TD Asset Management

Investor Knowledge 10 Minutes

AI and Electricity Demand:

Will the Power-Hungry Machines be Satiated?

Artificial Intelligence (AI) is a highly disruptive, general-purpose technology (GPT) with compute requirements that are likely to continue growing exponentially well into the 2030s. Alongside the rising share of electric vehicles on roads and the trend towards the onshoring of manufacturing, AI is beginning to create an electricity demand surge in the U.S. This surge has critical implications for infrastructure investment, especially as the AI-induced boom follows an extended period of flat U.S. electricity consumption.

AI models possess a seemingly insatiable thirst for electricity. To illustrate, between 2020 and 2022, annual electricity demand from Microsoft Corporation, Alphabet Inc. (Google), Amazon.com, Inc., and Meta Platforms, Inc. grew 58%. Most of this surge was driven by data center (DC) build, with Microsoft Corporation alone currently adding a new DC roughly every three days. DCs are power-hungry machines, and we expect their electricity demand to triple over the next decade.

The electricity demand boom will stress existing infrastructure, including generation capacity, transformers and the transmission and distribution (T&D) grid. Without massive investment, as well as transformational innovations (for example, regarding battery storage, small modular reactors, and more efficient semiconductors), there is a rising risk that electricity demand may race ahead of supply. This could create a chokepoint that impedes AI progress, with negative consequences for innovation, productivity, security, and not to mention equity markets.

At TD Asset Management ("TDAM", "we", "our") we see five important questions to ask about the impending electricity demand boom:

- **Why is electricity load growth increasing now, after having been flat from 2007-2022?**
- **Will power supply constraints impede AI progress?**
- **Is AI unwelcome news for CO2 emissions and climate change?**
- **What is the key risk of this boom?**
- **What are the implications for investors?**

1 Why is electricity load growth increasing now, after having been flat from 2007-2022?

We are still in the early days of AI diffusion, analogous to where the personal computer was in 1980, the internet in 1995 or electricity in 1900. This is important because technology companies often believe that throwing more compute and data at an AI model is the best way to improve performance. As a result, the compute used to train AI models has been increasing exponentially over the last fifteen years with no signs of slowing down.

Training a foundational AI model requires an enormous amount of compute and hence, electricity. However, that is not the end of the story as inference is also power hungry and becoming ever more so. To illustrate, a traditional Google query requires about 0.3 watt-hours (Wh). However, a comparable ChatGPT request would necessitate 2.9 Wh, a roughly ten-fold increase. Future use cases for AI will increasingly emphasize compute-intensive capabilities which consume multiples more energy **(Figure 1)**.

Figure 1: Power hungry - Inference energy used for various tasks (as a multiple of energy for a simple text classification, log scale)

Generating images requires more than 1,000 times the energy of text classification. The energy demands of sound and video generation will be thousands of times greater still.

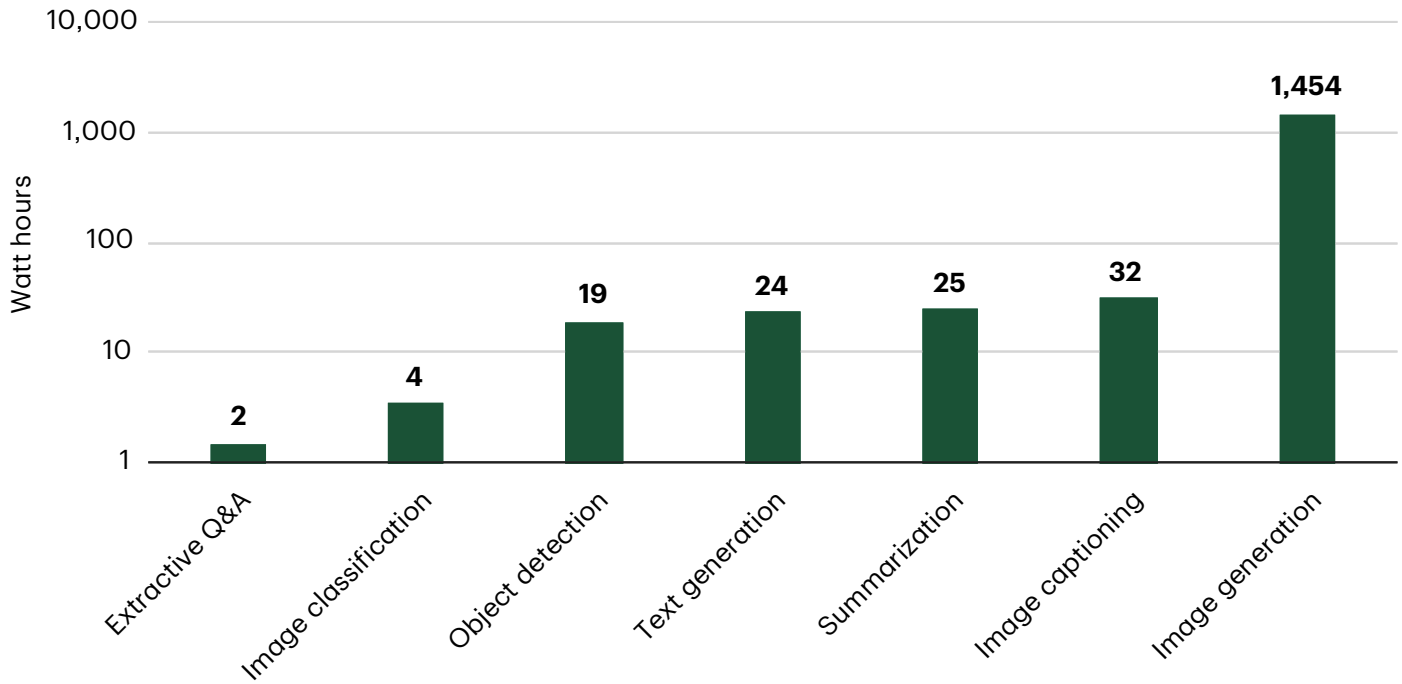

Source: "Power hungry processing: Watts driving the cost of AI deployment?" by A. Luccioni et al, Hugging Face and Carnegie Mellon, 2024.

Turning to DCs, there are currently more than 8,000 globally, with the market size of DCs projected to soar **(Figure 2)**.

Figure 2: Global DCs, forecasted market size (USD, bn)

Expected to almost triple, representing a compounded annual growth rate (CAGR) of 12.1%.

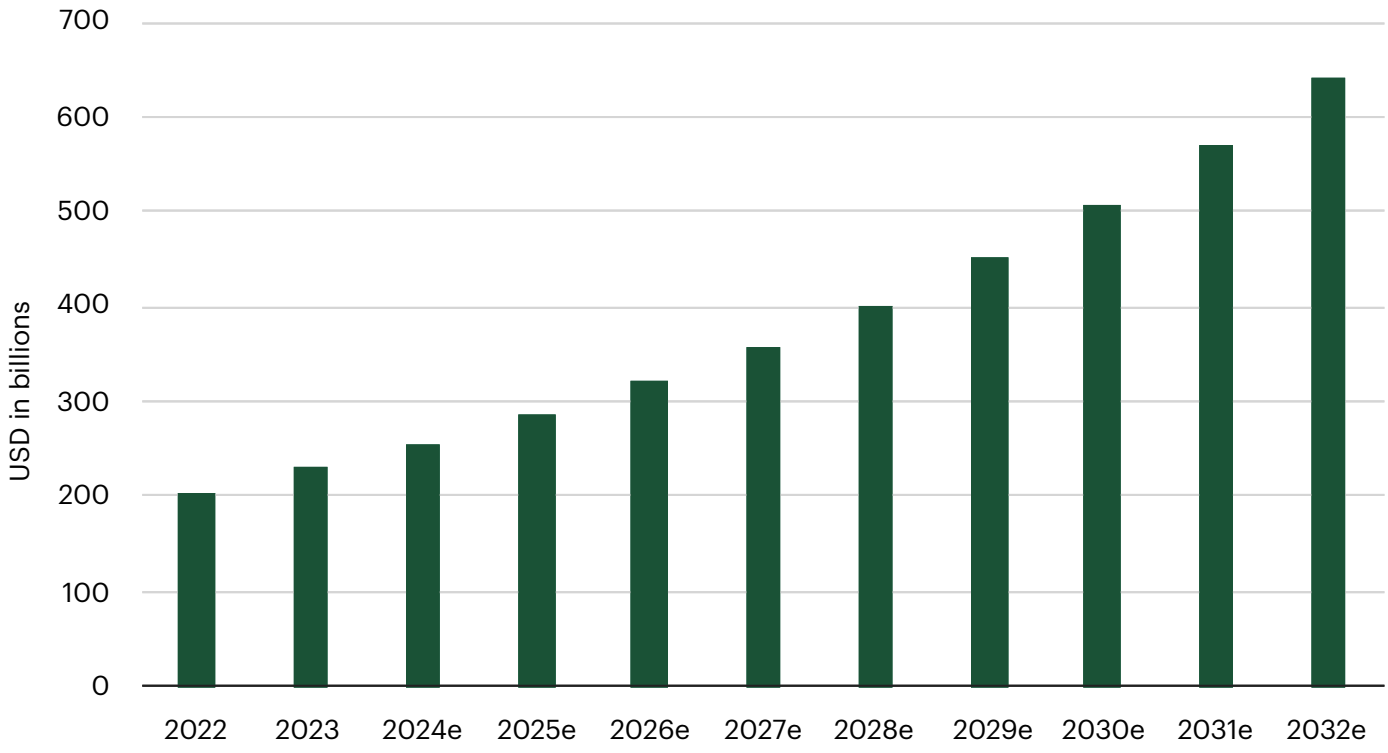

Source: Bloomberg L.P. As at May 30, 2024.

As DCs proliferate and become increasingly compute intensive, large scale cloud service providers, called hyperscalers, are all looking to put gigawatts of additional demand on the grid. To illustrate, the Electric Power Research Institute (EPRI) estimates that U.S. DCs currently consume about 100 Terawatt-hour (TWh) TWh/year with this increasing to 300 TWh/year by 2030.

Most importantly, there is an almost comically wide band, 150-510 TWh, around the EPRI's point forecast.

The massive uncertainty about future electricity demand is under-appreciated and represents a critical challenge.

Load growth exhibited a piddling CAGR of 0.4% in the decade ending 2022, but the U.S. Energy Information Administration (EIA) projects this to increase seven-fold, to 2.8%, from 2023-2030. In response, infrastructure investment is expected to boom **(Figure 3)**.

Infrastructure

Figure 3: Global electrical grid investment by region (USD, bn)

With a CAGR of 14%, the total is forecast to almost triple from 2022 to 2030.

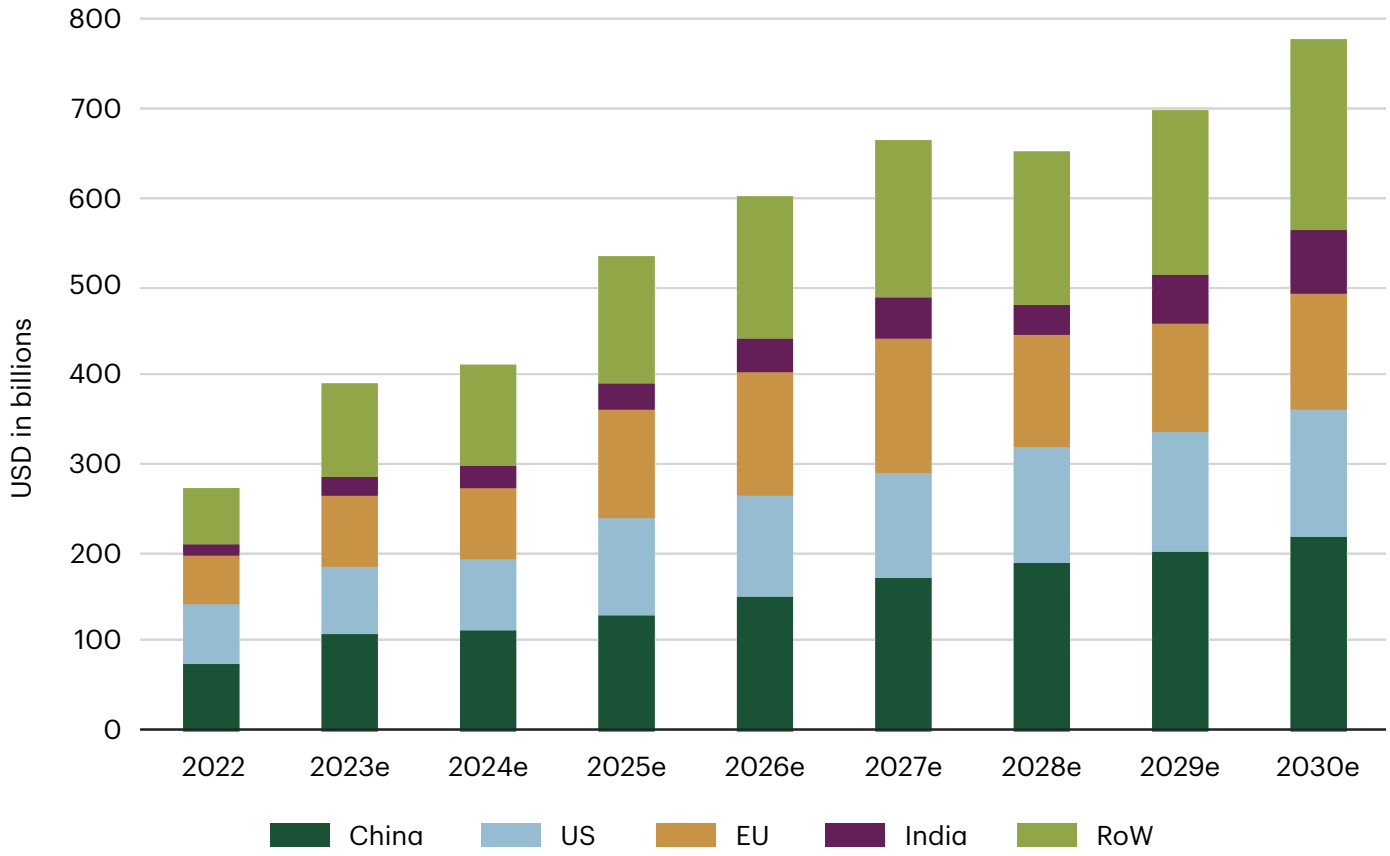

Source: Bloomberg L.P. As at March 31, 2024.

2 Where bits meet atoms: Will electricity constraints impede AI progress?

The possibility of insufficient electricity capacity offers a concrete example of where the fast-moving world of bits runs up against the slow-moving world of atoms. Think of the time it takes to add generation capacity, affix transformers, or build out the T&D grid. Moreover, utilities are by design slow-moving, regulated entities.

Reflecting this, many are worried about a looming power crunch. According to a survey by Barclays ("Grid of the future event: Highlight & survey results. February 20, 2024), 75% of respondents believe rising grid investment is a long-term secular trend. However, only a minority of those surveyed believe grid equipment and T&D capacity can keep pace.

3 Is AI unwelcome news for CO2 emissions and climate change?

Unfortunately, there are four reasons to believe AI is bad news for emissions, at least for the next decade. First, natural gas power plants are the best near-term solution to meet AI-driven demand growth, as they can be built quickly (under 24 months for a combined-cycle gas turbine plant) **(Figure 4)**.

Figure 4: The dash to gas – U.S. electricity generated by source (% share)

Natural gas and renewables will continue gaining share from coal, while nuclear's contribution has been stable for decades.

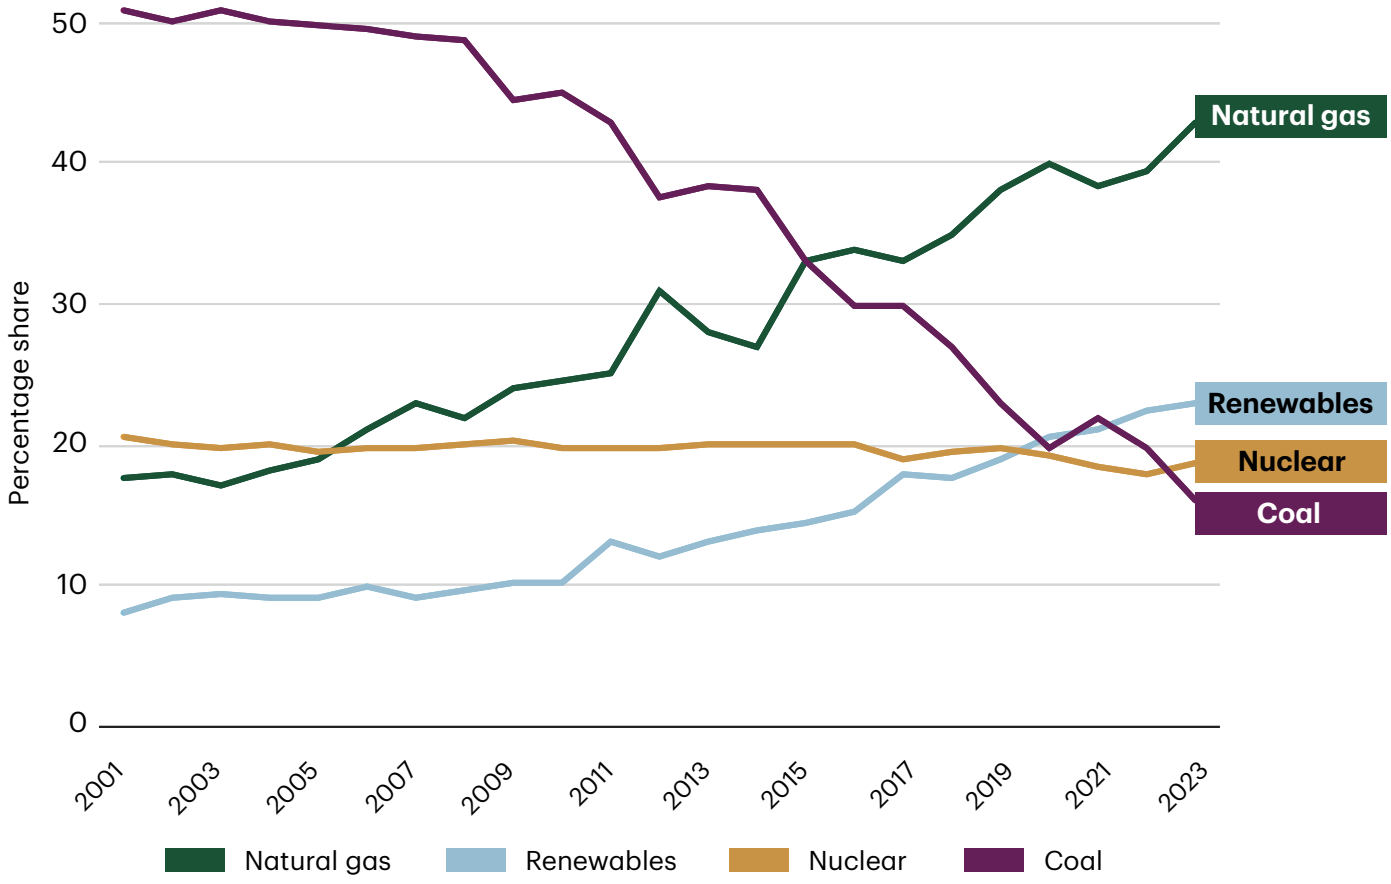

Source: Energy Information Administration (EIA). As at March, 2024.

Second, some coal plant closures might be delayed so that baseload growth can be met. Third, wind and solar are increasingly important but intermittency requires dramatic improvement in battery storage (from several hours to several days). Finally, although nuclear power is carbon free and has many fans, it has not represented a significant addition to U.S. generation capacity since the 1970s and 1980s.

4 The key risk: Enormous uncertainty about future electricity demand

Many anticipate increased grid congestion and constrained electricity supply. While we agree, we believe the key risk is actually the enormous uncertainty regarding future electricity demand.

There are critical uncertainties regarding technological progress, but a big concern for AI is how it proceeds to track cumulative costs over time. Earlier we cited the EPRI's estimate that the electricity consumption of U.S. DCs is set to increase from 100 TWh/year currently to 300 TWh/year by 2030, but with an enormously wide band, 150-510 TWh.

In light of such a wide band, what is the optimal investment path for utilities and infrastructure-related companies? There are clear risks to underinvesting, so should they just assume 510 TWh? However, there are also major costs associated with investing too much, too early, especially if this ends up being yet another case of phantom demand. Further, there is significant risk of extreme volatility in DC build and electricity demand over the next decade.

In particular, the current capital expenditure (capex) boom by hyperscalers appears unsustainable. The big three are forecast to spend $150 billion USD on capex this year with some commentators speculating that combined spending by 2030 could surpass $1 trillion USD **(Figure 5)**.[1] This capacity will likely be required eventually but it could be a decade later than some pundits expect.[2] A boom-bust cycle would create havoc for utilities and infrastructure-related companies because hyperscaler capex is a leading indicator of DC-driven electricity demand.

Figure 5: Capex by Microsoft Corporation (MSFT), Alphabet Inc. (GOOG) and Meta Platforms Inc. (META) (USD, bn)

No longer capital light – spending has increased more than fivefold since 2017.

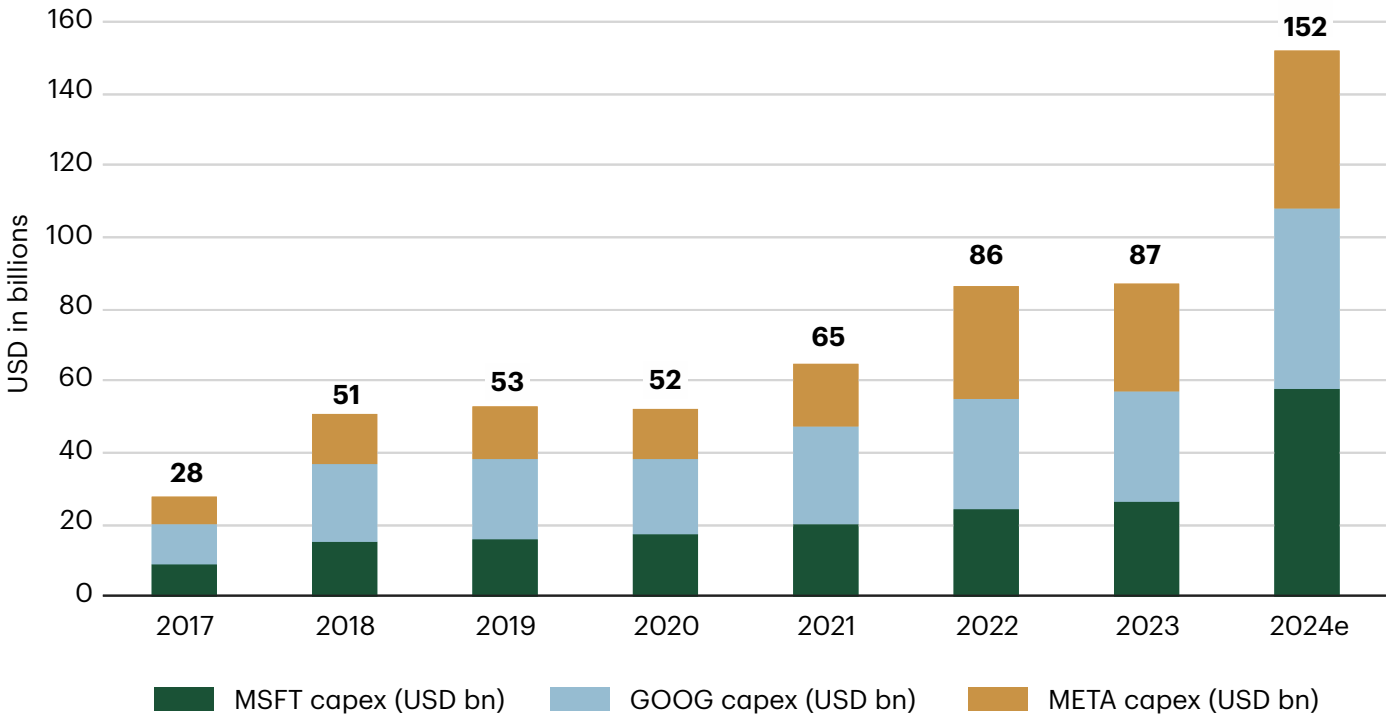

Source: Bloomberg Finance L.P. As at May 30, 2024.

[1] To illustrate the eye-popping scale of these investments, the 13-year Apollo program cost a total of $120 billion USD and the 4-year Manhattan project cost $30 billion USD (both in today's dollars).
[2] This is what occurred with the 1990s tech wave as well as earlier bouts of exuberance, such as the British railway boom in the 1840s.

There are two reasons to believe it might take longer for AI dreams to be realized. First, the diffusion of GPTs always takes decades (as it did with the steam engine, electricity, and personal computers/internet). Second, AI has not yet produced a so-called killer app, or an application so popular that the technology itself becomes indispensable to the majority of people. Right now we don't even know what that could look like and it might be many years before the big spenders earn a return on their massive capex outlays.

An important question for investors is then: How long will the hyperscalers' investment cycle last? We can never know for sure, but previous technology cycles lasted multiple years, suggesting we may be closer to 1996 than to 1999.

5 What are the implications for investors?

In this section we highlight the implications of increased electricity demand for utilities, companies exposed to electricity infrastructure, commodities, and infrastructure investments.

The increase in electricity demand has been great news for the Utilities sector. It has been the third best performing S&P 500 Index sector year-to-date, trailing only Telecommunications (includes GOOG and META) and Information Technology (includes MSFT and NVDA). Three companies from the Utilities sector delivered especially outsized returns **(Figure 6)**.

Figure 6: S&P 500 Index Top 10 companies (as a percentage, year-to-date (YTD))

It has been a technology-driven market, but three of the top ten are Utilities.

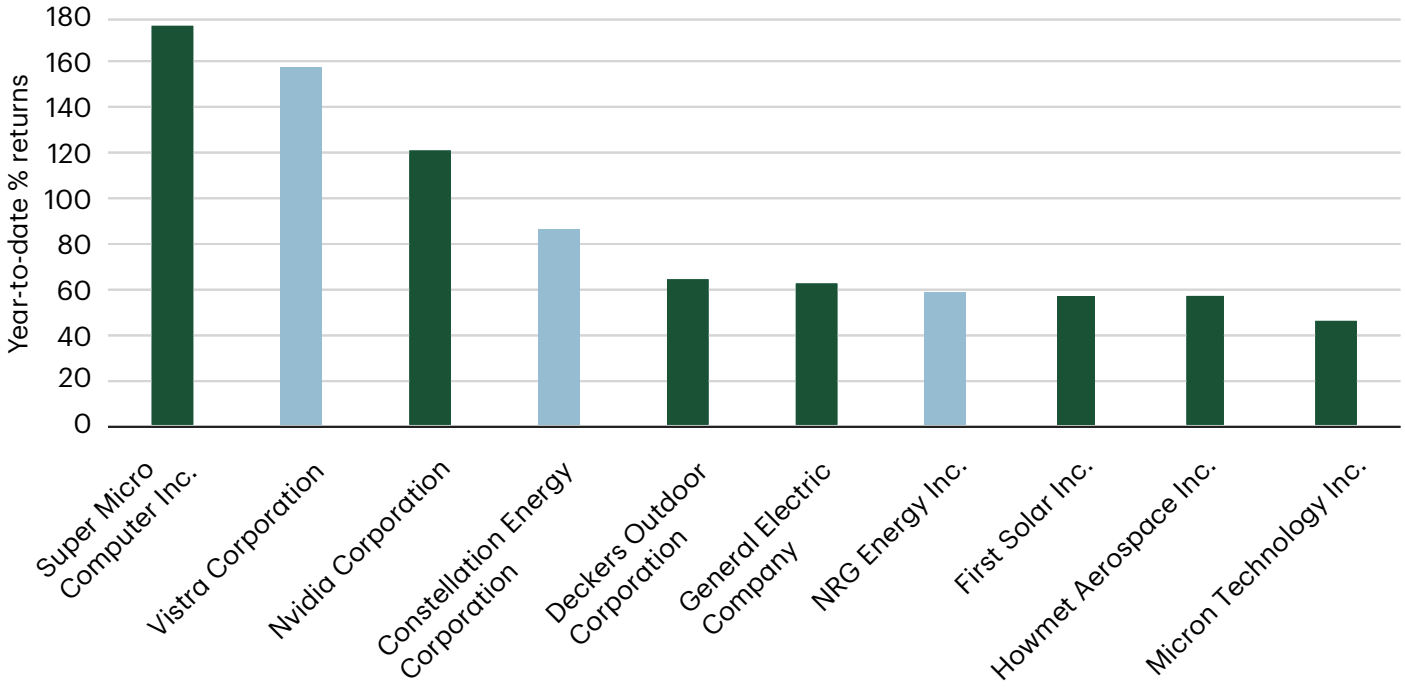

Source: Bloomberg Finance L.P., as of May 31, 2024.

Companies exposed to electricity infrastructure have also performed well. As shown in Figure 7, a hypothetical market-cap weighted index of eight companies exposed to energy infrastructure[3], illustrates how these companies outperformed the S&P 500 Index, since ChatGPT was released in late 2022. **(Figure 7)**.

[3] This is not an investable portfolio and is for illustrative purposes only. Hypothetical index comprised of: Eaton Corporation plc, Trane Technologies plc, Quanta Services Inc. Vertiv Holdings Co, nVent Electric plc, Equinix Inc., Digital Realty Trust Inc., and Amphenol Corporation.

Figure 7: An index of eight companies exposed to electricity infrastructure has outperformed since ChatGPT was released

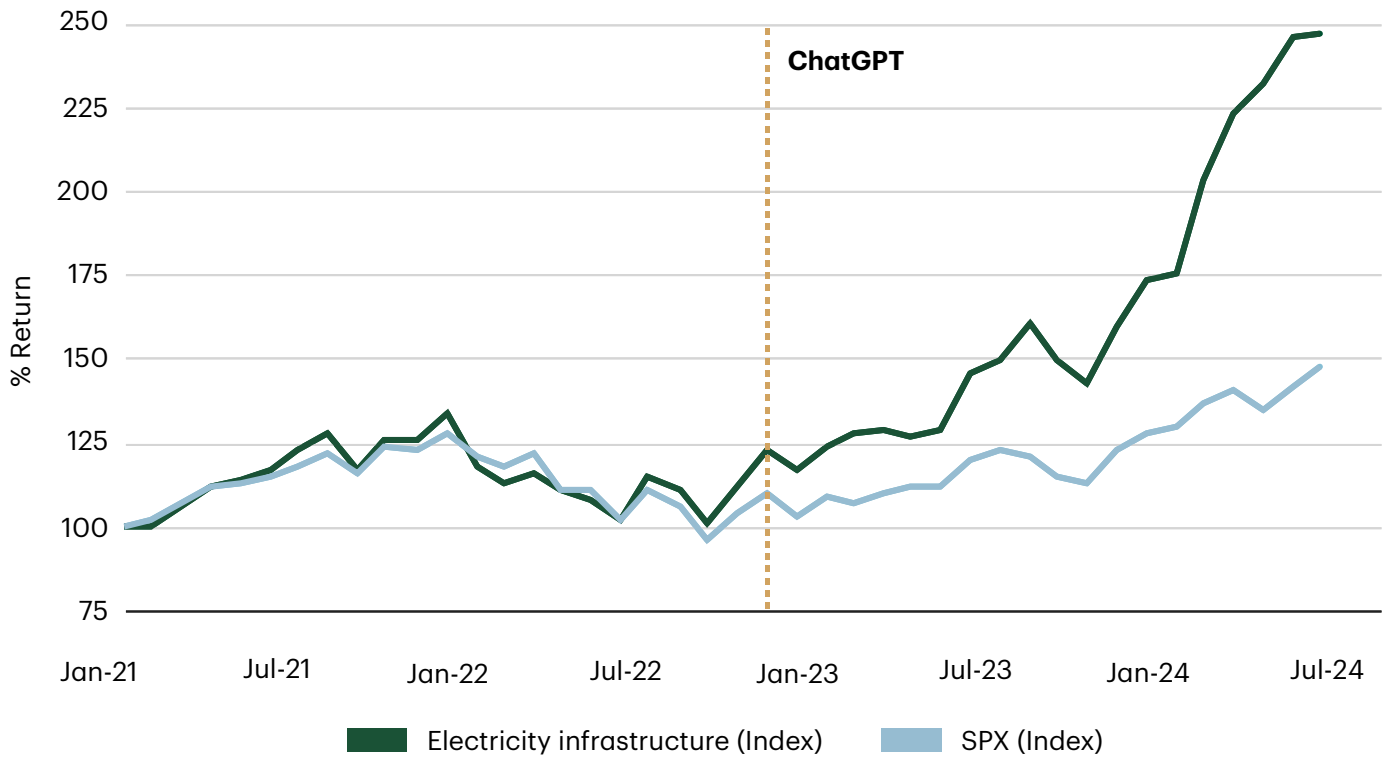

Source: TDAM, Bloomberg L.P., as at May 31, 2024.

Regarding commodities, the upgrading of the power grid, as well as the green transition, is beneficial for copper and aluminum. To illustrate, grid demand for copper is expected to exhibit a CAGR of over 12% to 2030. Another beneficiary is natural gas, which will continue to play a key role in electricity generation for at least the next decade.

Finally, the infrastructure asset class should benefit from significant investment in power. Moreover, we view the asset class as attractive because it has a relatively low correlation to equities, provides a hedge against inflation, and offers long-term, stable, risk-adjusted returns. Overall, we believe the investment case for infrastructure is very compelling.

TDAM offers a number of high-quality solutions that can provide exposure to the growth expected in electricity demand over the coming decade and beyond. This includes investments in companies within Utilities, electricity infrastructure, commodities, and Infrastructure. Contact your advisor to learn more about these opportunities and how they can potentially enhance your portfolio's performance.■

Attractive

High-quality solutions

Connect with TD Asset Management

(0724)