared infrastructure.[7] From an industrial policy perspective, this high level of concentration in self-built data centre capacity also has a strong geographical dimension, being overwhelmingly dominated by a small number of US technology firms.

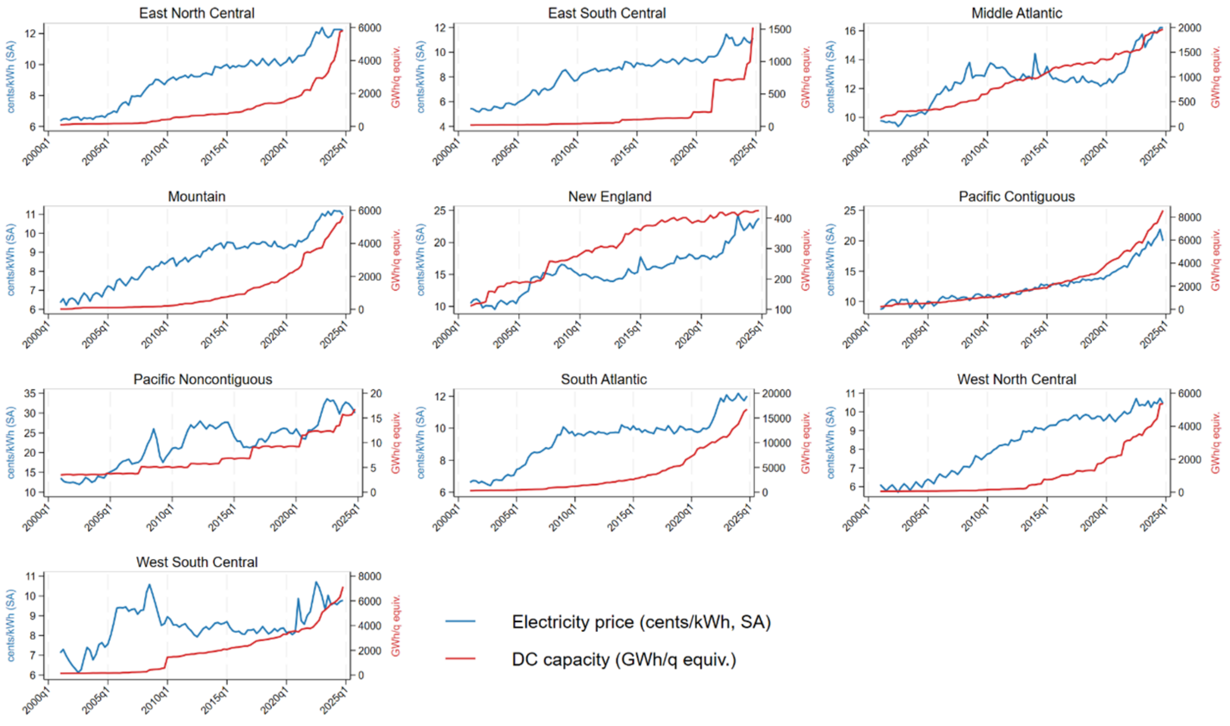

Note: Our elaborations on Bloomberg and EIA data. New England includes Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island, and Vermont; Middle Atlantic includes New Jersey, New York, and Pennsylvania; East North Central includes Illinois, Indiana, Michigan, Ohio, and Wisconsin; West North Central includes Iowa, Kansas, Minnesota, Missouri, Nebraska, North Dakota, and South Dakota; South Atlantic includes Delaware, the District of Columbia, Florida, Georgia, Maryland, North Carolina, South Carolina, Virginia, and West Virginia; East South Central includes Alabama, Kentucky, Mississippi, and Tennessee; West South Central includes Arkansas, Louisiana, Oklahoma, and Texas; Mountain includes Arizona, Colorado, Idaho, Montana, Nevada, New Mexico, Utah, and Wyoming; Pacific Contiguous includes California, Oregon, and Washington; and Pacific Noncontiguous includes Alaska and Hawaii

FIGURE 5: *Electricity prices of US census regions and data centre IT capacity*

3 Power hungry: The electricity footprint of data centres

Globally, data centre electricity consumption has grown by around 12% per year since 2017, more than four times faster than total electricity consumption (IEA, 2025). In 2025, according to IEA estimates, data centres accounted for around 485 TWh, or slightly more than 1.5% of global electricity consumption (IEA, 2026a). Similar figures emerge from other data providers, with Bloomberg pointing to almost 450 TWh (Bloomberg, 2026). To illustrate the scale, a 10 MW data centre may consume approximately 80 GWh annually,

---

[7] Meta does not generally sell cloud compute, but it generates revenue from data centres by providing social media services that are used to sell advertisements.

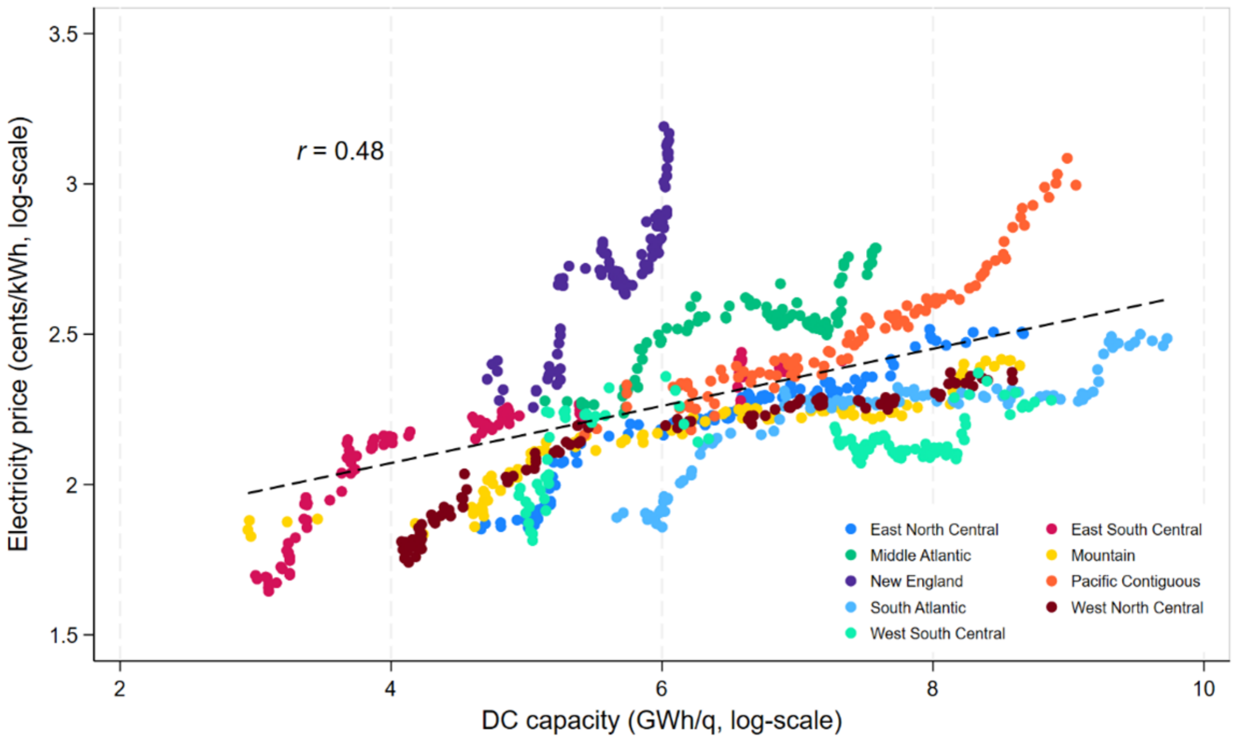

FIGURE 6: *Scatterplot of electricity prices and data centre capacity by US Census regions. Our elaborations on Bloomberg and EIA data. The graph reports the correlation coefficient (r) between the two variables. Sample: 2001-2024.*

which is broadly comparable to the annual electricity consumption of around 20,000 European households or 25,000 electric vehicles.[8] A hyperscale, AI-focused data centre can reach much higher levels of capacity and electricity use: for instance, the largest facility currently under development by Meta could consume, at full capacity, around 16 TWh per year, exceeding the total electricity consumption of Slovenia.

The United States, Europe, and China account for around 85% of global data centre electricity consumption (IEA, 2025). In the United States, electricity consumption from data centres amounts to around 180 TWh, approximately 40% of the global total, and represents more than 4% of total US electricity consumption (Shehabi et al., 2024). Europe accounts for about 15% of global data centre electricity consumption (less than 2% of total electricity consumption in Europe), while China's share is roughly 25% (around 1.1% of total electricity consumption in China). Looking ahead, global electricity consumption

---

[8] As a benchmark, we assume that a typical European household consumes about 3–4 MWh per year, while an electric vehicle consumes around 3 MWh annually.

from data centres is expected to double by 2030. Baseline estimates range between 935 and 950 TWh according to IEA (2026b) and Bloomberg (2026), respectively, and could exceed 1,000 TWh in a lift-off scenario. This level would be broadly comparable to current electricity consumption in Japan and would correspond to about 3.6% of global electricity demand. Available estimates for the United States, where most global information technology capacity is located, vary substantially depending on scenario assumptions and on the expected trajectory of information technology capacity growth. Despite some heterogeneity, they consistently indicate that data centres could account for approximately 7 percent to as much as 13 percent of total electricity consumption in the most bullish scenarios, see EPRI (2024), Shehabi et al. (2024), and De Roucy-Rochegonde and Buffard (2025).

While at the global level electricity demand from data centres remains relatively modest, there are regions in which these facilities are highly concentrated and their share of total electricity consumption is substantial. As shown in Figure 5, the gradual expansion of data centre capacity in recent years, and the corresponding pressure on local electricity supply, has occurred alongside a rise in electricity prices for households, which increased by 5% year over year in 2025 and are now 32% higher than in 2019 (IEA, 2026a). Figure 6 plots average retail electricity prices against data centre capacity across United States areas and shows a positive and statistically significant correlation. Although this cross-regional association does not establish a causal relationship, the pattern is nevertheless consistent with the view that concentrated data-center expansion can increase pressure on local power systems. Recent model-based estimates confirm this evidence and point to a potential increase in electricity prices by up to 9% in the US as the country is expected to experience the largest surge in AI-driven demand for electricity (Bogmans et al., 2026).

We rely on the granular dimension of our dataset to map where data centre electricity demand is most concentrated across the United States and European countries, and where pressures on power systems are likely to be stronger, ultimately affecting electricity prices. To do so, we convert data centre IT capacity into electricity consumption using the following formulation:

$$\text{Electricity consumption} = \text{Live IT capacity} \times \text{utilization factor} \times 24 \times 365 \times \text{PUE} \quad (1)$$

where live IT capacity denotes the power capacity of IT equipment that is online in a given period. The utilization factor captures the ratio of actual power consumption to

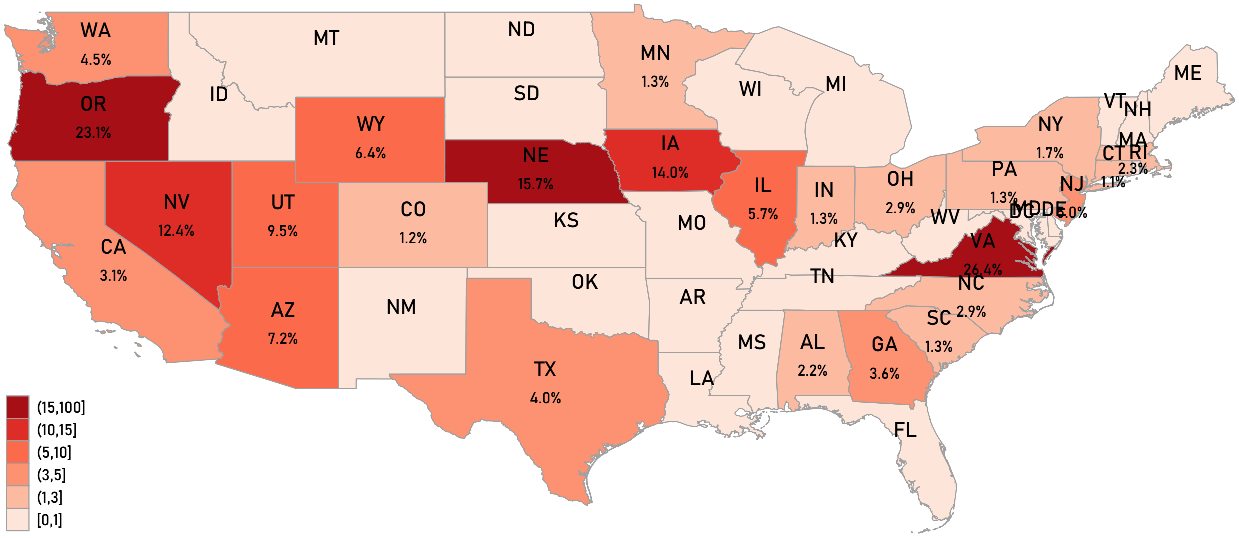

FIGURE 7: *Share of electricity demand from data centres in the US. Our elaborations on Bloomberg and EIA data. Estimates refer to available data as of 2024, the latest available year for US states electricity consumption. For readability we do not report shares smaller than 1%*

the rated capacity of IT equipment. Since equipment that is online does not necessarily operate at its maximum rated power, owing for example to idle servers or partial load operation, we assume an average utilization factor of 0.7. The PUE (power usage effectiveness) is defined as the ratio of total facility electricity consumption to IT equipment electricity consumption and captures all non-IT energy uses, such as cooling, lighting, and power conversion. We assume an average PUE equal to 1.3. We emphasize that both the utilization factor and PUE can vary significantly across facilities, so the figures reported in this section should not be interpreted as actual consumption or official measurements, but rather as plausible estimates based on a set of reasonable assumptions. Greater reporting of these parameters by major operators could help fill existing data gaps in the future and provide more refined estimates of electricity absorption (Masanet et al., 2020).

Figure 7 displays the share of electricity demand accounted for by data centres across US states. In 2024, six states had data centres accounting for around 10% or more of electricity demand, with Virginia leading at roughly 26%. In Virginia, the Ashburn area,

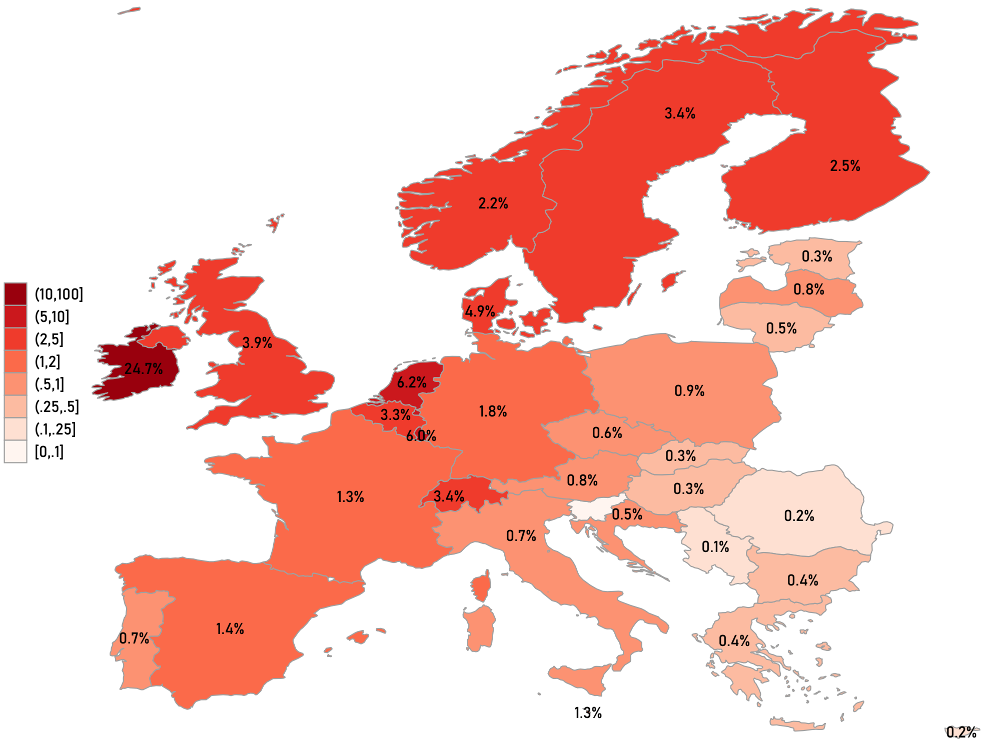

FIGURE 8: *Share of electricity demand from data centres in European countries. Our elaborations on Bloomberg and Ember data. Estimates refer to available data as of 2024, the latest available year for European countries electricity consumption.*

often referred to as "Data centre alley," hosts one of the world's largest concentrations of data centres. This concentration is closely tied to the region's role as a major internet exchange hub, through which a substantial share of global internet traffic is routed. Other states with substantial shares of data centre electricity consumption are Oregon, Nebraska, Nevada, Iowa, and Utah.[9] The corresponding analysis for Europe is presented in Figure 8. Here, Ireland stands out as a clear outlier with data centres accounting for almost a quarter of national electricity demand. Among the largest European economies, the UK is the only one exceeding the 3% threshold, while the four largest EU area countries (Germany, France, Italy, and Spain) remain below 2%.

To meet this massive electricity demand from data centres there are multiple sources and technologies, each characterized by different performance, cost and construction

---

[9] With the only exception of Oregon, our estimates are broadly similar to those reported in EPRI (2024) where scenario analyses are also reported for data centre electricity absorption in selected US states. Again, this discrepancy may reflect, among other factors, different assumptions regarding PUE and the utilization factor.

timelines. Currently, the generation technologies that can be deployed within a timeframe broadly consistent with data-center construction times are mainly solar PV and gas turbines, but with important differences: solar PV is quickly scalable but non-dispatchable, while gas turbines provide dispatchable capacity at the cost of greater exposure to gas prices and emissions. Large hyperscalers are adopting diversified strategies to ensure adequate electricity supply, including signing long-term power purchase agreements (PPAs), primarily from renewable sources, entering into nuclear partnerships, supporting the restart of inactive gas-fired plants, and investing in technological improvements aimed at enhancing energy efficiency and reducing energy intensity (EPRI, 2024; Masanet et al., 2020).

However, securing sufficient electricity generation is only one side of the challenge, as the expansion of data centre capacity also faces increasing risks from grid connection delays, particularly in regions of strong demand growth. When a new data centre applies for connection to the electricity grid, it enters an interconnection queue, where grid operators assess available capacity and determine whether network upgrades are required to accommodate the additional load; this process can take three to seven years or longer, see IEA (2025). In recent years, policymakers have adopted a range of approaches to address these challenges. On the one hand, some jurisdictions have sought to facilitate the grid connection of data centres (IEA, 2026a). In the United States, the Department of Energy and FERC have moved toward reforms aimed at ensuring the timely and orderly interconnection of large loads (e.g. U.S. Department of Energy, 2025). The United Kingdom introduced AI Growth Zones intended to reduce planning barriers and improve access to power for strategically important AI and data centre projects (Department for Science, Innovation and Technology, 2025). In Canada, the system operator has introduced interim large load integration measures, including reliability based limits and technical requirements for new large load connections. On the other hand, some countries have imposed moratoria or restrictions on new data centre connections. For example, Irish authorities introduced a moratorium on new grid connections around Dublin since 2021 and, in 2025, implemented new requirements obliging data centres seeking connection to provide dispatchable generation and/or storage capacity matching the facility’s maximum capacity. Similar restrictions - also reflecting land and water constraints - have been adopted in other jurisdictions, including the Netherlands and Singapore (Sandalow et al., 2025).

4 Conclusions and Policy Challenges

Data centres are the physical backbone of the digital economy, translating AI-related capital spending into computational capacity and, ultimately, into productive output. Our study has reviewed the recent market dynamics, scale, and geographic concentration of this infrastructure. Although data centres remain a relatively small component of global electricity demand in aggregate terms, their energy footprint is growing rapidly and is already sizable in some geographical areas, putting the future evolution of the data centre industry at the core of energy policy debates.

Several policy challenges arise from the future development of the data centre ecosystem. The first policy challenge is distributional. While the benefits of digital infrastructure accrue broadly, the costs (higher electricity prices, grid congestion, and network upgrades) fall disproportionately on households and businesses in host regions. As our subnational analysis shows, data centres already absorb a substantial share of local electricity demand in some areas, and political tensions over cost allocation are mounting. As such, the expansion of digital infrastructure cannot be treated as a purely private investment decision. Regulators need cost-allocation rules that prevent large-load customers from shifting network costs onto residential and small business consumers, while preserving incentives for efficient siting in areas with spare grid capacity and for operator investment in self-supply generation and renewable sources to reduce system costs.

The second policy challenge concerns the interaction between the digital and energy transitions. Under faster-growth scenarios, data-centre expansion could intensify competition for electricity and grid access, especially where infrastructure investment fails to keep pace with AI-driven demand. The key tension is temporal: data centres can often be developed far more quickly than the transmission, distribution, and generation infrastructure needed to serve them (Chen et al., 2025). If grid expansion and clean power deployment lag behind, AI-related electricity demand may raise system costs, increase reliance on fossil generation, or delay decarbonization objectives.

A third challenge relates to the planning of electricity-system development under greater uncertainty. Forecasts of data-centre electricity demand vary widely, depending on future AI adoption, hardware efficiency, and broader technological developments, with the interplay between these factors evolving rapidly (IEA, 2026b; Bogmans et al., 2026; Masanet et al., 2020). While efficiency gains may reduce the electricity required for each AI task, the expansion of IT applications, especially toward more energy-intensive uses such as autonomous agents, may increase pressure on the electricity system. Policymakers therefore need adaptive planning frameworks based on enhanced electricity-demand forecasting, better disclosure of facility-level consumption, local demand monitoring, and regular updates to grid-investment plans as AI technologies and business needs evolve.

The final two challenges are particularly relevant for Europe, where data-centre expansion intersects with concerns over technological dependence and energy security. Global IT capacity is concentrated both geographically and among a small number of U.S.-based hyperscalers. Europe accounts for less than a fifth of global capacity, implying heavy reliance on foreign-controlled infrastructure for cloud services, frontier AI training, and data governance. As AI capabilities become increasingly central to economic productivity and national security, reducing these dependencies will require coordinated policies on digital infrastructure investment, grid planning, and technological sovereignty (Mariniello, 2026).

Lastly, the fifth challenge is energy-security exposure. Because data centres require abundant, reliable, and competitively priced electricity, their location is sensitive to energy-security conditions. This is particularly relevant for Europe, where the energy crisis following Russia's invasion of Ukraine showed how geopolitical shocks can propagate through gas markets, electricity prices, and industrial competitiveness (Ferriani and Gazzani, 2023; Emiliozzi et al., 2025). For data centres, the implication is direct: regions with volatile electricity prices, import-dependent energy systems, or uncertain grid access may face higher operating costs for existing facilities and struggle to attract new capacity.

References

ACEMOGLU, D. (2025): "The simple macroeconomics of AI," *Economic Policy*, 40, 13–58.

AGHION, P., B. F. JONES, AND C. I. JONES (2017): "Artificial intelligence and economic growth," .

BARHOUMI, K., F. A. DE CARVALHO, M. GORBANYOV, Y. KIDO, D. KOLL, D. OSTOJIC, B. SHANG, N. TAMIRISA, S. TOMS, E. DABLA-NORRIS, A. NGUYEN, AND Y. ZHAO (2026): "Global Economic and Financial Implications of Artificial Intelligence," *IMF Notes*, 2026, 1.

BLOOMBERG (2026): "AI Data Centers Fuel Quicker Growth In Power Demand," .

BOGMANS, C., G. GANPUREV, P. GOMEZ-GONZALEZ, G. MELINA, A. PESCATORI, AND S. THUBE (2026): "Power hungry: How ai will drive energy demand," *Energy Economics*, 109278.

BONFIGLIOLI, A., R. CRINÒ, M. FILOMENA, AND G. GANCIA (2025): "Data, power and emissions: The environmental cost of AI," .

BRYNJOLFSSON, E., D. LI, AND L. RAYMOND (2025): "Generative AI at work," *The Quarterly Journal of Economics*, 140, 889–942.

CARPINELLI, L., F. NATOLI, AND M. TABOGA (2026): "Artificial Intelligence and the US Economy: An Accounting Perspective on Investment and Production," *Bank of Italy occasional paper*.

CHEN, X., X. WANG, A. COLACELLI, M. LEE, AND L. XIE (2025): "Electricity demand and grid impacts of AI data centers: Challenges and prospects," *arXiv preprint arXiv:2509.07218*.

COMUNALE, M. AND A. MANERA (2024): "The economic impacts and the regulation of AI: A review of the academic literature and policy actions," .

DE ROUCY-ROCHEGONDE, L. AND A. BUFFARD (2025): "AI, Data Centers and Energy Demand," .

DEPARTMENT FOR SCIENCE, INNOVATION AND TECHNOLOGY (2025): "Delivering AI Growth Zones," *UK Government policy paper*.

EMILIOZZI, S., F. FERRIANI, AND A. GAZZANI (2025): “The European energy crisis and the consequences for the global natural gas market,” *The Energy Journal*, 46, 119–145.

EPRI (2024): “Powering intelligence: Analyzing artificial intelligence and data center energy consumption,” *EPRI white paper*.

FERRIANI, F. AND A. GAZZANI (2023): “The invasion of Ukraine and the energy crisis: Comparative advantages in equity valuations,” *Finance Research Letters*, 58, 104604.

GIUNTELLA, O., J. KÖNIG, AND L. STELLA (2023): “Artificial intelligence and workers’ well-being,” .

IEA (2025): “Energy and AI,” .

——— (2026a): “Electricity 2026 - Analysis and forecast to 2030,” .

——— (2026b): “Key Questions on Energy and AI,” .

LANE, M., M. WILLIAMS, AND S. BROECKE (2023): “The impact of AI on the workplace: Main findings from the OECD AI surveys of employers and workers,” *OECD Social, employment and migration working papers*.

MARINIELLO, M. (2026): “Europe’s artificial intelligence strategy should be built on European strengths,” *Bruegel First glance*.

MASANET, E., A. SHEHABI, N. LEI, S. SMITH, AND J. KOOMEY (2020): “Recalibrating global data center energy-use estimates,” *Science*, 367, 984–986.

OMDIA (2025): “Global cloud infrastructure spending rose 21% in Q1 2025,” *OMDIA Newsroom*.

SANDALOW, D., M. WU, AND A. YUAN (2025): “SUSTAINABLE DATA CENTERS ROADMAP,” *ICEF publication*.

SHEHABI, A., A. NEWKIRK, S. J. SMITH, A. HUBBARD, N. LEI, M. A. B. SIDDIK, B. HOLECEK, J. KOOMEY, E. MASANET, AND D. SARTOR (2024): “2024 united states data center energy usage report,” .

SHEHABI, A., S. SMITH, D. SARTOR, R. BROWN, M. HERRLIN, J. KOOMEY, E. MASANET, N. HORNER, I. AZEVEDO, AND W. LINTNER (2016): “United states data center energy usage report,” .

STERN, N., M. ROMANI, R. PIERFEDERICI, M. BRAUN, D. BARRACLOUGH, S. LINGESWARAN, E. WEIRICH-BENET, AND N. NIEMANN (2025): "Green and intelligent: the role of AI in the climate transition," *npj Climate Action*, 4, 56.

U.S. DEPARTMENT OF ENERGY (2025): "Section 403 Proposal on the Interconnection of Large Loads to the Interstate Transmission System," October 23 Proposal to FERC.

A Additional results

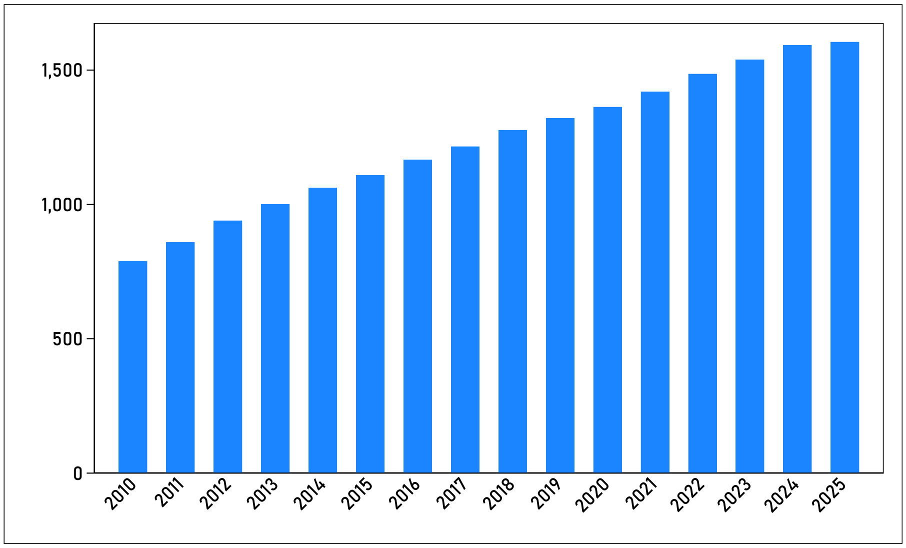

(a) *Number of data centres in the US*

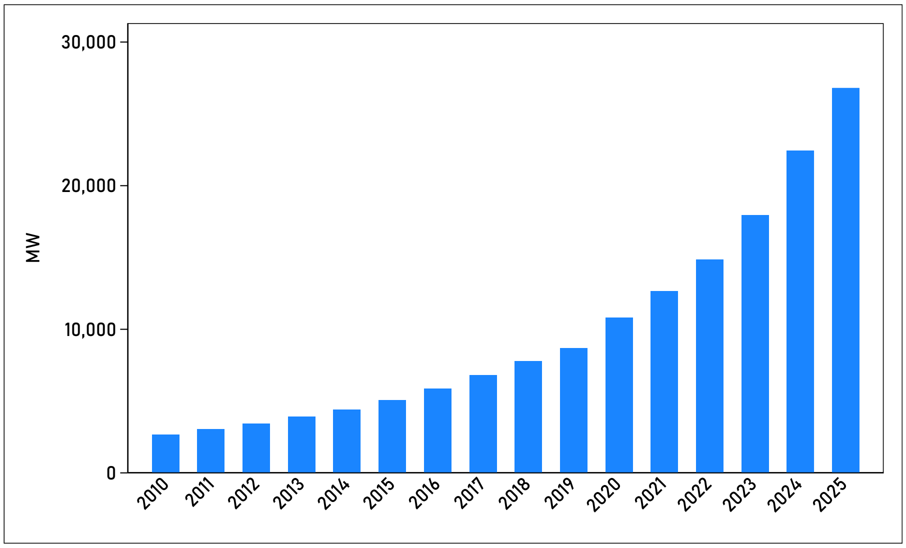

(b) *IT capacity (MW) in the US*

FIGURE A.1: *Number and IT capacity of data centres in the US; latest data available is 2025 Q2.*

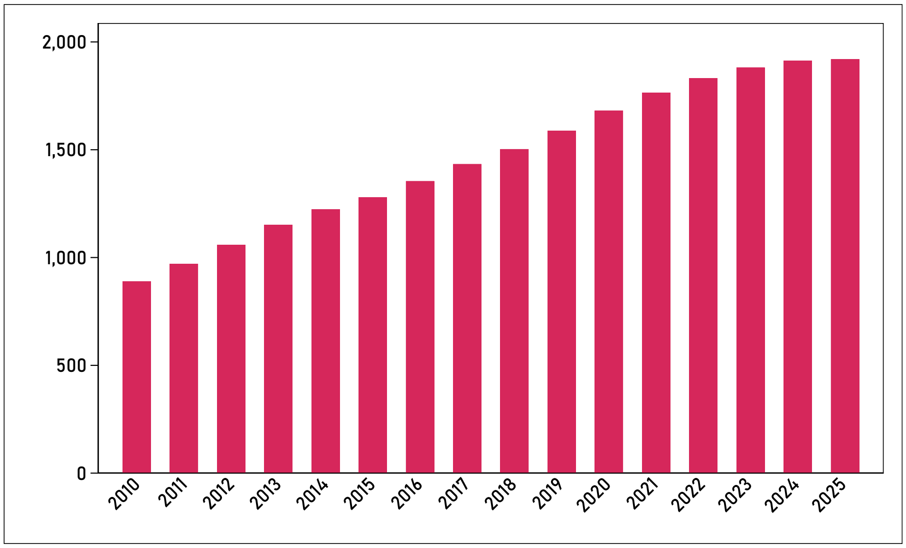

(a) *Number of data centres in Europe*

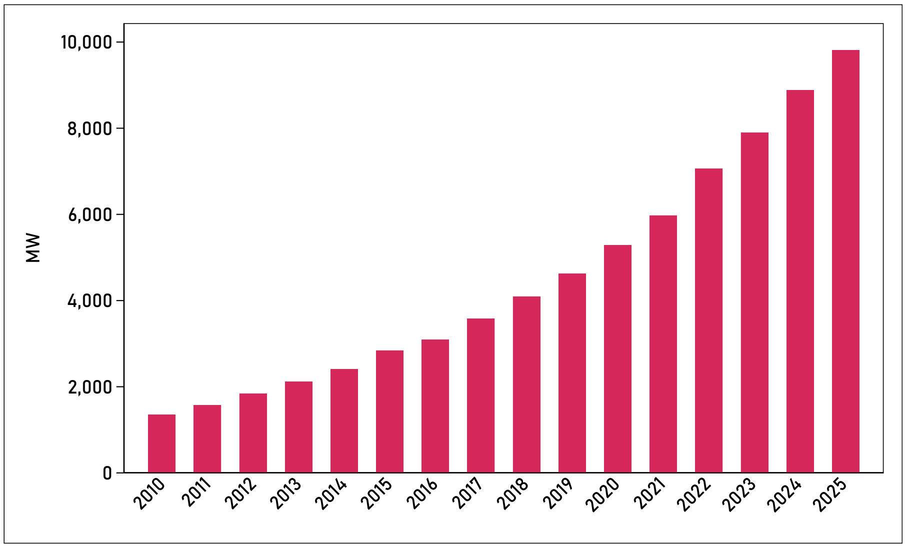

(b) *IT capacity (MW) in Europe*

FIGURE A.2: *Number and IT capacity of data centres in Europe; latest data available is 2025 Q2. Europe includes countries in the European Union, but also UK, Norway, and Switzerland.*

Number of data centres			IT capacity of data centres (GW)		
Country	**2010 Q1**	**2025 Q2**	**Country**	**2010 Q1**	**2025 Q2**
Germany	100	275	UK	0.41	1.59
France	117	262	Germany	0.17	1.26
UK	147	260	Ireland	0.04	1.23
Italy	76	135	Netherlands	0.12	1.11
Netherlands	53	128	France	0.15	0.87
Spain	48	127	Sweden	0.03	0.65
Switzerland	39	93	Spain	0.06	0.48
Sweden	34	77	Norway	0.01	0.43
Poland	26	64	Belgium	0.02	0.36
Austria	23	51	Finland	0.01	0.32
Ireland	15	51	Switzerland	0.03	0.31
Belgium	25	46	Italy	0.06	0.29
Denmark	15	46	Denmark	0.01	0.27
Finland	12	36	Iceland	0.00	0.23
Norway	7	36	Poland	0.02	0.21
Portugal	15	28	Austria	0.02	0.09
Czech Republic	13	25	Portugal	0.01	0.06
Greece	11	19	Luxembourg	0.01	0.05
Romania	9	18	Czech Republic	0.01	0.05
Latvia	10	17	Greece	0.01	0.03
Bulgaria	4	16	Bulgaria	0.00	0.02
Luxembourg	11	16	Hungary	0.02	0.02
Croatia	7	15	Romania	0.01	0.01
Lithuania	4	13	Croatia	0.00	0.01
Slovakia	7	12	Slovakia	0.00	0.01
Estonia	3	9	Lithuania	0.00	0.01
Cyprus	4	8	Latvia	0.00	0.01
Hungary	6	8	Serbia	0.00	0.01
Iceland	0	7	Malta	0.00	0.01

Table A.1: Number and IT capacity of data centres in the main European countries. Countries are ranked with respect to their relative position on 2025 Q2 for respectively the number and the total IT capacity of data centres.

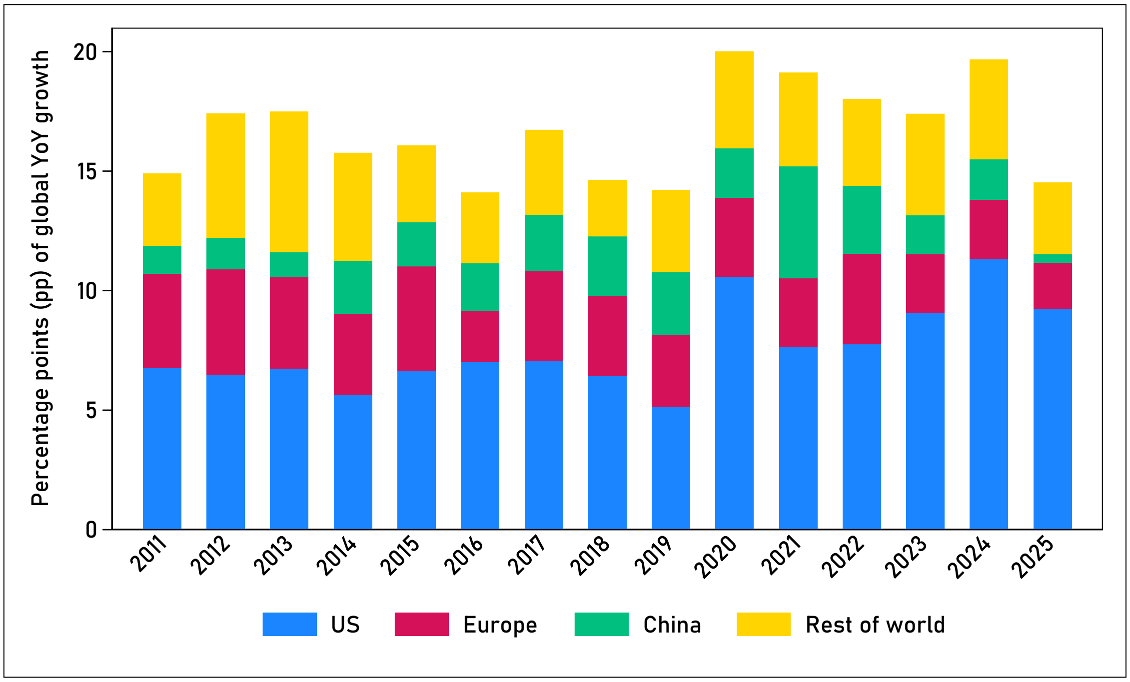

FIGURE A.3: *Regional contributions (pp) to IT capacity growth (MW). Europe includes countries in the European Union, but also UK, Norway, and Switzerland.*

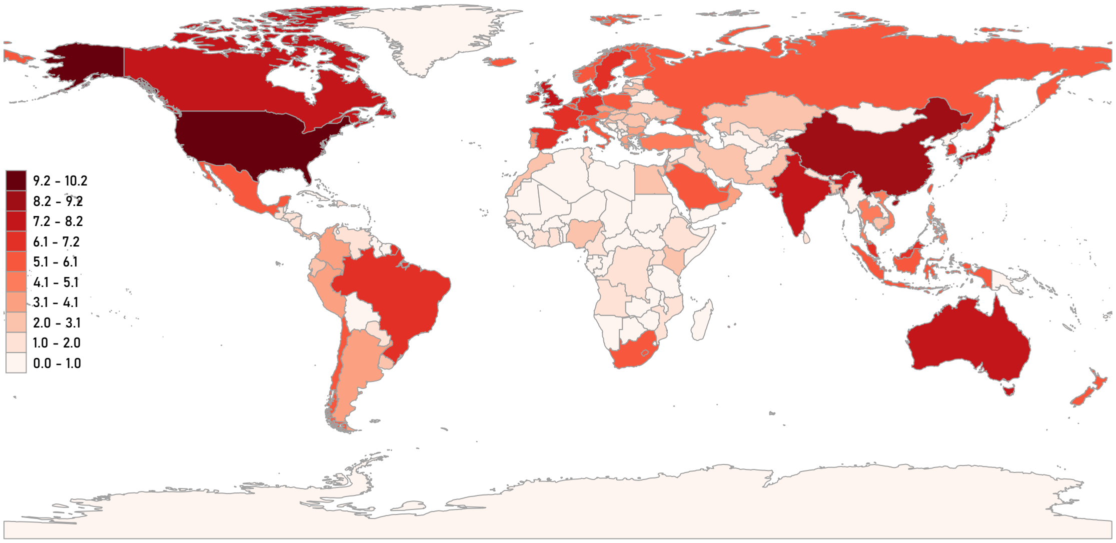

FIGURE A.4: *Country-level IT capacity; data as of 2025 Q2 in log scale*

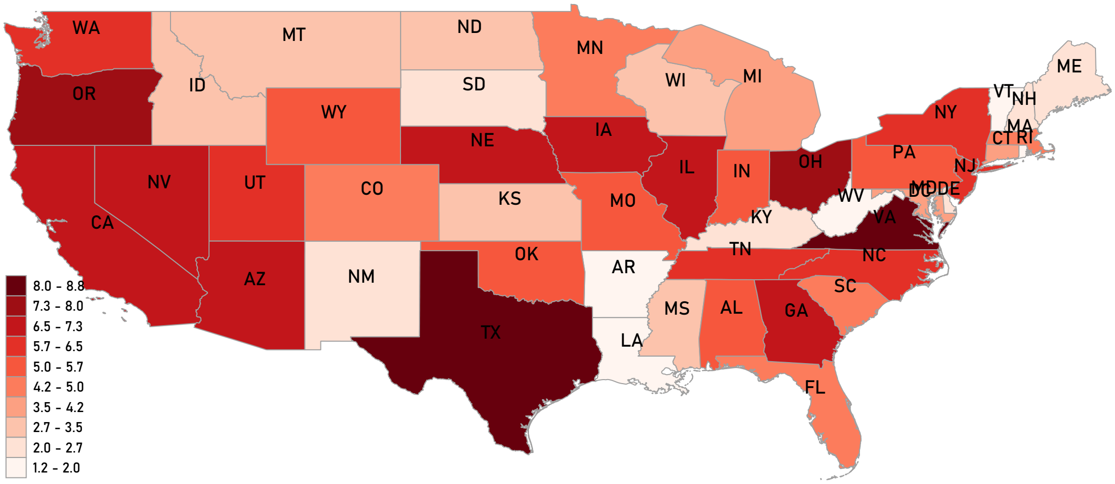

FIGURE A.5: *US state-level IT capacity; data as of 2025 Q2 in log scale*

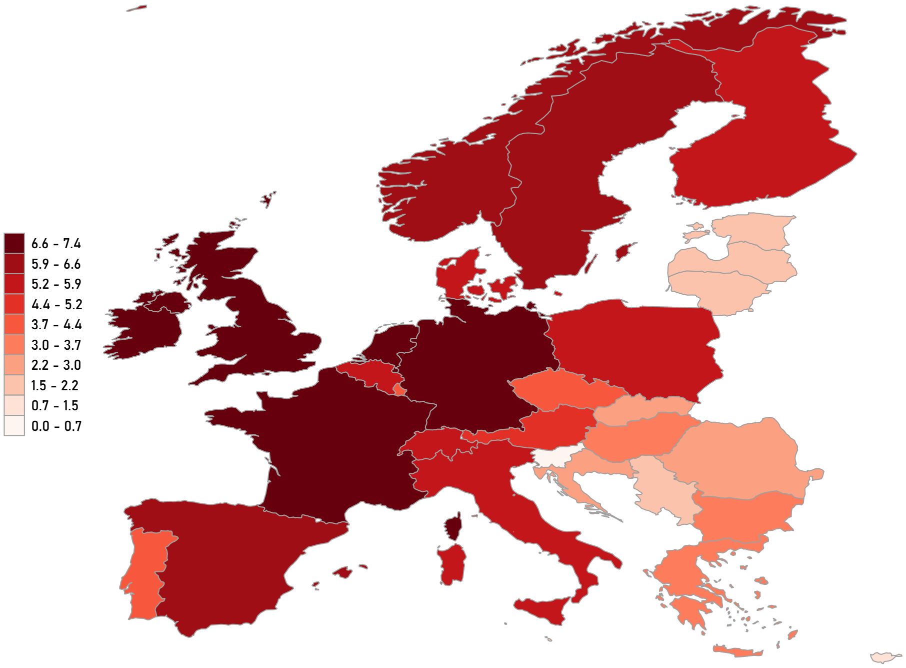

FIGURE A.6: *EU country-level IT capacity; data as of 2025 Q2 in log scale*

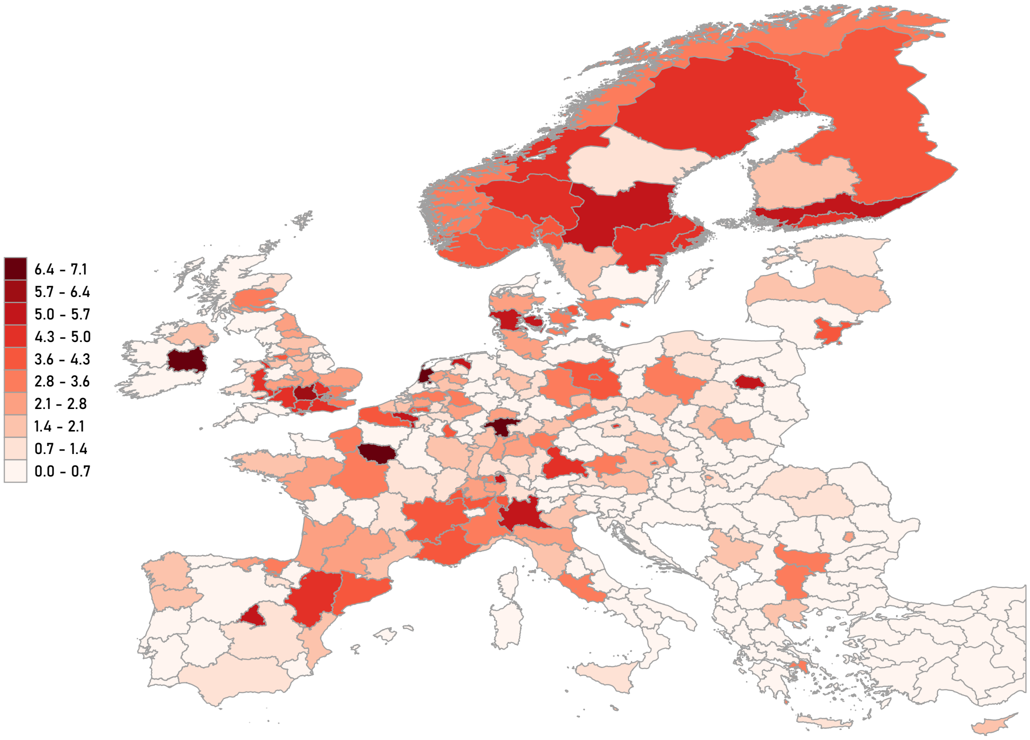

FIGURE A.7: *EU NUTS2-level IT capacity; data as of 2025 Q2 in log scale.*