SYSTEMIQ

February 2025

DATA CENTRES: A GENERATIONAL OPPORTUNITY TO REWRITE THE ENERGY PLAYBOOK?

This short paper explores the opportunities and risks associated with the burgeoning demand for data centre build out, with implications for the energy sector and investors. While this paper focuses specifically on the implications of data centre expansion, Systemiq's forthcoming paper examines the flip side of the coin: how AI can accelerate the climate transition by scaling innovation and driving profound systemic change.

EXECUTIVE SUMMARY

Demand for compute intensive applications and growth in the technology sector has necessitated the construction and deployment of a new generation of data centres. Here developers and operators are prioritising speed and scale of deployment, often at the expense of their own sustainability mandates.

However, demand from data centres represents the first meaningful net new energy demand in developed economies since the 1980, and this represents a powerful opportunity for energy companies and authorities to leverage this demand stimulus and proactively reimagine and design more sustainable, scalable and resilient energy systems, and in so doing push sustainable technologies down the cost curve.

There are many opportunities across multiple geographies for investors to capitalise on the necessity of data centre build out, appealing to various risk profiles from infrastructure to frontier technology. Many of these investable themes can support sustainability mandates, including efficient chip design, water management technology and sustainable construction technology.

The powerful demand stimulus from the technology sector for new data centres presents a generational opportunity to investors and stakeholders in the energy industry to make forceful interventions to both capitalise on the underlying trend and have a meaningful impact delivering a sustainable future.

INTRODUCTION

The remarkable growth in demand for data centres, and compute more generally, since the launch of GPT-3 and ChatGPT has captured the imagination of investors, as well as stakeholders in energy, chip design and others. The emergence of artificial intelligence (AI) as a mainstream technology has compounded the already growing demand for compute ('traditional' compute – think cloud, video, etc.), and many of the world's largest technology companies, private equity investors, and others are making historic commitments towards investments to reinforce this trend.

The net result of this is that for the first time in a generation, and certainly for the first time since 'energy transition' became a common narrative, there is net new energy demand in developed countries. Indeed the technologies (AI and related innovations) that are powered by this boom in energy demand have the potential to be unusually impactful to society; akin to the way that the internet changed the world, or even the Industrial Revolution. To this end, the growth in demand for data centres represents a generational opportunity to rewrite playbooks in energy and technology.

What even is a data centre, and how and why are they built?

A data centre is a large, air conditioned room full of computer hardware, with a sufficiently large power connection to power all the computer hardware, and cabling with sufficient bandwidth to ingest and release all of the data that is stored and processed there.

Data centres first became popular as a category of real estate asset in the mid-to-late nineties coinciding with the dot-com bubble. They offered internet era companies scalability and flexibility as well as cost benefits for distributing and servicing their products. As the technology sector grew, 'the cloud' became a natural home for software, services and networks that power many digital consumer and enterprise experiences. As digital technology has become ever more pervasive in everyone's daily lives, demand for compute has skyrocketed. Until about 2020 there was an equilibrium between the increasing demand for compute and incremental improvements in energy efficiency at data centres (improvements in computer hardware, cooling, or other), though that equilibrium has now been completely shattered. Demand for compute rapidly outstrips capacity that can be made available through incremental improvements in efficiency.

AI software which is just entering the mainstream and is likely to see exponential adoption over the next decade is significantly more compute/power hungry than traditional software. In order to service this demand for compute, technology companies must deploy ever increasing numbers of data centres.

Large data centres (100MW+ capacity) are typically built over a two to three year period; a developer first purchases land and offers customers (typically a hyperscaler or other large technology firm) a long term lease (ten-plus years) for the asset. Then the developer or end user might negotiate a long term PPA to provide a stable energy connection for the data centre asset. The long term economics of the data centre asset are usually set in stone before a shovel is placed in the ground, and are often very attractive relative to other types of real estate assets.

A developer might expect a 10+% yield, which would be highly stable and predictable, and would outperform other classes of commercial real estate. A good rule of thumb is that each 100MW of data centre capacity requires $1bn in investment (excluding the cost of computer hardware). The risk that the data centre developer effectively underwrites is the credit-worthiness of their customers (the hyperscalers and the like), rather than market risk for cloud compute or AI or similar.

In early 2025, there are approximately 1000 data centres worldwide that have 100MW of capacity or greater and that number is projected to grow by 100-200 every year until at least the 2030s. There have been several notable announcements of funding for these initiatives; some examples include 'Stargate', the $500bn venture partnership of OpenAI, Softbank and Oracle; Microsoft's $80bn commitment towards data centre development; and the $30bn Global AI Infrastructure Investment Partnership, headed by Blackrock and Microsoft. The energy sector has an important role to play in powering these ventures, and authorities of the jurisdictions that they are built in are responsible for the stewardship and safeguarding of these mega projects.

Incentives for data centre development and sustainability

Naturally, since data centres consume a lot of power and resources, there have been various attempts to integrate sustainability principles into their development. In some cases, there is strong alignment of incentives and sustainability can strengthen the market appeal of a data centre asset, for instance, where a data centre can be powered by incredibly cheap geothermal or hydropower. Sustainability extends beyond energy consumption; water use for cooling, embedded emissions from construction materials, and electronic waste management are also critical factors.

In many cases, however, it is apparent that sustainability concerns are somewhat orthogonal to industry and market incentives. The best evidence for this is through 2023-2024 when Microsoft, Google and other technology companies opted to sidestep their sustainability programmes to build out more data centre capacity predominantly for AI and in so doing dramatically increased their emissions, by 30% in the case of Microsoft. It is important to note that capital markets strongly rewarded this decision; Microsoft's capitalisation doubled over the same period.

The key motivator for rapid data centre build out is an economic imperative for growth; the importance of this cannot be underestimated both at the company level and a systemic level. At a company level, this is easy to grasp since most leadership are compensated in a way that correlates with the financial performance of their company. At a systemic level, it is important to note that 40% of the total capitalisation of the S&P 500 is tied up in a handful of technology stocks, meaning that the value of millions of pensions and other investments are inextricably linked to the continued performance and growth of the technology sector. The system wide economic consequences of a derailment of growth in the technology sector would cause profound hardship and distress in developed economies, akin to or even worse, than the 2008 financial crisis. This systemic importance has meant that it has become necessary for governments to aggressively support domestic technology build-out, the UK AI Action Plan is a good example of this, as is the $500bn Stargate project quickly brokered by the incoming Trump administration.

This is not to say that sustainability is unimportant, and authorities might be well minded to find ways to incentivise sustainable behaviours and initiatives, but it is important to properly contextualise the importance of sustainability and the real trade-offs in play. Data centres locations can be somewhat fungible, and developers can choose to build in "optimal locations" where readily available energy supply, cooler climate, and favourable regulation converge. When paired with the powerful imperative for economic growth, it becomes very important for authorities to find solutions that are globally competitive without resorting to heavy handed sustainability legislative approaches that could cause economic stagnation locally in the case that developers prioritise regions that they find are cheaper or less burdensome to construct and deploy data centre assets.

HOW THE DATA CENTRE SECTOR INTERACTS WITH THE POWER SECTOR

It is well known that data centres have huge demands on power systems. Data centre operators purchase long-term offtake for power from local utilities and pass the cost of power through to their customers in most cases. Given the strong coupling between power prices and the long-term economic viability of a data centre asset, it may be surprising to consider that for the most part **data centre developers and operators treat power companies as vendors and not as more involved/important partners.** This is a pragmatic choice; once power is delivered onsite, there is little additional capability that a power company can provide to an established data centre operator. Complicated structures/joint ventures etc. are generally unattractive to data centre developers and operators since they optimise for a reliable and predictable financial yield and are unwilling to introduce any uncertainty.

The data centre operator typically does not choose how the power they use is generated, and whether or not it is sustainable depends on the energy generation stack of the utility that they purchase power from. It is not uncommon for data centre operators to purchase renewable energy certificates, guarantees of origin, or emissions trading instruments to make their power procurement more sustainable (superficially or otherwise). Some stakeholders, in particular the hyperscalers, might have very strong corporate sustainability mandates and there are examples where they have proactively sought green generation. Since there is some degree of flexibility in determining where a data centre might be located to serve a given market, it is normal for developers to shop for the best deals on land, power, interconnects etc.

There is some evidence that certain data centre operators/hyperscalers are looking to move upstream and develop or manage their own energy generation assets, or integrate more tightly with energy companies, for two key reasons:

1. **Economics**; if a data centre operator can procure energy via a private wire and bypass the grid, they can avoid paying network charges and statutory fees associated with purchasing power on grid. This possibility is only viable in certain jurisdictions where these assets can be permitted.
2. **Pursuit of speed and scale**: here there are examples of extremely ambitious projects, such as Microsoft's proposal to reopen the Three Mile Island nuclear plant. The rationale for attempting to reopen this generation asset would be the possibility of bringing 1GW+ of new

generation capacity to market extremely quickly relative to market norms for similar levels of new power demand, especially since most interconnect infrastructure already exists.

However, this emerging trend of large assets vertically integrating upstream if not managed carefully could be damaging to efforts to decarbonise energy generation and build fair electricity systems that benefit populations broadly, since vertical integration is for the most part motivated by the commercial requirements of the developer rather than any desire to improve systems that broader society benefits from.

Speed, scale and reliability are perhaps the most important priorities for technology companies as they determine their investment priorities as we enter this new technological epoch. Given this scale of demand, there should be a clear expectation for data centre operators to contribute to new generation capacity in a sustainable way. While this is sometimes treated as a secondary concern as authorities recognise the need to their geographies to remain competitive as a hub for technology investment, **data centre operators could actively push for renewable energy supply given decreasing premiums for reliable green energy and, for some, their eagerness to stick to their decarbonisation targets.**

In contrast to this, there is a substantial opportunity for energy companies to innovate and introduce new technologies to market as a result of significant new demand from the data centre sector. **For the first time since 'energy transition' has been a mainstream narrative, there is net new energy demand at a system level. This could be a powerful catalyst for transforming the energy generation stack, pushing technologies down the cost curve and expediting permitting for new energy mega projects.** Indeed, there are projections that overall power demand will increase by 3%/year in the USA driven solely by demand from the technology sector – demand growth the like of which has not been seen since the 1980s.

Indeed, if electricity systems embraced the opportunity to modernise and become able to deliver cheap, sustainable, reliable power at scale, then there would be no reason for data centre developers to look to vertically integrate, and societies and industrial economies would benefit more broadly. That prerogative falls on energy companies and regional authorities more so than it does on technology companies and data centres.

What are energy companies doing?

Energy companies and legislators are generally taking pragmatic steps to manage new demand from data centres, looking for solutions that maximise utilisation of existing asset bases rather than developing new generation capacity at scale to minimise expense and risk. Indeed, any costs incurred supporting large scale projects that do not become viable are ultimately borne by retail customers.

Individual data centres are typically accommodated in a piecemeal fashion, one by one, by many grid networks with modest interventions that ensure that energy systems are capable of managing peak demand rather than taking a longer term, more aggressive view of demand. This approach of making modest interventions to accommodate individual assets limits risk and can push some technologies down the cost curve, but does not necessarily embrace the opportunity to make significant enhancements to electricity systems in the short term.

Several grid systems have received enquiries about supporting new data centre load that could more than double the base load of their network over the next decade, for instance AEP Ohio. To service this demand, significant investments in transmission infrastructure as well as generation capacity are necessary. There is some misalignment over who should bear responsibility for paying for new grid upgrades and how much of that cost is reasonable to transfer to data centre developers or technology companies.

As demand for compute continues to increase over the next decade, the overall base load requirement will become far greater than it is now in many electricity systems, and, during peak times the pressure on grid capacity will be far greater than it is now.

In conjunction with regulators, it might be wise for power companies to assess whether now is the opportune time to make more forceful interventions to future proof energy systems and ensure regional competitiveness – not just for data centres/compute/digital economy, but also more broadly – cheap power is well correlated with many economic and development indices. Where possible, energy companies should find contractual guarantees to underwrite the risks associated with infrastructure development and embrace the opportunity to significantly modernise both generation and transmission infrastructure.

To date it appears that most of the new projected demand caused by data centres is likely to be serviced by gas generation (there is a pipeline of over 50 GW of new gas powered generation assets planned in the United States) in parallel with strong investor enthusiasm for nuclear technologies including fusion, SMRs and conventional fission.

However, with strong foresight and planning, plentiful renewable and energy storage infrastructure could be deployed – at significant enough scale to push these technologies further down the cost curve. In this sense**, there is a generational opportunity for energy companies to reimagine and**

rebuild their stack, stimulated by net new, robust commercial demand as opposed to regulatory mandates or similar.

It is important for stakeholders involved in power to find ways to ensure that they can service demand and remain competitive, including **taking a medium to long term view of demand, supported by meaningful contractual commitments where appropriate**, to modernise energy systems at scale rather than seek piecemeal solutions to accommodate individual assets. Energy companies should embrace technological innovation, in particular **technologies that could make energy generation cheaper and more resilient**; nuclear technologies have received particular emphasis and support from stakeholders involved in data centres.

Finding solutions that expedite processes including permitting, construction and operation would greatly appeal to data centre developers; expedited processing could be complemented by some sustainability lever. Authorities should look to design incentive structures where appropriate, perhaps similar to the IRA program, to encourage stakeholders to make sustainable choices rather than regulating opportunity away that make their regions uncompetitive.

INVESTING IN SUSTAINABLE DATA CENTRES

The primary driver of growth in data centres and energy demand is the rapid expansion of the technology sector, particularly in compute-intensive applications. Investors can gain direct exposure to this trend through blue-chip technology equities, including hyperscalers, application companies, and chip manufacturers. Of course, second-order effects of the growth of the technology sector, in particular the booming growth in data centre development, represent the main subject matter of this paper and create additional investable opportunities with differing risk profiles and upside.

Infrastructure and energy investment opportunities

Data centre developers can capitalise on pent up demand from the technology by financing new projects with strong backing from the private sector stakeholders that have traditionally supported the sector (technology companies, private equity, data centre REITs etc.). Additionally, investors can benefit from a rather unique moment in which **legislators are overtly supportive** of data centre development (indeed the opposite was true not long ago - for example, Singapore imposed a moratorium on data centre development from 2019-2022).

Most major technology companies prioritize speed and scale as they expand their data centre footprint. This presents opportunities for investing in technologies and approaches that can **speed up data centre development** and/or reduce capital expenditures, and in so doing **create more**

sustainable, more standardised data centre designs. A strong recent example of a major facility that was launched with an expedited time frame was xAI and NVIDIA's three-week data centre deployment (instead of ~18 months) in Tennessee leveraging expertise from Tesla in supply chain and modular prefabrication. By optimising the construction of data centre buildings and the deployment of computer hardware, it will be possible to deliver assets which are less wasteful (in terms of time, capex and materials/emissions) and more resilient since standardised designs could be more readily refitted/modernised/relocated when appropriate.

It is also helpful to consider the **application of technologies that can reduce the embedded carbon of new assets,** such as low carbon cement, and authorities may wish to find ways to incentivise the use of low carbon construction techniques; in some cases there is a negligible cost difference between standard construction techniques and lower carbon alternatives. There may be opportunities to invest in the development of low carbon construction materials, but also in the service companies that can deliver low carbon construction projects at scale.

Energy generation investments remain critical to sustainable data centre expansion. As highlighted in the previous sections, investors face a fundamental choice: they can opt to deploy the cheapest and most readily available power option – often natural gas (although CCG-turbines also have extended delivery times) - or **invest in renewables plus storage** (supplemented with gas peakers), which, while currently slightly more expensive, offer a cleaner and more sustainable alternative. Nuclear technologies, particularly small modular reactors (SMRs), have been particularly hot with investors that have keen interest in the data centre segment. Breakthroughs in the underlying technology rather than deployment, such as perovskite-based solar cells, and next-generation battery technologies, like sodium-ion, represent interesting opportunities in the data centre's upstream value chain that promise to reshape the cost and performance dynamics of the wider (renewable) energy system. Investors and energy companies can design **integrated 'campuses'** where assets are bundled to create the perfect site for a data centre development which could then be sold for both a quicker and greater return than traditional energy assets; a good example is Talen's ongoing $650m deal with AWS in Pennsylvania. **This campus approach could be an effective way to bundle renewables and match them to demand.**

There are inconsistencies between geographies in terms of permitting, power markets and other important drivers, so opportunities to deliver solutions that drive sustainability may not necessarily be equally well received in every market (Europe vs America vs Middle East vs Asia, with added local nuance in each region). However, if authorities and energy companies proactively design sustainable systems it is possible to create environments where sustainable decision making is rewarded.

Operational and efficiency investments

Energy efficiency improvements within data centres present another avenue for investors. The most significant efficiency gains come from improved computer hardware, including next-generation chips designed for AI and high-performance computing. In most cases this is driven by technological development cycles of leading chip manufacturers, though there are several notable examples of companies designing and building custom silicon for tailored applications which may generate remarkable energy efficiency and performance improvements and therefore be highly investable.

Beyond compute efficiency, significant energy savings can be achieved through **optimized cooling**, networking, and auxiliary systems. The ownership structure of data centres plays a key role in efficiency incentives; for colocation facilities there is an owner/tenant split-incentive where there is no direct incentive for either the colocation firm or the customer to take a hard core approach to maximising efficiency. However, for owner operated data centres (e.g. hyperscalers) there is a stronger incentive to maximise energy efficiency and to invest into new (liquid or solid-state) cooling technologies, optimized **building design**, and initiatives to reduce water-use and recycling electronics.

Some technologies that are leveraged to improve energy efficiency include **waste heat recovery** and custom energy management software; there is relatively little scope for new asset light innovations to improve the efficiency of hyperscale data centres, but there is some scope for service providers that implement efficiency related technologies to scale. **Workloads deployment** can be optimised and dynamically shifted (time, region) in response to grid signals to maximise energy efficiency and minimize carbon footprints given a sufficient scale of distributed compute capacity.

Frontier and emerging opportunities

Some investors may look to more unconventional opportunities. For instance, Star Cloud (formerly known as Lumen Orbit) believe that data centres should be **deployed from space** rather than terrestrially, and necessarily their approach is solar powered. Though speculative, such concepts could redefine sustainable data centre operations in the long term.

A generational investment opportunity

Investors can engage with data centre-driven opportunities across a spectrum of risk profiles, geographies, and business models. Options range from infrastructure investment (data centres, power generation) to growth plays in enabling technologies like efficiency software, chip design and cooling systems, as well as higher-risk early-stage ventures. Many of these investments can align well with sustainability mandates and serve to both generate returns and progress various sustainable development goals. Given the robust demand driven by compute-intensive applications and broader technology sector growth, this moment represents a generational opportunity for investors to capitalize on the data centre boom while supporting the sustainable evolution of the industry.

CONCLUSION

There is significant demand for new data centre assets underpinned by growth in the technology sector and demand for compute intensive applications. The developers and operators of these assets are primarily motivated to maximise the speed and scale at which they can deploy these assets, and often do not prioritise sustainability in their decision making.

However, new data centre assets represent the first meaningful net new energy demand in developed economies since the 1980s, and indeed the first meaningful net new energy demand since the 'energy transition' became a mainstream narrative. This presents a powerful opportunity and catalyst to reimagine energy generation systems and enable stakeholders to proactively design and deploy more sustainable, scalable and resilient energy systems. This would require energy companies and authorities to take the initiative and think carefully about the big picture and end game, rather than integrating individual data centre assets to grid networks in a piecemeal fashion.

Investors can find alpha in a varied set of opportunities from infrastructure development to growth equity to frontier technology. Many of these opportunities can help fulfil sustainability mandates including technologies that promote sustainable building practices, water management and cleaner energy generation.

There is a certain inevitability to the growth of the technology sector and the AI boom that reinforces the robust demand for new data centres. There will be commentators that will provide modern day analogues of McKinsey in the 1980s predicting that the market for mobile phones would only be 900,000 users. Indeed there were commentators in the early innings of the internet era that drastically underestimated its size and impact. And so there will be at the dawn of this era defined by AI, compute and data centres. Investors would be wise to recognise the underlying mega trend and make the most of necessary data centre build out, which represents a generational opportunity to both capitalise on an important trend but also reimagine systems including but not limited to energy for the betterment of the world.