AI, Data Centres, and the Problem We Keep Avoiding

There is a growing sense that artificial intelligence is about to blow a hole in our energy and water systems. Headlines warn of runaway electricity demand, thirsty data centres, and infrastructure pushed beyond its limits. At the same time, others argue the opposite: that rapid efficiency gains will make these concerns largely irrelevant.

Both narratives are tempting. Neither is particularly helpful.

The real issue is not whether AI will be the catalyst of a global environmental catastrophe or whether technology will magically solve everything. The real issue is that we are making long-lived infrastructure decisions with surprisingly little certainty about the long-term impact.

That gap between ambition and evidence is where the risk sits.

Why AI feels different this time

Data centres are not new, the UK has been building them for decades, and for most of that time they grew quietly in the background. Improvements in server efficiency, cooling, and software meant that more computing could be delivered without visibly pushing up national electricity demand.

However, AI changes that dynamic considerably.

Unlike traditional digital workloads, AI is highly energy-intensive and unusually concentrated. Large training runs and inference clusters draw enormous amounts of power in relatively small physical spaces. A single AI-focused facility can place demands on local grids that would once have been associated with heavy industry.

National Grid modelling reflects this shift. Under some scenarios, UK data centre electricity demand could rise several-fold by 2030. Government analysis suggests installed capacity may need to more than double in the same timeframe, and internationally, the International Energy Agency has warned that AI could become one of the dominant drivers of data centre energy use by the end of the decade.

That sounds alarming, and in isolation it probably should. However, as with everything AI related, there things change fast and there’s more to consider when you dig deeper.

Forecasts built on assumptions, not certainties

If you look closely at the forecasts driving concern, a common theme emerges. They are highly sensitive to assumptions.

- How quickly will AI be adopted across the economy?
- Will ever-larger models remain economically attractive?
- How fast will hardware and software efficiency improve?
- Where will new facilities actually be built, and can they connect to the grid?

Small changes to any of these inputs produce wildly different outcomes. That is not a flaw in the modelling, but a reflection of how early we still are in understanding AI's real-world deployment.

The International Energy Agency has been explicit about this uncertainty with estimates of future data centre energy use varying by orders of magnitude depending on methodology. Even current consumption is difficult to measure consistently, because data centres are poorly classified in national statistics and reporting varies considerably.

This matters, because confidence often creeps in where evidence is thin or cherry-picked.

Water adds another layer of confusion

Water use is often discussed as if it were a single, easily measured number. In reality, it is one of the least transparent aspects of digital infrastructure.

Recent UK survey data suggests that many data centres use little or no water for cooling, relying instead on air-based or closed-loop systems. Where water is used, reported volumes are often comparable to other large commercial buildings.

That challenges the idea that data centres are inherently water-hungry.

But again, context matters. The survey data is voluntary and may not fully capture the largest facilities. Direct water use also ignores indirect impacts, such as water consumed in electricity generation or semiconductor manufacturing.

Most importantly, water availability is a regional issue. A facility that is unproblematic in one location can be deeply contentious in another. For example, in parts of eastern England, including Suffolk, water stress is already a planning constraint. There's a danger that if we treat data centre water use as a national consideration, it could obscure very real local limits.

Infrastructure decisions are being made anyway

Despite all this uncertainty, decisions are being made at pace.

Grid upgrades, planning approvals, and investment commitments lock in patterns of development for decades. Once built, data centres are not easy to move. Their environmental footprint becomes a fixed feature of the landscape for a long time.

This is not an argument for stopping AI or data centre development. Digital infrastructure is essential to economic resilience, public services, and innovation. The question is whether we are being honest about what we know, and just as importantly, what we do not.

Right now, ambition seems to be running far ahead of measurement.

What better decision-making would look like

The solution here is not perfect forecasting - that is simply unrealistic at this point.

What is realistic is better governance and transparency.

Mandatory reporting of energy and water use would dramatically improve transparency. Integrated planning between electricity, water, and digital infrastructure would reduce the risk of regional bottlenecks. Clearer classification of data centres in national statistics would allow policymakers and the public to track impacts over time rather than rely on one-off studies.

Crucially, uncertainty should be acknowledged explicitly in policy decisions, not smoothed over. Scenario planning, rather than single headline numbers, is far better suited to a technology that is still evolving this quickly.