A Hexalemma Framework for AI Datacentres and Grid Stress

Rebecca Le Jeune

Table of Contents

Abbreviation	
AI	Artificial Intelligence
CLD	Causal Loop Diagram
CNI	Critical National Infrastructure
DESNZ	Department for Energy Security and Net Zero
DSIT	Department for Science, Innovation and Technology
EAC	Environmental Audit Committee
EESI	Environmental and Energy Study Institute
ESG	Environmental, Social and Governance
ESO	Electricity System Operator (National Grid ESO, now NESO)
GDP	Gross Domestic Product
GPU	Graphics Processing Unit
GW	Gigawatt
IDCA	International Data Center Authority
IEA	International Energy Agency
IT	Information Technology
LILEE	Low Income Low Energy Efficiency (fuel poverty metric)
MOSL	Market Operator Services Limited (the UK water retail market operator)
ESO	National Energy System Operator
NPPF	National Planning Policy Framework
NPS	National Policy Statement
NSIP	Nationally Significant Infrastructure Project
PPA	Power Purchase Agreement
PUE	Power Usage Effectiveness
REC	Renewable Energy Certificate
RIIO	Revenue = Incentives + Innovation + Outputs (Ofgem's network price-control framework)
RSE	Reliable, Safe, Ethical (the soft value set of the hexalemma)
TNUoS	Transmission Network Use of System (charge)
TPU	Tensor Processing Unit
TWh	Terawatt-hour
UKCEH	UK Centre for Ecology and Hydrology
UPS	Uninterruptible Power Supply
VPPA	Virtual Power Purchase Agreement
WUE	Water Usage Effectiveness

Introduction

The hexalemma comes from an international research project that looked at inefficiencies in the Food/Energy/Water nexus in cities including Bristol. The starting point was a common systems management dilemma, resilience versus efficiency. Systems built for resilience, with lots of backups, cost more. Systems built for efficiency are cheaper but more fragile. The hexalemma shows that this either-or choice is just one of many tensions that systems face. Looking more closely at each dilemma reveals a deeper pattern. Two factors in tension can be either opposed or interdependent. This is called a "tension" because it involves both conflict and cooperation.

When a third factor is added to resilience and efficiency, it creates a trilemma: Resilient, Sustainable, and Efficient. The hexalemma expands on this by combining this outward-facing 'hard' trilemma with an inward-facing 'soft' trilemma focused on people's experience: Reliable, Safe, and Ethical. Together, these form two connected RSE groups.

1. Hard RSE (system-outward): Resilient, Sustainable, Efficient
2. Soft RSE (human-inward): Reliable, Safe, Ethical

These values fit into three main categories. Quality, which includes integrity and trust, pairs with Resilience which asks if we can rely on the system to last and withstand stress. Care, meaning non-harmful and healthy, pairs with Sustainable and Safe, does the system and its people avoid causing harm, and are they protected? Allocation, which covers scarcity and justice, pairs Efficient with Ethical, are scarce resources distributed by market forces or by fairness and justice?

With six values, there are fifteen possible pairs of tensions. Although all six are important, improving one often puts pressure on the others. These trade-offs are built into the system, not just caused by poor management or lack of political will.

The reason this matters for AI datacentre governance is that one of these six values, efficiency, has a property the other five do not. Efficiency gains tend to feed the very demand they were meant to contain. The Jevons paradox, first set out by William Stanley Jevons in 1865 and revived for ecological economics by Alcott (2005) is useful to us here. Jevons watched the steam engine become steadily more coal-efficient across the eighteenth and nineteenth centuries and noticed that total coal consumption did not fall but rose sixteenfold, far faster than population (Alcott, 2005, p. 17). His conclusion was that "it is wholly a confusion of ideas to suppose that the economical use of fuel is equivalent to a diminished consumption" (Jevons, 1865, cited in Alcott, 2005, p. 12). Greater efficiency lowers the cost of using a resource, which raises its profitability, which draws in new capital and opens up new uses, so absolute consumption climbs rather than falls (Alcott, 2005, p. 13-14).

The mechanism is not a quirk of coal. Alcott argues it holds wherever there is latent demand and a positive price response, which describes AI compute almost exactly (Alcott, 2005, p. 16). When the rebound is large enough that consumption ends up higher than before the efficiency improvement, the literature calls it "backfire" (Alcott, 2005, p. 11). This is the strong version of the argument, and it is the one that should worry anyone designing datacentre policy around Power Usage Effectiveness (PUE) thresholds and demand response. It reframes the central problem, meaning that the risk is not that the system is insufficiently efficient. The risk is that efficiency is the one value the sector knows how to optimise, and optimising it accelerates the growth that erodes the other five.

The fifteen tensions and the problem of maximisation

The hexalemma is useful because it looks at all fifteen tensions together. Usually, decision-makers focus on just a few values and ignore the rest. As Meadows (2008) points out, making systems more efficient often makes them less resilient. Over the last century, systems have moved from being resilient to being more efficient. This isn't anyone's personal fault; it's a built-in bias. Focusing on economic efficiency has reduced diversity, backup options, and adaptability.

In this paper, the mechanism works as follows. AI datacentre governance has focused on things like PUE[1] optimization, better liquid-cooling efficiency, corporate renewable PPAs[2], and

[1] Power Usage Effectiveness: The ratio of total facility energy to IT equipment energy. A PUE of 1.0 means all energy reaches the computing load; a PUE of 2.0 means half is lost to cooling, lighting, and other overheads. The industry standard metric for datacentre energy efficiency.

[2] PPA (Power Purchase Agreement): A long-term contract between an energy buyer (such as a datacenter operator) and a renewable energy generator.

marginal demand response[3]. All these actions move toward the efficient side of the hexalemma. The other five values are either not measured (like grid resilience to concentrated AI demand), hidden (the Renewable Fiction Loop[4] scores low on sustainability and ethics), or shifted to others (grid upgrade costs are spread across all consumers, which scores low on ethics).

The Time Dimension: Shocks, Stresses, and Tipping Points

Each of the six hexalemma values changes over time. The hexalemma working notes show this as a matrix whereby resilience slowly weakens under stress, then fails suddenly; sustainability declines gradually, then collapses; safety gets worse bit by bit before a disaster; efficiency can improve step by step or change quickly in a crisis; equity changes slowly, but sometimes shifts suddenly through major events.

The phrase from Hemingway's *The Sun Also Rises* (1926), "gradually, then suddenly," sums up the dynamic discussed here. Grid stress from AI datacentre growth is building up slowly. Each connection request, PPA, or planning approval seems manageable on its own. But feedback loops like the Demand Spiral[5], Renewable Fiction Loop, and Resilience Erosion Loop[6] are quietly increasing pressure.

---

[3] Demand response: The practice of adjusting electricity consumption in response to grid conditions, typically reducing or shifting load during periods of high demand or low supply. In the datacentre context, demand response is limited because most AI workloads are treated as requiring continuous, uninterruptible power.

[4] Renewable Fiction Loop: One of the three causal feedback loops developed in this paper. It describes the dynamic in which corporate renewable energy accounting (via PPAs and RECs) allows reported emissions to fall while actual grid carbon intensity remains unchanged or rises, because the procurement reshuffles credit for existing renewable generation rather than adding new capacity.

[5] Demand Spiral: One of the three causal feedback loops developed in this paper. It describes the dynamic in which efficiency improvements reduce the cost per unit of AI compute, which stimulates demand growth, which increases total energy consumption, which drives further grid stress, which prompts further efficiency investment, completing the loop

[6] Resilience Erosion Loop: One of the three causal feedback loops developed in this paper. It describes the dynamic in which geographic concentration of AI datacentre demand creates regional grid stress, which requires infrastructure investment, which is not funded adequately under short-horizon governance, leading to declining grid resilience over time.

The hexalemma's time dimension explains why this is risky as systems often fail not because of one big mistake, but because many small changes add up until stress turns into a crisis. This is the tipping point. Furthermore, this maps onto the Koutsandreas and Keppo (2026) finding about temporal inconsistency in governance, short-horizon decision-makers consistently under-invest in structural resilience, managing accumulating stress through incremental efficiency measures until the system reaches a structural threshold it cannot recover from without far costlier intervention.

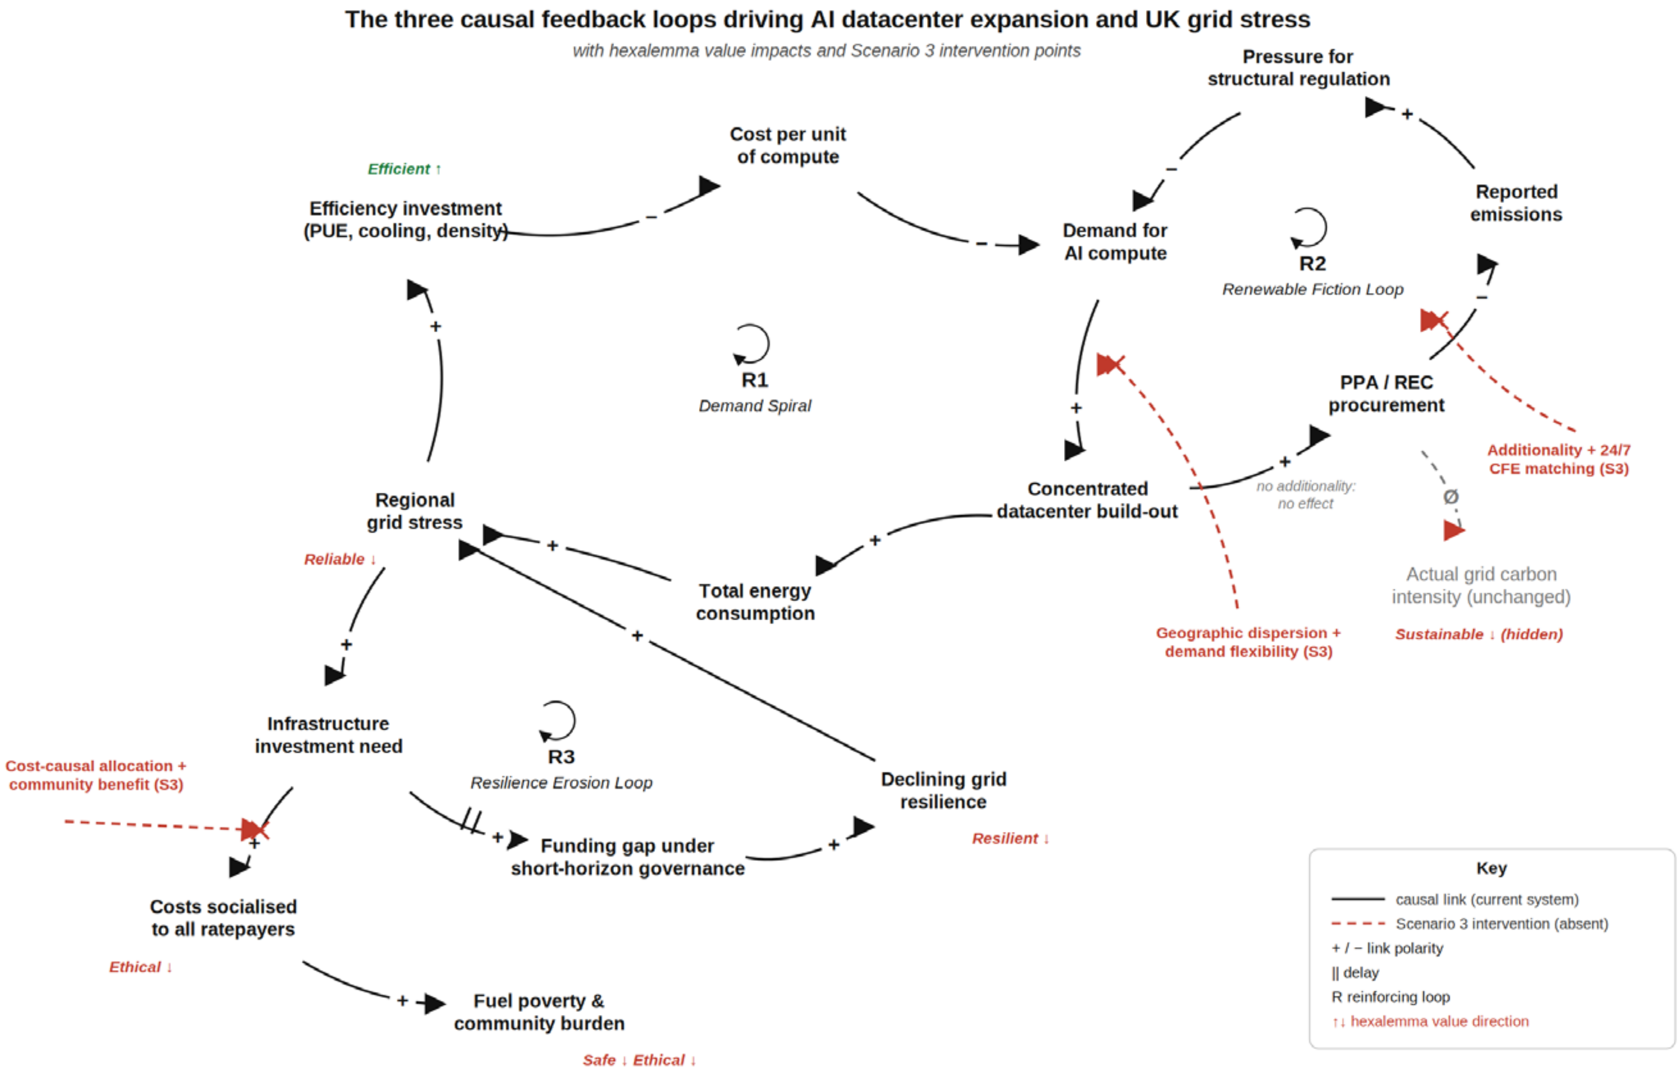

Figure 1: The Demand Spiral, Renewable Fiction Loop, and Resilience Erosion Loop as an integrated system

Figure 1 maps the three causal loops together. Each loop is reinforcing, meaning that movement around the loop strengthens the conditions that started it. The Demand Spiral (R1) shows how efficiency investment lowers the cost per unit of compute, which stimulates demand, drives concentrated datacentre build-out, raises total energy consumption and regional grid stress, and prompts further efficiency investment.

The Renewable Fiction Loop (R2) shows how Power Purchase Agreement (PPA) and the Renewable Energy Certificate (REC) procurement allow reported emissions to fall, which weakens the pressure for structural regulation and removes the constraint on demand growth, while leaving actual grid carbon intensity unchanged because no additionality is required.

The Resilience Erosion Loop (R3) shows how grid stress generates infrastructure investment need that is not met under short-horizon governance, producing a funding gap, declining grid resilience, and further stress, with the costs socialised across all ratepayers. The figure also records what each loop does to the hexalemma, efficient is the only value moving upward, while resilient, sustainable, reliable, safe, and ethical all decline or are hidden. The red dashed arrows mark where the Scenario 3 interventions would cut each loop. None of these interventions currently exists in UK governance, which is why the loops run uninterrupted. The significance of presenting the loops together rather than separately is that they share variables. Concentrated build-out feeds both R1 and R2, and grid stress links R1 to R3. This is why single-loop fixes fail. An intervention that addresses one loop in isolation leaves the shared variables under pressure from the others.

The hexalemma as evaluative lens on causal loop diagram intervention points

A key strength of the hexalemma is that it can assess any proposed change, from small tweaks like pricing signals to major system overhauls. It helps show how each change affects all six values and makes the trade-offs clear. This method works well with CLD[7] methodology

[7] A systems thinking tool that maps variables, causal links, and feedback polarity to show the feedback structure driving system behaviour over time. This paper constructs three CLDs (the Demand Spiral, the Renewable Fiction Loop, and the Resilience Erosion Loop) to map the structural dynamics of AI datacentre demand and grid stress.

instead of competing with it. CLDs show the feedback loops that cause problems in the system. The hexalemma then checks how proposed changes affect all six values, making it clear what is gained and what is lost in each regulatory scenario. In the paper's scenario analysis to 2040, this is the hexalemma's main role. The three scenarios differ not by how bad things get, but by the way governance interacts with the feedback loops identified by the CLDs.

Scenario 1, Unmanaged Growth, is the default trajectory. No new regulatory intervention on datacentre demand, siting, or grid impact beyond what is already in motion. AI Growth Zones[8] proceed with streamlined planning and prioritised connections. RIIO-3[9] funds grid investment but does not differentiate between demand drivers. PPA-based sustainability accounting continues without additionality requirements. Ofgem, National Energy System Operator (NESO), Department for Science, Innovation and Technology (DSIT), and the Planning Inspectorate continue to operate in institutional silos. The Demand Spiral, the Renewable Fiction Loop, and the Resilience Erosion Loop all continue to operate without structural interruption. This is not a worst case. It is what happens if no further governance decisions are made beyond those already committed to. The hexalemma scoring under this scenario is straightforward: efficient scores high in the short run; all other five axes score low and declining over time. The "success to the successful" dynamic identified in the Efficient/Ethical tension compounds across the system

Scenario 2, Efficiency-Led Reform, represents governance responding within the existing paradigm. Efficiency improvements are mandated, stricter PUE thresholds, mandatory WUE[10] reporting, some form of additionality requirement for PPAs, possibly following the EAC[11] inquiry recommendations. Grid connection reform continues along the "first ready, first

---

[8] Designated areas in the UK where datacenter development receives streamlined planning processes, prioritised grid connections, and other policy support. The initial zones are located in Oxfordshire, South and North Wales, North Lincolnshire, and the North East. Established under the UK government's AI Opportunities Action Plan (January 2025).

[9] RIIO-3 (Revenue = Incentives + Innovation + Outputs): Ofgem's regulatory framework for setting the price controls that determine how much electricity and gas network companies can charge consumers. RIIO-3 covers the five-year period from 1 April 2026 to 31 March 2031 and has approved £28.1 billion of initial investment, of which £10.3 billion is allocated to electricity transmission.

[10] Water Usage Effectiveness

[11] Environmental Audit Committee

connected" model. RIIO-3 investment proceeds. Some demand-side signals are introduced, perhaps a differentiated TNUoS12 charge for hyperscale loads. But geographic concentration is not addressed. The fundamental architecture of cost socialisation has not changed. The Demand Spiral is slowed but not interrupted because efficiency improvements reduce the cost per compute, which feeds back into demand growth. This is the Jevons paradox (see Alcott, 2005)13 dynamic, the Demand Spiral loop maps. The fix creates conditions for the problem to return on a larger scale.

The Renewable Fiction Loop is partially disrupted by additionality requirements but not eliminated, because annual matching continues to obscure the temporal mismatch between renewable generation and datacenter consumption. The Resilience Erosion Loop is partially addressed by RIIO-3 investment, but it is reactive, lagging the demand curve because the governance horizon (five-year price control periods) and is shorter than the required infrastructure planning horizon. This is the soft-constraint trap that Koutsandreas and Keppo (2026) model directly. In their analysis, soft constraints such as carbon prices present decision-makers with low near-term costs, which preserves the flexibility to defer structural investment; the deferral looks rational at each step because the future is heavily discounted, but it locks the system into a more expensive and less resilient trajectory and passes the transformation cost onto later periods.

Hard constraints behave differently, creating strong dependencies across years that remove the option to postpone. Efficiency-led reform is built almost entirely from soft instruments, PUE thresholds, reporting standards, differentiated charges, so it leaves precisely this room for deferral. The Resilience Erosion Loop therefore continues to run; each price control period addresses the stress already visible rather than the stress accumulating ahead of it.

The assessment of this scenario reveals a mixed outcome. While efficiency continues to perform strongly and sustainability shows some improvement as reporting becomes more

[12] Transmission Network Use of System

[13] The observation, first made by William Stanley Jevons in 1865 regarding coal consumption, that improvements in the efficiency of resource use tend to increase rather than decrease total consumption, because lower per-unit cost stimulates greater demand. In the datacentre context, efficiency improvements that reduce cost per AI compute unit stimulate demand growth that outweighs the per-unit saving.

grounded in actual conditions, there is little progress in resilience, reliability, safety, or ethics. This is primarily because the underlying issues, such as the concentration of resources in specific locations, the way costs are distributed, and the fragmentation of institutions, remain unaddressed. As a result, this scenario exemplifies a governance approach that implements surface-level solutions without tackling the root causes, echoing what is often described as the 'fixes that fail' pattern.

Scenario 3, referred to as Structural Transition, represents a significant shift in governance by prioritising all six values, not just efficiency. The measures proposed are intentionally robust, moving away from voluntary or flexible approaches. These include requirements for continuous carbon-free energy matching, rules to ensure data centres are more geographically dispersed, and a system where those who drive the need for grid upgrades bear the associated costs. There are also obligations for community benefits to address the tension between efficiency and ethics, and a new integrated planning body would be responsible for evaluating the combined impact of data centres on the grid, involving Ofgem, NESO, DSIT, and the Planning Inspectorate.

Additional measures include the ability to curtail non-essential AI workloads during periods of stress and enforceable water usage limits based on regional conditions. The importance of these hard constraints lies in their ability to create dependencies across years, making it impossible to simply postpone necessary changes. As Koutsandreas and Keppo (2026) demonstrate, it is the firmness of these constraints, rather than the ambition of the policies themselves, that determines whether the problematic feedback loops are effectively addressed.

The scenario does not eliminate the tensions. The hexalemma's claim is that they are structurally unavoidable. What it does is move governance from privileging efficiency at the expense of the other five to holding all six in view and making the trade-offs explicit. The scoring reflects a conscious trade-off. Efficiency drops to moderate, because internalising externalities raises cost per compute, siting and environmental limits slow deployment, and flexibility requirements relax the always-on assumption. The other five rise because the feedback loops eroding them are interrupted at source rather than left to compound. The hexalemma does not declare Scenario 3 optimal. It frames the real question, which is what governance is prepared to accept lower efficiency, higher cost, and slower deployment in exchange for gains on the five values efficiency currently overrides.

The hexalemma's five most analytically central tensions for this paper

- Resilient/Efficient: The structural heart of the paper's argument. PUE-focused policy is a move toward efficiency that systematically erodes grid resilience. This is Meadows' (2008) trade-off, explicitly stated and institutionally embedded in how data centres are currently regulated.
- Sustainable/Efficient: Formalised as the Renewable Fiction Loop. Efficiency in reported emissions (via PPAs) diverges from actual grid carbon intensity. The system scores well on reporting efficiency but poorly on genuine sustainability. This is a dualism masquerading as a duality.
- Resilient/Ethical: Datacenter geographic concentration drives regional grid stress, which in turn drives infrastructure costs, which are socialised to all consumers, including those in regions that receive none of the economic benefits. The resilience problem and the justice problem are the same structural problem seen from different axes.
- Sustainable/Reliable: Renewable intermittency compounds the inflexibility of AI demand. Datacenters require near-constant power; variable renewables require demand flexibility. The grid is asked to bridge this mismatch. Scoring well on sustainability requires accepting lower reliability scores under current technology conditions unless substantial storage and dispatchable backup are added, at cost, and with their own sustainability implications.
- Efficient/Ethical: The energy justice dimension, the Schumacher Institute is institutionally positioned to foreground. Efficiency gains at the system level are being captured by a small number of large technology companies, while the costs, grid upgrade costs, carbon costs, and land use costs are distributed far more widely. This is the allocation dimension of the hexalemma, scarcity and justice.

Systems thinking theory: contribution and added value

Systems thinking has a substantial theoretical pedigree. Arnold and Wade (2015) define it as "a set of synergistic analytic skills used to improve the capability of identifying and understanding systems, predicting their behaviours, and devising modifications to them in order to produce desired effects." This is the definition the paper draws on, and it is worth being precise about what it entails. Systems thinking is the opposite of linear thinking; it involves recognising that repeated events and patterns derive from systemic structures, which in turn derive from underlying mental models; it focuses on relationships rather than components and attends to self-organisation and emergence (Bhuyan and Broom, 2025).

The specific tools systems thinking employs include causal loop diagrams with feedback and delays, stock-and-flow diagrams, behaviour-over-time graphs, and scenario-based modelling. A causal loop diagram provides a powerful visual tool for scoping out the feedback structure in a problem, showing variables, causal links, and link polarity to describe the feedback structure driving system behaviour over time. The most complex problems involve multiple feedback loops that together participate in an ongoing tug-of-war to shape system behaviour, and the feedback structure and delays in a system are major determinants of how the system behaves and why it is so difficult to manage (Deaton and MacDonald, 2025).

The added value of applying systems thinking to the AI datacenter/grid stress problem is that the energy sector is an intrinsically dynamic and complex system whose behaviour is not solely controlled by its constituent components but rather a consequence of dynamic interactions among them. Despite this, the interconnections within the energy sector have received little attention in research and policy, and current policies promoting energy development are likely to be "fixes that fail" that ultimately undermine sustainability. This is the precise analytical gap the paper addresses (Laimon et al., 2022).

The "fixes that fail" archetype is a standard systems-thinking diagnosis, a symptomatic solution (efficiency improvements, PPA adoption) relieves pressure in the short run while leaving the underlying structure unchanged or reinforcing it, so the problem recurs at a larger scale. The Demand Spiral is structurally a "fixes that fail" loop, efficiency investments reduce the cost per compute, AI capability grows, demand grows further, and the fix generates the problem it was meant to solve.

More broadly, energy demand surged 2.2% in 2024, driven by electrification, including from data centre expansion, placing additional pressure on grid infrastructure, while energy security progress remained minimal, due in part to declines in power system flexibility. Infrastructure investments must prioritise grid modernisation, decentralisation, and digitalisation to enhance system flexibility and support the integration of renewables. What that policy prescription from the World Economic Forum (see Bocca and Ashraf, 2025) does not supply, and what the paper's systems thinking framework does, is a structural account of why that investment is not happening at the required rate and scale, and what feedback dynamics are preventing it. That is what the CLD methodology contributes, and what the temporal inconsistency framework (Koutsandreas and Keppo 2026) formalises theoretically.

The Bhandari et al. (2025) paper on water utilities is a useful methodological parallel here. It demonstrates exactly the move this paper makes, applying a systems thinking framework and CLD-based visual mapping to a complex socio-technical infrastructure domain where existing governance is siloed, data management is fragmented, and the interconnections among influencing factors are structurally underexplored. Bhandari et al. explicitly argue that systems thinking, mapping, and modelling can identify risks by visualising feedback loops, dependencies, and potential points of failure, and can identify leverage points, nodes where targeted interventions yield significant systemic impact. The same logic applies to the datacenter/grid system, the CLD loops identify the structural feedback generating grid stress, and the hexalemma evaluates proposed leverage points against all six values.

The Lord and Widdicks (2026) blog from UKCEH reinforces this from a practitioner context. Reflecting on a conference session on whole-systems evidence, the authors report a consistent message: datacentres need to be understood as part of wider systems rather than in isolation, embedded within a broader network of environmental and societal systems that brings additional considerations including material demand, heat generation, land use, and seasonal variability. Lord raises the rebound dynamic directly, as digital technologies become more efficient, accessible, and embedded in everyday life, demand for computation can grow rapidly, offsetting or even outweighing expected environmental gains. This is the Jevons paradox operating in the digital domain (Alcott, 2005), and it is the dynamic the Demand Spiral loop maps. The whole-systems practitioners identify the pattern; this paper supplies the structural explanation.

Why this theoretical stack works together

The paper's theoretical architecture is coherent and non-redundant. Arnold and Wade (2015) define systems thinking and specify the analytic skills it comprises, four of which, recognising interconnections, identifying feedback, understanding system structure, and understanding dynamic behaviour, are precisely what a causal loop diagram operationalises. Qudrat-Ullah (2025) provides the methodological grounding for that operationalisation, reviewing how CLDs map reinforcing and balancing loops to identify leverage points in energy and climate systems, and identifying as a standing gap the underexplored feedback between climate impacts and energy-system viability. This paper extends that methodology into a domain his review's scope leaves untouched: few studies have applied CLD-based systems analysis to AI datacentre demand and grid stress in a forward-looking, scenario-based form. Koutsandreas and Keppo (2026) supply the mechanism explaining why governance fails structurally rather than contingently. Modelling decision-makers who discount the future at escalating rates, they show that what is judged optimal shifts as the planning horizon moves, so decisions taken early under short-term bias lock the system into sub-optimal pathways, raise total system costs, and push transformation costs onto later periods and future generations. The hexalemma then adds a layer that asks not just what the system is doing and why governance fails, but which values are systematically sacrificed and which privileged, and how any proposed intervention scores across all six axes.

The hexalemma also gives the paper its Schumacher Institute positioning. The institute's critique of efficiency-led growth as the primary measure of system success is what the hexalemma formalises, pursuit of the efficient corner at the expense of the other five is the structural problem. Schumacher (1973) argued that economies treating efficiency as the goal rather than a tool produce systems that are brittle and unjust. Meadows (2008) supplies the mechanism, resilience depends on the redundancy that reductionist efficiency strips out and is often the first thing sacrificed for short-term productivity. The datacentre growth model is extractive in exactly this sense, with efficiency gains concentrated in a few firms while the costs of grid infrastructure, water depletion, carbon emissions, and community disruption fall on populations and ecosystems with no stake in the value created. The paper draws on this tradition without being parochial. The analysis stands on its own systems-theoretic and empirical grounds, and the hexalemma makes the trade-offs explicit and measurable.

The fifteen hexalemma trade-offs applied to AI datacentre expansion

The governance gap that the hexalemma analysis addresses is visible in the UK's own regulatory trajectory. The UK AI Opportunities Action Plan (January 2025) committed to expanding sovereign compute capacity by at least twentyfold by 2030 and establishing AI Growth Zones with streamlined planning, but it contains no binding energy demand management framework, no additionality requirements for datacentre renewable procurement, and no mechanism for assessing cumulative grid impact across the proposed zones. The EU AI Act, while establishing risk-based governance for AI systems, similarly does not integrate AI energy demand into binding regulatory structures. Both frameworks acknowledge the problem, and neither addresses it systemically.

Within the UK, the regulatory response is fragmented across multiple bodies. Ofgem regulates energy pricing and network investment through the RIIO framework. NESO manages grid connections and system planning. DSIT drives AI industrial strategy and the Growth Zones programme. The Planning Inspectorate handles Nationally Significant Infrastructure Projects. No single body is responsible for assessing the systemic interaction between AI data centre demand growth and grid integrity, which is precisely the kind of governance fragmentation that the CLDs in this paper are designed to make visible.

The House of Commons Environmental Audit Committee opened its inquiry into the environmental impact of UK datacentres on 27 February 2026, with a report expected later this year (UK Parliament, 2026a). The inquiry terms of reference ask whether planning processes adequately account for environmental impacts and how much energy and water data centres are likely to use. Significantly, the Committee also published a letter from Secretary of State for Energy Security Ed Miliband in which he acknowledged that future datacenter demand is "inherently uncertain" (UK Parliament, 2026b), a statement that itself illustrates the temporal inconsistency problem: governance decisions are being made now (AI Growth Zones, NSIP designation, RIIO-3 investment) on the basis of demand projections that the responsible minister describes as uncertain.

The AI Growth Zones programme itself was established under the AI Opportunities Action Plan (GOV.UK, 2025), which committed to expanding sovereign compute capacity by at least twenty-fold by 2030 but contains no binding energy demand management framework. The scale and planning implications of this expansion are synthesised in the House of Commons Library briefing on datacentre planning policy, sustainability, and resilience (Clark, 2026).

The UK's starting position on the energy trilemma is already strained. The ARUP Five Minute Energy Guide to the Energy Trilemma shows the UK's trilemma weighting, based on GDP and natural resources, tilted toward energy security (43 per cent) over sustainability (36 per cent) and affordability (21 per cent), with the UK ranking 10th on security, 13th on sustainability, and 10th on affordability in the World Energy Council's Trilemma Index (ARUP, 2017; World Energy Council, 2019). AI-driven demand growth puts pressure on all three dimensions simultaneously, security is threatened by concentrated demand exceeding regional grid capacity; sustainability is threatened by the accounting displacement mechanisms documented in the Renewable Fiction Loop; and affordability is threatened by the network cost increases being passed to all consumers through RIIO-3 and the Ofgem price cap. The hexalemma extends this analysis by adding the three dimensions the trilemma omits, reliability, safety, and the ethical distribution of costs and benefits. The World Energy Council's own 2021 Index found the UK performing strongly overall but noted that energy poverty extends to all geographies including the most developed nations (World Energy Council, 2021), a finding that the hexalemma's ethical axis is designed to make structurally visible.

A preliminary note on structure; with six axes, there are fifteen pairwise tensions. The hexalemma framework distinguishes these tensions as either dualisms (genuinely opposed, either/or) or dualities (interdependent, both/and, in which the values can in principle reinforce one another if conditions are right). In the datacentre context, most tensions manifest primarily as dualisms under current governance arrangements, values that should ideally cooperate are structurally forced apart by the way the sector is governed, financed, and sited. Where a tension has duality potential, that is noted, because it is often the basis for the policy alternatives identified in the scenario analysis.

The following table maps each tension pair to its predominant dualistic or duality under current UK governance conditions. Where duality potential exists, the conditions required for it are noted. This classification is not fixed as it is a function of the current regulatory and market architecture, and several tensions would shift from dualism to duality under the governance arrangements explored in Scenario 3.

Table 1. The fifteen hexalemma tensions in the AI datacentre context: classification as dualism or duality under current UK governance and associated causal loop.

#	Tension pair	Set	What is in tension in the datacentre context	Direction under current UK governance	CLD loop
1	**Resilient / Efficient**	Hard	PUE-driven design rewards dense, co-located clusters; the efficiency optimum strips out the	Dualism (duality potential via demand-	Demand Spiral
2	**Resilient / Sustainable**	Hard	Short-run resilience maintained through carbon-intensive backup and delayed coal	Dualism short-run; duality long-run	Renewable Fiction
3	**Sustainable / Efficient**	Hard	Efficient carbon accounting (RECs, annual matching) lets reported emissions fall while	Dualism presented as duality	Renewable Fiction
4	**Reliable / Safe**	Soft	Infrastructure guaranteeing uninterrupted compute (water cooling, diesel backup) imposes	Dualism	Demand Spiral
5	**Reliable / Ethical**	Soft	Reliable compute for a global user base is cross subsidised by local ratepayers who neither	Dualism	Resilience Erosion
6	**Safe / Ethical**	Soft	Facility-level water and energy data withheld as commercial information; communities	Dualism (mutually reinforcing failure)	-
7	**Resilient / Reliable**	Cross	Behind-the-meter generation makes individual facilities grid-independent but removes	Dualism (duality potential via grid-service	Demand Spiral
8	**Resilient / Safe**	Cross	Grid stress from concentrated demand is a public safety problem falling hardest on	Dualism	Resilience Erosion
9	**Resilient / Ethical**	Cross	Resilience investment costs are socialised across the rate base while benefits concentrate in	Dualism	Resilience Erosion
10	**Sustainable / Reliable**	Cross	Renewables are intermittent; AI load needs firm continuous power; storage and demand	Dualism (hardest to resolve)	-
11	**Sustainable / Safe**	Cross	Water cooling lowers PUE and scores well on sustainability metrics while consuming	Dualism	Renewable Fiction
12	**Sustainable / Ethical**	Cross	Voluntary corporate accounting defines what counts as sustainable, serving investors while	Dualism	Renewable Fiction
13	**Efficient / Reliable**	Cross	Operator-level efficiency (queue bypass, off-grid generation) raises individual reliability	Dualism	Demand Spiral
14	**Efficient / Safe**	Cross	Maximising rack density and utilisation concentrates heat, fire and power-quality risk not	Dualism	Demand Spiral
15	**Efficient / Ethical**	Cross	Efficiency gains accrue to a few firms; externalised costs spread across communities,	Dualism (the core allocation tension)	All three loops

Hard set: tensions among the system-facing values (Resilient, Sustainable, Efficient). Soft set: tensions among the human-facing values (Reliable, Safe, Ethical). Cross set: tensions pairing one value from each. Duality potential, where noted, indicates the conditions under which the tension could become generative; these conditions form the basis of the Scenario 3 interventions

1. Resilient/Efficient

This is the core tension that led to the development of the hexalemma. Meadows (2008) explains that increasing efficiency often reduces resilience. Over the last century, resilient systems have been replaced by more efficient ones-larger, less diverse, and with less redundancy. The datacentre industry is a clear example of this trend today.

The primary efficiency metric governing datacentre design and regulation is Power Usage Effectiveness (PUE), the ratio of total facility energy consumption to IT equipment energy consumption. A PUE of 1.0 would mean all energy reaches the computing load with none lost to cooling, lighting, or other overheads. The largest hyperscale operators now report fleet-average PUEs approaching 1.09 (Google) and 1.08 (Meta), representing genuine engineering achievements in thermal management. Google reports a comprehensive trailing twelve-month PUE of 1.09 across its large-scale data centres, with some sites operating below 1.06, while Meta's 2024 Sustainability Report states an average PUE of 1.08, indicating that IT computing accounts for most of the electricity consumption in these facilities (Chen et al., 2025).

The issue is not that optimising for PUE is bad - Meadows (2008) notes it is usually a good thing. The problem is what gets left out. Focusing on PUE leads to large, tightly packed, and co-located computing clusters. These setups are efficient and cost-effective, but they also mean data centres are concentrated in a few places, rely heavily on constant grid power, and cannot easily adjust their energy use. Hyperscale AI training clusters cannot quickly reduce their load to help balance the grid. Their power use is steady and hard to change. Most AI data centres today use a firm service model, so the grid must always be ready to supply whatever power they need. Both AI training and inference consume substantial resources and require extensive, intensive computing cycles across many connected GPUs or TPUs (Chen and Zheng, 2026).

A single large training run synchronises tens of thousands of these GPUs, and GPU-based racks can draw up to six times the power of conventional racks, raising both cooling intensity and peak load (Chen et al., 2026). Because the accelerators must draw power continuously and in step, the load cannot be ramped down for grid balancing the way conventional industrial demand can. The firm service model is therefore not a contractual preference but a physical property of the hardware, the GPU cluster is the reason the demand is inflexible.

This focus on efficiency has major effects on the grid. Places like Oregon, Virginia, and Ireland can see Power Stress Index14 values above 0.25, indicating that their grids are vulnerable. In contrast, more diverse systems can better handle new demands. Efficient data centres tend to cluster in a few areas, often because of land prices, fibre access, and existing infrastructure. This concentration puts extra stress on local grids. If data centres were spread out and less optimised, their PUE would be higher, but the overall grid would be less stressed. In short, efficiency and resilience often pull in opposite directions (Chen et al., 2026).

The UK illustrates this dynamic with clarity. The previous first-come, first-served grid connection model led to the connections queue growing tenfold in five years, reaching over 700 GW across transmission and distribution, approximately four times what Great Britain requires by 2050 (NESO, 2025). NESO's December 2025 reform removed more than 300 GW of stalled projects and advanced 283 GW, but the reformed pipeline still stands at 238 GW (NESO, 2025). Within this queue, NESO identified approximately 140 datacentre facilities totalling 50 GW of demand, more than the entire peak electricity demand of the British grid, which stood at 45 GW on 11 February 2026 (Ofgem, 2026a). The queue for grid connections reportedly increased by 460 per cent during the first half of 2025 (IDCA, 2026).

National Grid has projected that datacentre demand could rise from approximately 2.5 per cent of UK electricity consumption to 6 per cent in 2026 and potentially 9 per cent by 2035 (Data Center Dynamics, 2025), while NESO projects that datacentre electricity consumption could quadruple by 2030 (Clark, 2026). National Grid ESO executives have warned that there will be blackouts in the South East by 2028 if transmission constraints are not addressed (Energy Institute, 2024).

The siting pattern in the UK mirrors the US concentration logic. Most existing datacentres, and half of those in planning, are concentrated around London and the South East, a region already classified as severely water-stressed and subject to significant transmission constraints (Global Action Plan, 2026). The government's NPPF consultation (December 2025) proposes

[14] The Power Stress Index (PSI), from Chen et al. (2026), is the ratio of projected datacenter demand to regional generation capacity. Above 0.25 is moderate stress, above 0.40 high stress, in line with IEA and EPRI guidance that industrial loads exceeding 25 to 30 per cent of regional generation can affect reliability and reserve margins.

classifying datacentres as "essential infrastructure" for flood risk purposes and co-locating large-scale generators and power users, further embedding concentration logic into planning policy rather than challenging it (Tucker, 2025). Ofgem's RIIO-3 framework for 2026 to 2031 confirms £10.3 billion of investment in electricity transmission, funding 80 transmission projects (Ofgem, 2025a). However, the investment is driven primarily by the Clean Power 2030 target and general grid modernisation, not by a systemic assessment of datacentre-driven demand.

The RIIO-3 framework intends to "fund the efficient transition to a clean power system while protecting consumers" (Ofgem, 2025a), but it does not differentiate between demand growth from electrification of transport and heating (which serves a broad public good) and demand growth from concentrated hyperscale AI infrastructure (which serves a narrow set of commercial actors). The cost of both is socialised identically through network charges. Ofgem's own regulatory mapping exercise found no direct lever for managing the interaction between AI demand and grid resilience as a distinct governance problem (Ofgem, 2025).

By the mid-2030s, the country will probably need thousands of miles of new transmission lines. The real choice is between building a fragile, overextended system or creating a more resilient grid that uses a variety of solutions. Right now, the push for efficiency is making the system more brittle. True resilience needs backup options, diversity, and extra capacity, even though these are seen as inefficient from a strict engineering perspective (Gaster, 2025).

Right now, this tension mostly works as a dualism, focusing on PUE means we avoid the redundancy and spread-out locations that make systems resilient. There is potential for a better balance if we had more distributed and flexible computing, which could help the grid. But making that happen would need a very different set of rules and system designs than we have today.

2. Resilient/Sustainable

This tension can be seen as either a dualism or a duality, depending on the timeframe. In the short term, building resilience often involves investing in backup power sources such as diesel generators or gas plants, which can undermine sustainability. Over the long term, a truly sustainable grid becomes more resilient, since renewable energy sources spread across different locations are less vulnerable to disruption than centralised fossil fuel systems. This is why the tension is the most clearly time-dependent of the hard set, as distributed renewables-and-storage grid are both more sustainable and more resilient than centralised fossil backup, so what reads as a dualism on a five-year horizon can become a duality on a fifteen-year one.

For datacentres, this tension mostly shows up as a dualism, and in a concerning way, actions taken for sustainability are weakening grid resilience, not the other way around. Most companies focus on buying renewable energy through PPAs to meet sustainability goals. This approach looks good on sustainability reports but does little to improve grid resilience.

The main problem is that PPAs, even when well-designed, do not provide extra power where and when AI datacentres need it. This means datacentre operators might push out other users without helping the overall power system, especially since datacentres already make up 43% of clean power purchase agreements in 2024. For example, if a large company signs a PPA for wind power in one area but relies on steady power from a different grid mix elsewhere, the sustainability numbers look good, but grid resilience remains the same or even worsens (Lyu and Tang, 2025).

The rush to expand AI capacity is even causing some coal plants to delay retirement, reversing earlier decarbonization plans. In the US, refurbishing a coal plant to keep it running past 40 years can cost up to $1.3 billion, and operating costs have risen by 28% from 2021 to 2024 because of higher fuel and capital expenses. This is the clearest example of the Resilient/Sustainable dualism: grid operators keep unsustainable backup power just to maintain reliability as AI demand grows (Yañez-Barnuevo, 2026).

In the UK, this tension takes a unique regulatory form. On 27 February 2026, the House of Commons Environmental Audit Committee began an inquiry into the environmental impact of UK data centres, with a report due later this year (UK Parliament, 2026a). Secretary of State for Energy Security Ed Miliband admitted that future data centre demand is "inherently uncertain" (UK Parliament, 2026b). The inquiry is looking at whether planning processes properly consider environmental impacts and how much energy and water data centres will use. Importantly, the Committee also found that data centres were excluded from the government's projections for the Seventh Carbon Budget (CB7) (UK Parliament, 2026b). This gap in national sustainability accounting is similar to the PPA displacement issue at the corporate level, if datacentre demand is not included in carbon budget forecasts, the national sustainability picture is incomplete, just as annual PPA matching leaves gaps at the company level.

Six NGOs wrote to UK Technology Secretary Liz Kendall warning that rising electricity demand from new AI datacentres could increase the UK's carbon emissions, calling for developers to demonstrate that their projects will not cause an increase in overall UK emissions or worsen local water scarcity (Adams et al., 2026). The grid's generation mix, still over a quarter gas-fired, is set out under Tension 10.

A further dimension of this tension is the geography of renewable intermittency. The sustainability case for renewables is strong and widely agreed. Another aspect of this tension is the uneven availability of renewables. While renewables are sustainable overall, wind and solar power can be unpredictable, but AI datacentres need steady power. The grid must fill this gap using storage, better transmission, or backup power, all of which need investment. The current PPA-based sustainability approach does not sufficiently encourage these solutions. These would score well on both axes, but the conditions for it do not currently exist at scale, and the governance framework does not require it.

3. Sustainable/Efficient

This is the formal version of the Renewable Fiction Loop. The main issue is that current accounting rules create a split. Efforts to make carbon reporting more efficient, using tools like PPAs, RECs, and annual matching, often move further away from real sustainability, which should be measured by grid carbon intensity and whether new renewables are added.

This process is well-known. A big tech company promises to use 100% renewable energy, mainly by buying RECs or Virtual Power Purchase Agreements (VPPAs). These let the company claim renewable energy from anywhere on the grid, at any time, even if the actual electricity powering the datacentre is not carbon-free. The company can report that it met its sustainability goal at a low cost, but the real makeup of the grid does not change.

Meta uses RECs from VPPAs and other offsets to keep up its 100% renewable energy claims and say it is carbon neutral, even if it does not add new renewable energy or fully cut its emissions. For example, in Louisiana, Meta is building a huge data centre, and the local utility, Entergy, plans to build three methane gas power plants to meet the extra demand.

Google's emissions rose 51% from 2019 to 2024, even as it continued to make renewable energy claims. Because of efficient reporting methods such as annual matching and aggregate accounting, the company appears strong on sustainability reports, even though the grid's carbon intensity is increasing (Yagüe and Ashtine, 2025).

The biggest issue here is additionality. For a renewable energy purchase to count as additional, it must lead to new renewable projects that would not have happened otherwise. Without this, corporate PPAs mostly just shift credit for existing renewables around, instead of increasing the total amount of clean energy. In 2024, internet giants accounted for 43% of clean PPAs, suggesting data centre operators might be pushing out other buyers without adding any real benefit to the power system (Lyu and Tang, 2025). The UK's regulatory system clearly shows this problem.

Ofgem reviewed current energy laws against AI principles and found that most principles are only loosely connected to the rules. They are not specifically included in any licence or law but are only hinted at through broader requirements. There is no rule in the Electricity Act 1989 that requires additionality for data centre renewable energy purchases. There is also no rule for 24/7 carbon matching or for verifying whether corporate PPAs are just shifting clean power from other users rather than adding new generation. The Renewable Fiction Loop exists in the UK not because regulators are neglectful, but because the rules were made for a different time and different types of companies (Ofgem, 2025a).

The Climate Change Committee and the Environmental Audit Committee are both expected to examine the use of unbundled renewable energy certificates, which essentially relabel existing clean power as new capacity. In the UK, contracts for large data centres are increasingly requiring tenants to demonstrate that their energy purchases add new megawatts to the grid, mirroring trends in the US. However, these are private contracts, not legal requirements. The gap between what companies agree to in contracts and what the law requires is where the Renewable Fiction Loop continues in the UK (Murray, 2026).

A House of Commons Library briefing from 27 May 2026 shows the situation: data centres now use about 2.5 per cent of the UK's electricity, and this could increase by four times by 2030. Big datacentre companies have promised to meet voluntary sustainability goals by 2030 and to support renewables through PPAs, but the briefing notes that some studies question whether PPAs help reduce emissions. The word "voluntary" is important here: it shows the difference between efficient reporting and real sustainability. Voluntary commitments look good for efficiency, but whether they help sustainability depends on what they deliver to the grid. Right now, the UK has no rule to check this (Clark, 2026).

In this way, efficient corporate sustainability reporting can undermine real sustainability. It lets companies meet ESG goals and satisfy investors at low cost, without making the big changes-like 24/7 carbon-free power, adding new renewables, or adjusting demand to match renewable supply-that would lower grid carbon intensity. On paper, efficiency and sustainability seem to go hand in hand, but they are often at odds on the grid.

4. Reliable/Safe

Now focusing on the soft RSE set: this tension is about the quality and care aspects of how systems affect people. For data centres, this means the need for reliable AI infrastructure can pose safety risks to nearby communities.

Reliability in this context is about whether the system consistently works as expected. For datacentre operators, it means always having power, no downtime, and steady computing performance. These factors support service agreements with cloud customers and explain the large investments in building these facilities. Backup generators, uninterruptible power supplies, and extra cooling systems all help ensure reliability. Safety is about preventing harm to people, the environment, and public health. The systems that make datacentres reliable can also create safety risks that are not always visible in standard metrics. For example, backup generators often run on diesel or gas, releasing pollutants into nearby communities. Water-cooled data centres, chosen for their energy efficiency, use large amounts of freshwater and return it to local water sources at higher temperatures, with added chemicals.

Data centres cost the US economy $25 billion a year in hidden health and environmental damage, according to a National Bureau of Economic Research working paper (Bove, 2026). Google's data centre water consumption increased from 4.3 billion gallons in 2021 to 6.1 billion gallons in 2024. Their largest data centre in Council Bluffs, Iowa, withdrew an average of 3.9 million gallons of water and consumed 2.8 million gallons per day. Cooling water that returns to the water supply has a higher concentration of dissolved solids, such as calcium, chloride, and silica, which can affect the taste of drinking water, lower crop yields, and be toxic to aquatic life. Legionella bacteria present further risks (Lakhanpal, 2026).

The reliability imperative also creates power quality risks that affect surrounding communities. High-frequency switching devices and power electronic equipment in data centres introduce significant harmonic distortion to the grid, degrading local power quality and posing risks to nearby electrical facilities. Sensor data from more than 700,000 homes reveal a strong correlation between proximity to data centres and declining power quality, with effects observed within 20 miles of major data centre clusters (Chen et al., 2025).

In the UK, the water safety dimension of this tension is particularly acute. Thames Water, the UK's largest water utility, has stated that a large data centre may use between 4 and 19 million litres of water per day, equivalent to the daily demand of over 50,000 households (WRc, 2026). This consumption falls within the jurisdiction of a utility carrying more than £13 billion in debt, with ageing infrastructure responsible for losing over 630 million litres of water per day through leaks. Thames Water has considered restricting datacentre water use in London through flow restrictors and peak-demand pricing (Moss, 2023).

Eighty-four per cent of proposed UK data centres are to be built in water-stressed areas, predominantly in London and the South East (Global Action Plan, 2026; Water Magazine, 2026). Affinity Water, which serves parts of London and the South East where a high proportion of datacentres are sited, operates in an area where the regional water supply is already under pressure. The WRc report for MOSL confirms that water use is concentrated almost exclusively in cooling systems, and that approximately 5 per cent of UK facilities may be using hybrid cooling systems that significantly increase consumption during summer, precisely when water stress is highest (WRc, 2026).

The 1 GW Elsham Tech Park project in North Lincolnshire, approved in March 2026, illustrates this tension directly. Anglian Water, which serves the region, has stated that its daily supply limit for non-domestic customers is 20,000 litres per building. Thames Water estimates that hyperscale data centre water consumption is 4 to 19 million litres per day. The gap between what the local water infrastructure can supply and what a hyperscale facility requires is not a marginal engineering challenge. It is a structural mismatch between the reliability demands of AI infrastructure and the safety of the water supply serving the host community (Foxglove, 2025).

This creates a structural divide whereby the engineering choices that make data centres reliable-like water cooling, diesel backup, and large UPS systems-also impose safety costs on the communities where these centres are built. These costs are not included in the facility's financial calculations and do not show up in standard efficiency or carbon metrics.

5. Reliable/Ethical

Reliability means that the infrastructure can be trusted and works as it should. Here, ethics is about fairness, who benefits, who pays, and whether decisions are made openly and fairly. In datacentres, the challenge is to keep AI services reliable for everyone while making sure costs and benefits are shared fairly.

This challenge is clearest when it comes to sharing costs. Making the grid reliable for AI data centres requires building new substations, upgrading power lines, and adding more generation capacity. These costs are often spread across all utility customers, even those who do not use AI services or want data centres nearby.

A new analysis tallies $4.3 billion in 2024 costs across just seven states, stemming from utilities connecting data centres to the grid with expensive new lines while passing those costs to consumers under old rules, with individual connections often costing between $25 million and $100 million (Jacobs, 2025). Utilities requested more than $29 billion in rate increases in the first half of 2025, double the amount requested in the first half of 2024. Electric rate increases were expected to affect 40 million customers nationwide. Utility shut offs skyrocketed to 3.5 million in 2024 and may have reached 4 million in 2025, further straining disadvantaged households.

Electricity prices rose an average of 13% nationwide in 2025. Two in three Americans say utility bills are a source of financial stress, three in four are worried about further increases, and 80 million American households are already struggling to pay their utility bills (Levin, 2026).

In the UK, the cost socialisation mechanism operates through the Ofgem price cap and the RIIO regulatory framework. The energy price cap for a typical dual-fuel household stood at £1,758 in Q1 2026, fell to £1,641 in Q2 2026, and will rise by 13 per cent to £1,862 in Q3 2026 (Ofgem, 2025b; Ofgem, 2026b; Ofgem, 2026c), driven in part by network cost increases under the new RIIO-3 price control period. RIIO-3, covering 2026 to 2031, has approved £28.1 billion of initial investment into electricity and gas networks, of which £10.3 billion is allocated to electricity transmission (Ofgem, 2025c; TechUK, 2025).

Network charges currently constitute 20 to 25 per cent of electricity bills (Flame Energy, 2025). The transmission operator revenue requirement increases by 149 per cent compared to 2025/26 levels (Powerhub Solutions, 2025). Household transmission charges are projected to rise from approximately £23 per year in 2026 to approximately £44 per year by 2031 (National Grid, 2026), with overall network charges (electricity and gas combined) increasing from approximately £222 to approximately £330 per year over the same period (Powerhub Solutions, 2025).

These increases fall on a population already under significant energy cost pressure. The 2026 Fuel Poverty Statistics (covering 2025 data) report that 2.36 million households (9.4 per cent) in England are in fuel poverty under the Low Income Low Energy Efficiency metric, with an aggregate fuel poverty gap of £896 million. Under the broader 10 per cent of income threshold, 7.63 million households (30.4 per cent) exceeded this level in 2025 (GOV.UK, 2026b). Across the UK nearly 5 million in deep fuel poverty, spending more than 20 per cent (End Fuel Poverty Coalition, 2025).

The ethical dimension is sharpened by the fact that the grid investment driving these increases is partly responsive to the surge in datacentre demand connection requests. The 50 GW of datacentre demand identified by NESO in the UK connections queue requires transmission infrastructure that will be funded through network charges socialised across all consumers. The

RIIO-3 framework does not differentiate between demand growth that serves a broad public good and that which serves concentrated commercial interests. The cost is identical on the bill regardless of who caused it or who benefits from it.

The ethical dimension is further sharpened by the geography of who bears these costs. The communities most exposed to rate increases driven by datacentre demand are often not those receiving the economic benefits, high-skill employment, tax revenue, or digital services. Rural communities in Virginia, Ohio, and Oregon are experiencing rate increases associated with hyperscale development while the computational services delivered from those facilities are consumed globally.

In the UK, the AI Growth Zones in Oxfordshire, North Lincolnshire, South and North Wales, and the North East are being positioned as priority sites for fast-tracked datacentre development. The Infrastructure Planning (Business or Commercial Projects) (Amendment) Regulations 2025 will allow datacentre projects to enter the Nationally Significant Infrastructure Projects consenting process, bypassing local planning authority determination. Communities in these zones will bear the infrastructure costs, grid stress, and environmental impacts while the economic value of AI compute is captured globally. This is the allocation problem at the core of the hexalemma, reliable AI compute for global digital markets is being cross subsidised by local ratepayers who have no meaningful say in the decisions driving that outcome (Burges Salmon, February 2026; DSIT Policy Paper, "Delivering AI Growth Zones," 2025).

The political response is beginning to emerge. In January 2026, US legislation was introduced to update federal utility policy to ensure that massive electricity users, rather than everyday ratepayers, bear the costs of the grid infrastructure they require. Whether this constitutes a genuine resolution of the tension or a rebalancing of costs within the same structural framework remains to be seen. In the UK, no equivalent legislative proposal has been introduced. The RIIO-3 framework socialises costs identically across demand drivers. The question of whether datacentre operators should bear a differentiated share of grid investment costs has not yet entered the UK regulatory debate.

The Resilience Erosion Loop works like this: when AI demand is high in certain places, it drives up infrastructure spending, and those costs are shared by all consumers. This creates an ethical problem that is built into the system, not just caused by political choices.

6. Safe/Ethical

Safety is about preventing harm, while ethics focuses on fairness and proper decision-making. In datacentres, these issues often show up as a lack of transparency. Communities cannot protect themselves from risks they do not know about or participate in decisions that are not open to them.

The issue of water secrecy is a good example. In October 2025, California Governor Newsom vetoed a bill that would have required datacentre operators to share water-use estimates and annual reports with local water suppliers, after industry groups opposed it. Datacentre development is moving faster than the public's understanding of water's role in AI, and regulations are even further behind (Edwards, 2025).

This leads to a double injustice. Communities face real risks, including water shortages, poor power quality, and air pollution. Without clear information about how these facilities affect them, people cannot object to plans, hold operators accountable, or seek legal help. The harm is greater because people do not have the information they need to understand or challenge it. Communities are increasingly pushing back against datacentre siting. A Google facility in South Carolina faced sustained opposition over its proposed groundwater use, prompting a negotiated shift to surface water alternatives. In Memphis, a newly built AI data centre raised alarms among residents over daily withdrawals from ageing public water infrastructure (Privette, Barros and Cai, 2026).

In the UK, the transparency failure operates through a different institutional mechanism but produces the same structural outcome. Foxglove and Global Action Plan's investigation into UK datacentre water consumption found that around half of UK water companies were unable to say how much water they supply to datacentres because they did not know themselves. As The Times reported, the tally is almost certainly a large underestimate, covering only about half of the UK's datacentres, and the figures for Thames Water, by far the largest water supplier to datacentres, are three years old. If the water companies themselves cannot quantify the resource draw, the communities they serve have no basis on which to assess or contest the safety implications of datacentre siting (Foxglove, 2025).

UK community opposition is emerging along lines that directly illustrate this tension. The 1 GW Elsham Tech Park project in North Lincolnshire was approved in March 2026 despite formal opposition from Foxglove and Global Action Plan. Campaigners argued that the council failed to properly weigh the environmental harms, ignored its own policy requiring 20 per cent of energy to be generated through on-site renewables, and accepted the developer's carbon budget figures, which underestimated the datacentre’s impact on the UK's carbon budget by a factor of five (Adshead, 2026). In Hertfordshire, a large datacentre complex near Abbots Langley was approved by Deputy Prime Minister Angela Rayner despite opposition from local planners. The government subsequently conceded the decision should be quashed after Foxglove's legal challenge demonstrated insufficient environmental safeguards (Robinson, 2026b).

The UK government's planning reforms compound this tension. Data centres were designated as Critical National Infrastructure (CNI) in September 2024, with a government civil servant warning at the time that this designation would allow authorities to override local opposition to data centre sites. The subsequent move toward NSIP designation for large datacenter projects (Robinson, 2025), and the proposal in the Planning and Infrastructure Bill 2025 to shield certain infrastructure decisions from judicial review, progressively removes the procedural mechanisms through which affected communities can challenge decisions that affect their safety. The ethical case for transparency, that people affected by industrial decisions should have access to the information needed to participate in those decisions, conflicts not only with the operational interest in keeping facility-level data proprietary but with a planning governance trajectory that is actively reducing the points of democratic engagement (Robinson, 2026a).

Where communities push back, they frequently discover that the planning and regulatory frameworks governing data centre siting treat energy and water consumption as commercial information, withheld from public disclosure. The ethical case for transparency, that people affected by industrial decisions should have access to the information needed to participate in those decisions, conflicts with the operational interest in keeping facility-level data proprietary. In the UK, this is sharpened by the fact that the Environmental Audit Committee's inquiry has had to request information from the Secretary of State on whether data centre demand was even included in national carbon budget projections. If the government itself does not have complete data on the sector's resource draw, the prospect of communities accessing meaningful information for planning objections is structurally remote.

This tension is clear. Safety is about avoiding harm, and ethics is about who gets to make decisions. When people at risk lack information or a real voice, both safety and ethical problems get worse.

7. Resilient/Reliable

This challenge has two parts: the grid's ability to handle stress and the reliability of individual data centres. Ideally, these should work together, but the current growth of AI infrastructure is causing conflict between them.

To improve reliability, data centres invest in their own power sources such as gas turbines, diesel generators, and batteries. This makes each facility less dependent on the grid, but it also means they do not help balance the grid and can even make it less resilient by adding loads that are hard to reduce during emergencies.

Approximately 50 GW of behind-the-meter US gas generation projects were announced in 2025 alone, as datacentres increasingly move to generate electricity on-site to sidestep the interconnection queue. Each of these projects improves the reliability of the individual facility while contributing to the fragmentation of grid management and reducing the pool of flexible demand that grid operators can call on for stability services (Li et al., 2026). The interconnection queue itself illustrates this tension. The median time required from a preliminary interconnection request to commercial operation has steadily increased to over five years. These large load interconnection requests have completely overwhelmed traditional grid capacity expansion pipelines. Datacentres, unwilling to wait for grid connections that serve system-wide resilience, are building behind-the-meter solutions that serve individual facility reliability at the cost of systemic coherence (Chen and Zheng, 2026).

In the UK, the dynamics of the connection queue produce the same structural outcome through a different institutional pathway. The previous first-come, first-served model contributed to a queue of over 700 GW (NESO, 2025). NESO's December 2025 reform replaced this with a "first ready, first connected" approach, which was necessary precisely because the accumulation of individual connection decisions had produced a collectively incoherent system (Norton Rose Fulbright, 2025). But the reform itself creates a new expression of this tension, projects designated as "strategically important," which in practice means datacentres aligned with the AI Growth Zones programme, are prioritised for connection. This serves the reliability of the prioritised facilities while potentially displacing renewable generation and storage projects that serve system-wide resilience (Norton Rose Fulbright, 2025b).

The UK government's AI Opportunities Action Plan committed to expanding sovereign compute capacity by at least twentyfold by 2030. The speed implied by that commitment is structurally at odds with the pace of grid resilience investment. RIIO-3 covers 2026 to 2031 and funds 80 transmission projects over that period. The AI Growth Zones are now being designated. The gap between the pace of demand approval and the pace of resilience infrastructure is not a planning failure that can be resolved administratively. It is a structural expression of this tension; the reliability imperative of the individual sector (AI compute) is being served on a timeline that the collective system's (the grid's) resilience cannot match.

The idea of dualism matters here. If a data centre could adjust its demand, shift flexible AI tasks to times when renewable energy is available, and use its backup systems to support the grid, it could be both reliable and help grid resilience. This is possible, but it needs rules that reward grid support and an understanding that not all require constant, instantaneous power.

8. Resilient/Safe

Grid resilience and public safety are more connected than many people realise. When AI demand puts stress on the grid, it is not just a technical issue; it becomes a public safety concern, especially for those most at risk.

Close to 15 million households have medical devices that require electricity to operate, and in 2023, nearly a third of these households reported being affected by power outages. When electricity rates spike, households must choose between groceries and air conditioning (Vasudevan, 2025).

The Resilience Erosion Loop in the paper's CLD framework captures the mechanism. Geographic concentration of AI demand creates regional grid stress, which requires infrastructure investment, which is not adequately funded under short-horizon governance, leading to declining grid resilience over time. The safety consequences fall most heavily on populations that depend on grid reliability for medical equipment, climate control, and domestic energy needs. The connection between grid resilience and public safety is most visible during extreme weather events, which are themselves increasing in frequency. The 80% of major US power outages between 2000 and 2023 attributable to weather-related events (Climate Central, 2024), indicates how existing grid brittleness is already a safety risk. AI datacenter demand is adding structural load to a system already under environmental stress .

In the UK, the safety consequences of grid stress fall on a population already in energy precarity. The 2026 Fuel Poverty Statistics report that 7.63 million English households (30.4 per cent) spend more than 10 per cent of their income on energy (GOV.UK, 2026b). National Grid ESO executives have warned of blackouts in the South East by 2028 if transmission constraints are not addressed (Energy Institute, 2024). These are not separate problems. The households most exposed to the health consequences of power outages, those dependent on electrically powered medical devices, those in poorly insulated housing requiring continuous heating or cooling, and those unable to afford backup energy sources, are the same households least able to absorb the cost increases driven by RIIO-3 network investment. The Resilience Erosion Loop generates the grid stress; fuel poverty determines who is harmed by it.

The UK's position is further complicated by the age and condition of its grid infrastructure. National Grid Electricity Transmission's RIIO-T3 submission acknowledges that, without investment, bills would be £40 per year higher on average over the RIIO-T3 period due to constraint costs arising from insufficient network capacity (National Grid, 2026). The constraint costs are themselves a resilience indicator and they represent the price the system pays for insufficient capacity and are borne by consumers as a hidden safety-adjacent cost. When constraint costs rise because datacentre demand has consumed available transmission headroom, the grid's capacity to respond to supply disruptions during extreme weather is correspondingly reduced.

This challenge shows that technical problems with grid resilience can pose real safety risks to people. The hexalemma framework helps reveal these links. Focusing only on the technical side misses the safety risks, while looking only at community safety overlooks the grid problems behind them.

9. Resilient/Ethical

This is another example of shared costs, now from the perspective of resilience. Making the grid more resilient is expensive. The ethical question is who should pay for these improvements, and why. The current structure of grid investment governance in most jurisdictions socialises the cost of major infrastructure upgrades across the rate base. All consumers pay, regardless of who caused the need for investment. In the datacentre context, this means that regional grid upgrades driven by hyperscale AI demand are largely funded by the existing customer base, including residential and small commercial customers who use far less electricity and benefit far less from AI services. In 2024, utilities in seven mid-Atlantic PJM states passed $4.3 billion in transmission costs to ratepayers (Jacobs, 2025), while data centre load drove up the cost of PJM's most recent capacity auction by $7.3 billion, an 82 per cent increase, to $16.1 billion (Howland, 2025). These costs are ultimately passed on to consumers through their energy bills.

This is an ethical problem in how resources are shared. Grid capacity and infrastructure are limited, but large tech companies use their influence to secure connections, while the costs are spread across all customers under old rules. The industry claims everyone benefits from better digital services and economic growth, but most of the benefits go to a few companies and their shareholders.

In the UK, the government's framing reinforces this asymmetry. The AI Opportunities Action Plan positions data centre expansion as central to economic growth and international competitiveness. The AI Growth Zones programme offers streamlined planning and prioritised grid connections. Designating data centres as CNI provides legal protection. At no point in this policy architecture is there a mechanism to ensure that the communities bearing the grid and water costs, and the planning disruption, receive a proportionate share of the economic value generated. The ethical question is not whether AI infrastructure has economic value. It is whether the current distribution of costs and benefits is defensible and, on what basis, the communities paying the infrastructure costs have consented to that arrangement.

This tension and Tension 5 describe the same structural pattern from different angles in the hexalemma. Both result from the Resilience Erosion Loop’s cost socialisation mechanism.

10. Sustainable/Reliable

Renewable energy is not always available, but AI datacentres need steady, reliable power. This is a key challenge in the debate about how data centres fit into the energy transition, and it is one of the most important issues in the hexalemma analysis. A system powered mainly by wind and solar produces variable output that follows the weather, not demand. AI model training is one of the most computationally intensive and time-consuming tasks in modern computing. A large training run can take weeks and cannot be interrupted without losing progress. Inference workloads that serve real-time queries are even less flexible.

The conventional response to this tension is storage and demand response. Battery storage smooths intermittency; demand response allows some datacentre loads to shift temporally to match renewable availability. Both exist and are developing rapidly. But neither is yet at the scale required to bridge the gap between AI datacentre load profiles and renewable generation profiles, particularly as AI demand grows faster than both storage deployment and renewable generation capacity.
AI growth alone is projected to quadruple AI-optimised datacenter electricity demand by 2030. Without a dramatic shift in how these centres are powered, this trajectory has the potential to inflate our global carbon footprint, bringing us closer to breaching 1.5°C of global warming (Yagüe and Ashtine, 2025).

The industry response has been to turn to firm, low-carbon sources, with nuclear receiving renewed investment because it provides reliable, carbon-free baseload. Across 2024 and 2025, the major cloud providers committed billions to nuclear partnerships. Amazon anchored a roughly $500 million financing round in X-energy in 2024 (Waltz, 2024), Google committed early-stage capital to Elementl Power for 1.8 GW across three US sites in 2025, and Microsoft signed a 20-year agreement to purchase the output of the restarted Three Mile Island Unit 1, with Constellation Energy as operator (2024) (see Clancy, 2025). Twenty-two gigawatts of nuclear projects are now in development globally. Nuclear addresses the Sustainable/Reliable tension more effectively than variable renewables but introduces its own tensions on the resilient and safe axes (long construction timelines, cost overruns, waste storage), as well as ethical questions about who decides where facilities are built (Crosley, 2026).

In the UK this tension has a specific shape set by the grid's generation mix. In 2025 wind was the single largest source of British electricity, but gas still supplied more than a quarter of it and the grid drew a further 15 per cent from imports (NESO, 2026). That matters because the government is pursuing two goals at once on the same timeline. Clean Power 2030 aims for a largely decarbonised grid by the end of the decade, while the AI Growth Zones programme invites a steep rise in datacentre demand, which NESO projects could quadruple by 2030.

The hexalemma makes the conflict between them explicit. Datacentre load is firm and inflexible, it needs power at the moment of consumption, whereas the renewables meant to decarbonise the grid are variable. Every unit of datacentre demand that variable renewables cannot meet in real time has to be met instead by gas, imports, or storage. So pushing hard on reliability, guaranteeing always-on power for AI, pulls directly against the sustainability target, because the gap is filled by exactly the firm sources Clean Power 2030 is trying to retire. This is the Sustainable/Reliable tension in its sharpest form, and the two halves are currently being planned in separate institutional silos. The clearest sign of that is the Environmental Audit Committee's finding that datacentre demand was left out of the Seventh Carbon Budget projections altogether (UK Parliament, 2026b).

Nuclear is the obvious resolution in principle, firm low-carbon power that matches the AI load profile, but in practice the timelines do not line up. Hinkley Point C is over budget and behind schedule, Sizewell C is years from delivery, and no small modular reactor is yet operational in the UK (Loughran, 2026; Francis, 2026). The AI Growth Zones are being designated now; the nuclear capacity that could power them cleanly and reliably is a decade or more away. The tension is real, structural, and currently unresolved, and the hexalemma's contribution is to show that it cannot be answered by optimising reliability or sustainability alone.

11. Sustainable/Safe

Sustainability and safety both involve caring for people and the planet. Ideally, they should go hand in hand. But in datacentres, efforts to look sustainable can sometimes create safety risks, and steps to protect communities can sometimes make sustainability harder to achieve.

The clearest example is water. Water-based cooling (evaporative cooling towers) is more energy-efficient than air-based cooling. It lowers the PUE, improving the sustainability metric. But it consumes enormous quantities of freshwater in regions that may be water-stressed, creating direct safety risks for local communities and ecosystems. Northern Virginia data centres collectively consumed close to 2 billion gallons of water in 2023, a 63% increase from 2019. Loudoun County, with approximately 200 operational data centres, used around 900 million gallons of water in 2023, leading Loudoun Water to rely heavily on potable water for data centres rather than reclaimed water (Yañez-Barnuevo, 2026).

The main sustainability metric, PUE, encourages cooling methods that put the most strain on local water supplies. Air cooling is safer for water but uses more energy. Liquid immersion cooling is efficient and uses little water, but it is much more expensive. When the market focuses only on PUE, it often chooses the option that looks best on paper but poses the greatest water safety risks to local communities.

In the UK, this tension interacts with the Renewable Fiction Loop. The PUE metric incentivises water-based cooling, which scores well in sustainability reporting metrics but causes water stress in host communities. As set out under Tension 4, the great majority of proposed UK data centres are sited in water-stressed areas, and Thames Water's consumption and debt position compound the safety risk.

The WRc report for MOSL confirms that water use in UK datacentres is concentrated almost exclusively in cooling systems, with humidity control having only a marginal impact. The technology employed is decisive. In England, most data centres use cooling systems with little or no water consumption, or hybrid systems that use water only under high-temperature conditions. Approximately 5 per cent of facilities may be using hybrid cooling systems that significantly increase consumption during summer, precisely when water stress is highest and when the safety risk to local supply is greatest. The average Water Usage Effectiveness at the European scale stands at 0.58 litres per kWh, although differences between member states are considerable (WRc, 2026; Smart Water Magazine, 2026).

Planning trends make the UK situation even more complex. Google's Waltham Cross campus in Hertfordshire was approved only after agreeing to zero evaporative loss, suggesting that future projects may need to demonstrate water neutrality (Mordor Intelligence, 2026).

However, this was a site-specific condition, not a regulatory requirement. There is no national standard for datacenter water use, no mandatory Water Usage Effectiveness reporting, and no way to assess total water demand across multiple datacenter projects in the same water-stressed area (WRc, 2026). Individual planning approvals may seem manageable, but together they can exceed the region's water supply. This is the same "accumulation of small changes, each of little consequence" dynamic that the hexalemma's temporal dimension identifies as a path to tipping points.

The opposite problem also happens. When communities ask for additional safety measures, such as using air cooling, limiting water use, or requiring companies to share water data, the industry says this makes it harder to meet sustainability goals because it requires more energy. The focus on caring for the planet (using less energy) is often put ahead of caring for people (protecting water), because of how the rules are set up.

12. Sustainable/Ethical

The way large tech companies make sustainability promises raises important ethical questions about who gets to define what is sustainable and whose interests are actually being served by the current rules.

The PPA and REC model of renewable procurement allows companies to claim renewable energy use while drawing power from grids that remain substantially fossil fuelled. This serves the sustainability reporting needs of investors, ESG rating agencies, and regulators. It does not serve the sustainability interests of communities bearing the environmental costs of fossil-fuelled generation, the air quality impacts, the health consequences, and the climate emissions that are generated to power facilities claiming renewable status.

The ethical dimension concerns legitimacy- who has the right to define sustainability in this context? The current answer is effectively the companies themselves, through voluntary commitments and corporate reporting, with third-party verification by rating agencies whose methodologies may or may not reflect the concerns of affected communities. California's veto of water transparency legislation, following opposition from industry groups, underscores how data centre operators are actively resisting the accountability frameworks that would allow sustainability claims to be externally verified (Edwards, 2025).

The criticism that economic efficiency often comes at the cost of environmental and social values is especially true for an industry that claims to be sustainable while its real impact keeps growing. In 2024, global data centres used approximately 415 TWh of electricity, around 1.5 per cent of the world's total, and the IEA projects this could more than double to 945 TWh by 2030 (IEA, 2025). The tension between sustainability and ethics partly comes down to whether current accounting rules are meant to help the environment or just look good to investors.

The Climate Change Committee and the Environmental Audit Committee are both expected to scrutinise the practice of purchasing unbundled renewable energy certificates, which effectively rebadges existing clean power as additional capacity. The House of Commons Library briefing of 27 May 2026 notes that major datacentre operators have committed to achieving voluntary sustainability targets by 2030 and support renewables projects through PPAs, while adding that "some studies have questioned whether PPAs contribute to decarbonisation." The framing is instructive, the sustainability claims are presented as voluntary corporate commitments, and the evidence-based questioning of their effectiveness is noted but not resolved. There is no regulatory mechanism in the UK for making that resolution binding (Clark, 2026).

The cost differential facing UK consumers, among the highest in the developed world, is discussed under Tension 15, but it is the cost differential that is itself an ethical issue, for example, UK consumers are paying some of the highest electricity prices in the developed world, and the network investment required to connect datacentre demand at scale will further increase those prices under RIIO-3. The sustainability commitments of the technology companies, expressed through PPAs that do not add new generation and RECs that rebadge existing supply, do not reduce the cost burden on UK consumers. They reduce the reporting burden on the companies. The efficiency of the sustainability accounting mechanism and the ethical distribution of its costs point in opposite directions.

This is the legitimacy problem in the Renewable Fiction Loop. Current accounting rules allow companies to report lower emissions even as the grid’s actual carbon intensity rises. This is not just a technical issue; it is a system that helps companies and investors look good in sustainability reports while hiding the real costs borne by communities and the grid.

13. Efficient/Reliable

Efficiency and reliability might seem to go together, efficient systems should be reliable, and reliable systems should not waste resources. But in datacentres, the push for efficiency can make things reliable for one operator while making the grid less reliable for everyone else. On-site gas projects announced in 2025 make individual data centres more independent of grid conditions. However, this also removes large loads from the pool of demand that grid operators can manage flexibly. A datacentre powered by its own gas turbines is not available for demand response, cannot be called on to reduce load during a supply emergency, and does not support grid-balancing services that all other consumers depend on. The efficiency gain for the individual operator comes at the cost of reduced system reliability.

The same issue comes up with how data centres connect to the grid. Some data centres skip the waiting line by generating their own power, ensuring they receive reliable electricity sooner. But the queue is there to make sure the grid can handle new demands safely. When datacentres bypass it, they improve their own efficiency but make it harder to plan for everyone else's reliability.

This tension is structurally related to the Demand Spiral. As individual data centres optimise for efficiency by bypassing the grid connection queue through behind-the-meter generation, they remove large loads from the pool of flexible demand that grid operators need to maintain system stability. In the UK, the same dynamic drove the connection-queue crisis and NESO's December 2025 "first ready, first connected" reform discussed under Tension 7, each firm's decision to bypass the queue or build behind-the-meter was individually efficient, but the aggregate overwhelmed grid planning.

The UK government's response illustrates the tension rather than resolving it. The AI Growth Zones programme prioritises data centre connections as "strategically important projects." Ofgem's phased approach to demand connection reform starts with strengthening financial commitments and weeding out non-viable data centres from the queue. But the viable ones, the ones that survive the reform process, are precisely the large, firm-power-demanding hyperscale facilities whose connection most affects grid reliability for other users. The reform makes the queue more efficient. It does not resolve the underlying tension between the efficiency of individual connections and the reliability of the collective system.

14. Efficient/Safe

Running datacentres more efficiently can sometimes make them less safe in ways that standard measurements do not capture. Fitting more servers into smaller spaces saves money but generates a lot of heat, increasing the risk of fires, cooling failures, and chemical hazards, since many cooling fluids are hazardous. These risks need careful management. On a bigger scale, making data centres more efficient by automating and managing them remotely means there are fewer people on site to spot and fix new safety problems. While automation saves on staffing costs and boosts efficiency, it also reduces the human awareness that can catch issues machines might miss. At the same time, harmonic distortions can build up and spread into the local power system, posing risks to industrial and residential equipment, including malfunctions, accelerated wear, overheating, and electrical fires. The efficiency of high-density computing creates power-quality problems that impose safety costs on neighbouring properties, but these costs do not appear in the data centre operator's efficiency metrics (Chen et al., 2025).

This issue is part of the Demand Spiral at the facility level. As AI workloads increase, each server rack uses much more power- 40 to 70 kW compared to 8 to 12 kW for regular loads. This makes it harder to manage heat and power quality. Generative AI is pushing these limits even further. While the efficiency gains are real, the safety risks are often borne by nearby communities rather than factored into efficiency measures.

In the UK, the NPPF consultation published on 16 December 2025 proposes classifying data centres as "essential infrastructure" in the flood risk vulnerability classification (Ministry of Housing, Communities and Local Government, 2025; Tucker, 2025). This is a designation that prioritises efficient siting, placing datacentres where land is available and grid connections are feasible, over safety-of-siting considerations. Flood risk classification determines which locations are available for development. Classifying datacentres as "essential infrastructure" rather than "highly vulnerable" makes flood-prone sites in Flood Zone 3 available that would otherwise be restricted, on the logic that the economic importance of the infrastructure justifies the risk (O'Neill-Williams, 2025). The safety question, whether the communities adjacent to those flood-prone sites bear additional risk from the presence of a facility containing diesel backup generators, chemical coolants, and high-voltage power equipment in a flood zone, is not addressed by the classification. The efficiency of siting and the safety of siting are being resolved in favour of efficiency by a planning designation that treats the two as unrelated.

15. Efficient/Ethical

This is the allocation tension at the centre of the hexalemma's third column. Scarce resources are being allocated primarily by market efficiency mechanisms, with consequences for fairness and justice that those mechanisms do not price in.

The efficiency of the datacentre industry, measured by PUE, cost per compute unit, speed of deployment, has been achieved partly through the externalisation of costs. Grid upgrade costs are passed to ratepayers. Water costs are not charged at marginal environmental cost. Planning approvals are obtained under frameworks designed for smaller industrial loads. The efficiency gains accrue to technology companies and their shareholders; the externalised costs are distributed across communities, ratepayers, and future generations.

Average electricity prices in the United States rose to 19 cents per kWh by the end of 2025, about 27 per cent more than in 2019, and residential prices are projected to increase by up to 40 per cent by 2030 compared with 2025, while utility shut-offs rose to 3.5 million in 2024 (Yañez-Barnuevo, 2026).

In the UK, the efficiency/ethical tension has a distinctive shape determined by the structure of energy pricing and the relative cost position of the UK market. A 500 MW annual electricity bill in the UK sits at around £900 million, compared with £700 million in France, £438 million in Spain, and £219 million in the United States (Clark, 2026). UK electricity prices are among the highest in the developed world. The Ofgem price cap for Q3 2026 rises 13 per cent to £1,862 for a typical dual-fuel household (Ofgem, 2026c). Network charges constitute 20 to 25 per cent of electricity bills and are rising under RIIO-3 as the transmission operator revenue requirement increases 149 per cent compared to 2025/26 levels (Powerhub Solutions, 2025).

The fuel poverty data sharpens the ethical dimension. In England, 2.36 million households are in fuel poverty under the LILEE metric, with over 30 per cent who spend more than 10 per cent of their income on energy (GOV.UK, 2026b). Across the UK, the University of York found 12 million households (43 per cent) struggling with energy costs, with nearly 5 million in deep fuel poverty (End Fuel Poverty Coalition, 2025). The government's own Fuel Poverty Strategy acknowledges the affordability crisis and commits to lifting a million more households out of fuel poverty by 2030 (GOV.UK, 2026c). The network cost increases required to connect 50 GW of datacentre demand to the grid are socialised across this population through the same TNUoS charges that fund all transmission investment. The efficiency gains of AI compute, measured in cost per token, cost per training run, cost per inference query, are captured by a small number of technology companies. The infrastructure costs are distributed across a population in which nearly a third of households already spend a disproportionate share of their income on energy.

The UK government's AI Growth Zones programme compounds this. The zones in Oxfordshire, North Lincolnshire, South and North Wales, and the North East receive streamlined planning, prioritised grid connections, and potential NSIP designation (Department for Science, Innovation and Technology, 2025). The communities in these zones will bear the local costs, grid stress, water consumption, planning disruption, construction impact, visual and noise intrusion. The economic value of the AI compute hosted in those zones will be captured by technology companies headquartered elsewhere, by investors globally, and by consumers of AI services worldwide. techUK has itself asked whether AI Growth Zones will deliver local value or just national headlines, a question that directly frames the efficient/ethical tension in UK policy terms (Vieau and Dorociak, 2025).

The framing of this tension within the hexalemma notes that the neoliberal capitalist agenda contains reinforcing mechanisms, success to the successful, that arise from the pursuit of efficiency. The technology companies that benefit most from efficient AI compute are the same companies best placed to influence the regulatory frameworks governing how efficiency gains are measured and how costs are allocated. The efficiency/ethical tension is, in this context, not merely a technical trade-off but a political economy question about who has power over the terms on which resources are allocated.

In the UK, this political economy dimension is visible in the planning reforms themselves. Datacentres were designated CNI in September 2024 after lobbying from the technology sector. The move toward NSIP designation for large datacentre projects, and the proposal in the Planning and Infrastructure Bill 2025 to limit judicial review of infrastructure decisions, progressively reduces the institutional channels through which the ethical distribution of costs and benefits can be contested (Robinson, 2026a). A government civil servant warned at the time of the CNI designation that it would allow authorities to override local opposition. The efficiency of the planning process is being improved. The ethical legitimacy of the outcomes it produces is not being assessed (Adshead, 2026).

Conclusion and discussion

Economic efficiency pursued as a primary system value systematically generates social and ecological costs that remain invisible to the efficiency metric. The reason efficiency holds this position is that it carries a property the other five values lack. Optimising it feeds the very demand it was meant to contain, which is the Jevons paradox set out by Alcott (2005), where a fall in the cost of using a resource raises its profitability, draws in new capital, and opens new uses until absolute consumption climbs. AI compute fits these conditions almost exactly, and this single mechanism is what makes efficiency hazardous as a governing metric while it remains useful as an engineering one.

The central finding of this paper is that under current UK governance this pattern is total. As Table 1 records, all fifteen hexalemma tensions resolve as dualisms, and every one currently runs in efficiency's favour. Values that could in principle reinforce one another are structurally forced apart by the way the datacentre sector is financed, sited, and regulated. These are one problem seen from fifteen angles, and the integrated causal loop diagram locates its mechanism. The Demand Spiral, the Renewable Fiction Loop, and the Resilience Erosion Loop are each self-reinforcing, and the classification table maps each tension back to the loop that holds it in place. Tensions 1, 13, and 14 sit on the Demand Spiral, tensions 3 and 12 on the Renewable Fiction Loop, tensions 5, 8, and 9 on the Resilience Erosion Loop, and tension 15, the core allocation tension, draws on all three at once. The diagram shows why the dualisms stay stable. The three loops share variables, with concentrated build-out feeding both the Demand Spiral and the Renewable Fiction Loop, and grid stress linking the Demand Spiral to the Resilience Erosion Loop, so a tighter PUE threshold, a voluntary PPA, or a faster connection queue feeds the same loop it targets and leaves the feedback structure intact.

The Renewable Fiction Loop shows this dynamic in its sharpest form. PPA and REC procurement allow reported emissions to fall while actual grid carbon intensity stays unchanged, because no additionality is required, so the system scores well on reporting efficiency and poorly on genuine sustainability. The Resilience Erosion Loop carries the harm outward. Grid stress from geographic concentration drives infrastructure costs that are socialised across all ratepayers through the same TNUoS charges that fund all transmission, including the populations and regions that receive none of the economic benefit. The resilience problem and the justice problem are the same structural problem seen from two axes, and this is the Efficient/Ethical allocation tension that sits at the centre of the table.

The reason this resolution persists is temporal. As Koutsandreas and Keppo (2026) demonstrate, short-horizon decision-makers manage accumulating stress through incremental efficiency measures, and the soft constraints that dominate UK policy, the PUE thresholds, reporting standards, and differentiated charges, preserve the room to defer structural investment. Each deferral looks rational at every step because the future is heavily discounted, and the system locks into a more expensive and less resilient trajectory. Hard constraints behave differently, creating cross-year dependencies that remove the option to postpone. The phrase from Hemingway's The Sun Also Rises (1926), gradually then suddenly, captures the dynamic. Stress builds connection by connection until it crosses a structural threshold the system cannot recover from without far costlier intervention.

Framing the problem through the hexalemma gives this critique analytical structure and makes it addressable through the scenario analysis. The question becomes what a governance framework would look like that scored well on efficient and ethical at the same time. The scenarios answer it. Unmanaged growth leaves all three loops running and reproduces exactly the pattern the table records, efficient high in the short run while the other five axes decline. Efficiency-led reform fails in an instructive way, because mandating efficiency feeds the Demand Spiral it targets, which is the Jevons paradox operating at the level of governance and at the level of the engine (Alcott, 2005). Structural transition is the one scenario that moves a tension from dualism toward duality, and Table 1 marks the four where that movement is open, tensions 1, 2, and 7 carrying explicit duality potential and tension 10 standing as the hardest to shift. Structural transition reaches those openings through hard constraints, by making the trade-offs explicit and choosing among them consciously. The hexalemma's claim holds that the tensions are real and will persist. What the framework exposes is that they are currently resolved by default, in efficiency's favour, with the choice made by the structure and owned by no one accountable for it.

This is where the analysis meets the Schumacher Institute's Economic Thinking theme. The datacentre sector's growth model is extractive in the precise sense the institute critiques, with efficiency gains and value creation concentrated in a small number of firms while the costs of grid infrastructure, water depletion, carbon emissions, and community disruption fall on populations and ecosystems that hold a meaningful stake in the value being created. The move from extraction to maintenance and care carries concrete governance content in this context. It is the substance of the structural transition scenario, built from demand-side constraints, cost-internalisation requirements, community benefit obligations, and integrated planning that would carry the system out of the efficient corner of the hexalemma toward a higher score across all six axes. Whether UK governance holds the institutional capacity for that shift, while it continues to optimise the one value it knows how to measure and socialise the rest, is the central energy-governance question this analysis leaves open. The hexalemma leaves the answer to politics. Its work is to make the question impossible to avoid.

Reference list

Adams, C., Campbell, D., Dearden, N., Downes, R., Graham, S. and O’Leary, A. (2026). Letter to the Secretary of State for Science, Innovation and Technology re National Policy Statement on data centres,. [online] Available at: https://www.foxglove.org.uk/wp-content/uploads/2026/03/2026_02_05_PUB-Letter-to-Technology-Secretary-re-data-centres-NPS.pdf

Adshead, A. (2026). Enormous AI Growth Zone Datacentre Gets Planning Approval. [online] ComputerWeekly.com. Available at: https://www.computerweekly.com/news/366639995/Enormous-AI-growth-zone-datacentre-gets-planning-approval

Alcott, B. (2005). Jevons’ paradox. Ecological Economics, [online] 54(1), pp.9–21. doi:https://doi.org/10.1016/j.ecolecon.2005.03.020.

Arnold, R.D. and Wade, J.P. (2015). A definition of systems thinking: A systems approach. Procedia Computer Science, 44(44), pp.669–678. doi:https://doi.org/10.1016/j.procs.2015.03.050.

ARUP (2017). Five Minute guide: Energy Trilemma Five Minute Guide Energy Trilemma. [online] Available at: https://www.arup.com/globalassets/downloads/insights/five-minute-guide-energy-trilemma.pdf.

Bhuyan, S. and Broom, K.D. (2025). Systems Thinking in Healthcare. Springer Publishing.

Bocca, R. and Ashraf, M. (2025). What we can learn from the 2025 Energy Transition Index. [online] World Economic Forum. Available at: https://www.weforum.org/stories/2025/06/energy-transition-resilience-five-actions-to-increase/.
Bove, T. (2026). Data Centers Are Dealing Hidden Damage to Environmental and Public Health—costing the Economy $25 Billion Every Year. [online] Fortune. Available at: https://fortune.com/2026/04/21/data-centers-environmental-health-costs-25-billion/.
Chen, D., Zhou, Z., Cai, Y., Qin, J., Katchova, A. and Chen, L. (2026). Concentrated Siting of AI Data Centers Drives Regional power-system Stress under Rising Global Compute Demand. [online] Available at: https://arxiv.org/pdf/2604.06198.
Chen, X., Wang, X., Colacelli, A., Lee, M. and Xie, L. (2025). Electricity Demand and Grid Impacts of AI Data Centers: Challenges and Prospects. [online] Arxiv.org. Available at: https://arxiv.org/html/2509.07218v3.
Chen, Y. and Zheng, X. (2026). To Defer or To Shift? The Role of AI Data Center Flexibility on Grid Interconnection. [online] Available at: https://arxiv.org/pdf/2604.05376.
Clancy, H. (2025). 7 companies helping Amazon, Google, Meta and Microsoft go nuclear. [online] Trellis. Available at: https://trellis.net/article/amazon-google-meta-and-microsoft-go-nuclear/.
Clark, A. (2026). Data centres: Planning policy, sustainability, and Resilience Summary. [online] House of Commons Library. Available at: https://researchbriefings.files.parliament.uk/documents/CBP-10315/CBP-10315.pdf.
Climate Central (2024). Weather-related Power Outages Rising | Climate Central. [online] www.climatecentral.org. Available at: https://www.climatecentral.org/climate-matters/weather-related-power-outages-rising.
Crosley, B. (2026). Power Purchase Agreements (PPAs) for AI Data Centers: Renewable Energy Strategies. [online] Introl. Available at: https://introl.com/blog/power-purchase-agreements-ai-data-centers-renewable-energy-strategies.
Deaton, M. and MacDonald, R. (2025). Systems Thinking: a New Way to Tackle Problems. [online] Jmu.edu. Available at: https://pressbooks.lib.jmu.edu/sdlearningguide/chapter/systems-thinking-a-new-way-to-tackle-problems/.
Department for Science, Innovation and Technology (2025). Delivering AI Growth Zones. [online] GOV.UK. Available at: https://www.gov.uk/government/publications/delivering-ai-growth-zones/delivering-ai-growth-zones.

Edwards, E. (2025). Data Center Water Secrecy Hurts Communities (and the Industry Itself) - Bluefield Research. [online] Bluefield Research. Available at: https://www.bluefieldresearch.com/data-center-water-secrecy-hurts-communities-and-the-industry-itself/.
End Fuel Poverty Coalition (2025). Fuel Poverty Statistics Show 12m UK Households Struggling with Energy Costs. [online] End Fuel Poverty Coalition. Available at: https://www.endfuelpoverty.org.uk/fuel-poverty-statistics-show-12-million-households-struggling/.
Energy Institute (2024). New Energy World magazine. [online] Energy Institute. Available at: https://knowledge.energyinst.org/new-energy-world/article?id=139072
Flame Energy (2025). RIIO-3: What It Means for UK Businesses and Energy Costs - Flame Energy. [online] Flame Energy. Available at: https://flameenergy.co.uk/news/riio-3-what-it-means-for-uk-businesses-and-energy-costs/
Foxglove (2025). Drinking Water and Data centres: 10 Billion Litres of H2O, Floods of Sewage and Our Legal Challenge to the Biggest Centre in Planning Permission in the UK. [online] Foxglove. Available at: https://www.foxglove.org.uk/2025/07/28/water-data-centres-floods-sewage/.
Francis, N. (2026). UK Nuclear Build: From Mega-Projects to Modular Reactors. [online] Constructiondigital.com. Available at: https://constructiondigital.com/news/uk-nuclear-build-from-mega-projects-to-modular-reactors
Gaster, R. (2025). The United States Needs Data Centers, and Data Centers Need Energy, but That Is Not Necessarily a Problem. [online] Itif.org. Available at: https://itif.org/publications/2025/11/24/united-states-needs-data-centers-data-centers-need-energy-but-that-is-not-necessarily-a-problem/.
Global Action Plan (2026). Not a drop to drink: How Britain’s data centre surge threatens water security. Infectious Diseases in Clinical Practice, 3(5), pp.381–382. doi:https://doi.org/10.1097/00019048-199409000-00018.
GOV.UK (2025). AI Opportunities Action Plan. [online] GOV.UK. Available at: https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan.
GOV.UK (2026a). About This Release Information on Energy production, trade, and Consumption in the UK for Total Energy and by Specific Fuels. [online] Available at: https://assets.publishing.service.gov.uk/media/69cd1451b5210036050bc625/Energy_Trends_March_2026.pdf.

GOV.UK (2026b). Annual Fuel Poverty Statistics in England, 2026 (2025 data). [online] Available at: https://assets.publishing.service.gov.uk/media/69c3af123ed0546101e0dc3e/Main_Report__2026_Fuel_Poverty_Statistics_Publication_.pdf.
GOV.UK (2026c). Review of the Fuel Poverty Strategy: Summary of responses and government response (accessible webpage). [online] GOV.UK. Available at: https://www.gov.uk/government/consultations/review-of-the-fuel-poverty-strategy/outcome/review-of-the-fuel-poverty-strategy-summary-of-responses-and-government-response-accessible-webpage.
Howland, E. (2025). Data Centers ‘Primary Reason’ for High PJM Capacity prices: Market Monitor. [online] Utility Dive. Available at: https://www.utilitydive.com/news/data-centers-pjm-capacity-auction-market-monitor/801780/.
IDCA (2026). Global Data Center Report. [online] Idc-a.org. Available at: https://www.idc-a.org/insights/Sk2PFZCost98Eqhk2FUA
IEA (2025). Energy and AI – Analysis - IEA. [online] IEA. Available at: https://www.iea.org/reports/energy-and-ai.
Jacobs, M. (2025). Data Centers Are Already Increasing Your Energy Bills. We Have the Receipts. [online] The Equation. Available at: https://blog.ucs.org/mike-jacobs/data-centers-are-already-increasing-your-energy-bills/.
Jealous, B. (2025). Demand Data Centers Keep Americans Safe. [online] Chicago Sun-Times. Available at: https://chicago.suntimes.com/columnists/2025/12/16/data-centers-health-clean-water-oregon-ohio-ben-jealous
Koutsandreas, D. and Keppo, I. (2024). How Can Energy System Models Align with the Real-World Tragedy of the Horizon? the Impact of Decision-Makers’ Short-Term Biases and Temporally-Inconsistent Preferences. doi:https://doi.org/10.2139/ssrn.4997635.
Laimon, M., Yusaf, T., Mai, T., Goh, S. and Alrefae, W. (2022). A Systems Thinking Approach to Address Sustainability Challenges to the Energy Sector. International Journal of Thermofluids, [online] 15, p.100161. doi:https://doi.org/10.1016/j.ijft.2022.100161.
Lakhanpal, V. (2026). Data Center Water Use - MOST Policy Initiative. [online] Mostpolicyinitiative.org. Available at: https://mostpolicyinitiative.org/science-note/data-center-water-use/.
Levin, M. (2026). Rep. Mike Levin Introduces New Bill to Stop Data Centers from Driving up Electricity Prices for Consumers | U.S. Congressman Mike Levin. [online] House.gov.

Available at: https://levin.house.gov/media/press-releases/rep-mike-levin-introduces-new-bill-to-stop-data-centers-from-driving-up-electricity-prices-for-consumers.
Leyla Acaroglu (2019). Systems Thinking 101. [online] Unschools. Available at: https://unschools.co/systems-thinking-101/.
Li, L., Lu, B., Hechtman, J. and Cowan, D. (2026). Roadmap: The AI data center stack. [online] Bessemer Venture Partners. Available at: https://www.bvp.com/atlas/roadmap-the-ai-data-center-stack.
Lord, C. and Widdicks, K. (2026). Data Centres and sustainability: Why Power and Water Are Only Part of the Story | UK Centre for Ecology & Hydrology. [online] Ceh.ac.uk. Available at: https://www.ceh.ac.uk/news-and-media/blogs/data-centres-and-sustainability-why-power-and-water-are-only-part-story
Loughran, J. (2026). Hinkley Point C Faces Further Delays as Costs Continue to Mount. [online] Engineering and Technology Magazine. Available at: https://eandt.theiet.org/2026/02/23/hinkley-point-c-faces-further-delays-costs-continue-mount.
Lyu, J. and Tang, S. (2025). Power Hungry Data Centers Are Driving Green Energy Demand | BloombergNEF. [online] BloombergNEF. Available at: https://about.bnef.com/insights/clean-energy/power-hungry-data-centers-are-driving-green-energy-demand/.
Meadows, D. (2008). Thinking in Systems. [online] earthscan. Available at: https://research.fit.edu/media/site-specific/researchfitedu/coast-climate-adaptation-library/climate-communications/psychology-amp-behavior/Meadows-2008.-Thinking-in-Systems.pdf.
Ministry of Housing, Communities and Local Government (2025a). National Planning Policy Framework Plan-making and national decision-making policies. [online] Available at: https://assets.publishing.service.gov.uk/media/697b71c52ff8d10a830d5d4a/Draft_NPPF_December_2025.pdf.
Ministry of Housing, Communities and Local Government (2025b). Proposed reforms to the National Planning Policy Framework and other changes to the planning system. [online] Available at: https://assets.publishing.service.gov.uk/media/697b6bc6aacd0dc9777b4fd2/December_2025_NPPF_Consultation.pdf.
Monitoring Analytics (2026). Analysis of the 2026/2027 RPM Base Residual Auction Part A. [online] Available at:

https://www.monitoringanalytics.com/reports/Reports/2025/IMM_Analysis_of_the_20262027_RPM_Base_Residual_Auction_Part_A_20251001.pdf.

Mordor Intelligence (2025). United Kingdom Water Consumption Market Size & Share Analysis - Industry Research Report - Growth Trends. [online] Mordorintelligence.com. Available at: https://www.mordorintelligence.com/industry-reports/united-kingdom-data-center-water-consumption-market.

Moss, S. (2026). Thames Water May Restrict London Data Centers' Water Use. [online] Datacenterdynamics.com. Available at: https://www.datacenterdynamics.com/en/news/thames-water-may-restrict-london-data-centers-water-use/

Murray, J. (2026). Cleaner in, Useful out: the Maturing Sustainability Stack for UK Data Centres. [online] Bird & Bird. Available at: https://www.twobirds.com/en/insights/2026/uk/cleaner-in,-useful-out-the-maturing-sustainability-stack-for-uk-data-centres.

National Grid (2026). Value for money. [online] RIIO-T3. Available at: https://www.riiot3.nationalgrid.com/value-money

NESO (2025). NESO implements electricity grid connection reforms to unlock investment in Great Britain | National Energy System Operator. [online] National Energy System Operator (NESO). Available at: https://www.neso.energy/neso-implements-electricity-grid-connection-reforms-unlock-investment-great-britain.

Norton Rose Fulbright (2025a). 'First Ready, First Connected': Reform to GB Electricity Grid Connections. [online] Nortonrosefulbright.com. Available at: https://www.nortonrosefulbright.com/en/knowledge/publications/5567f27c/first-ready-first-connected-reform-to-gb-electricity-grid-connections

Norton Rose Fulbright (2025b). TMO4+ Connection Reform Proposals Receive Stamp of Approval. [online] Nortonrosefulbright.com. Available at: https://www.nortonrosefulbright.com/en/knowledge/publications/0101e3b9/tmo4-connection-reform-proposals-receive-stamp-of-approval

O'Neill-Williams, B. (2025). What the Critical Infrastructure Label Means for Data Centre Flood Risk - Blog - Cundall. [online] Cundall. Available at: https://www.cundall.com/ideas/blog/what-the-critical-infrastructure-label-means-for-data-centre-flood-risk.

Ofgem (2025a). RIIO-3 Final Determinations. [online] Ofgem. Available at: https://www.ofgem.gov.uk/energy-policy-and-regulation/policy-and-regulatory-programmes/network-price-controls-2026-2031-riio-3

Ofgem (2025b). Changes to energy price cap between 1 January and 31 March 2026. [online] Ofgem. Available at: https://www.ofgem.gov.uk/news/changes-energy-price-cap-between-1-january-and-31-march-2026.

Ofgem (2025c). Ofgem unlocks £28 billion investment to maintain a safe, secure and resilient energy grid and to upgrade and expand capacity to meet growing demands. [online] Ofgem. Available at: https://www.ofgem.gov.uk/press-release/ofgem-unlocks-ps28-billion-investment-maintain-safe-secure-and-resilient-energy-grid-and-upgrade-and-expand-capacity-meet-growing-demands.

Ofgem (2026a). Call for input: demand connections reform. [online] Ofgem. Available at: https://www.ofgem.gov.uk/sites/default/files/2026-02/2026-02-12-Demand-Connections-Call-for-Input.pdf.

Ofgem (2026b). Changes to energy price cap between 1 April and 30 June 2026. [online] Ofgem. Available at: https://www.ofgem.gov.uk/news/changes-energy-price-cap-between-1-april-and-30-june-2026.

Ofgem (2026c). Changes to energy price cap between 1 July and 30 September 2026. [online] Ofgem. Available at: https://www.ofgem.gov.uk/news/changes-energy-price-cap-between-1-july-and-30-september-2026.

Powerhub Solutions (2025). RIIO-ET3 & the 149% Rise in UK Transmission: Costs What It Means For Businesses - Powerhub Solutions. [online] Powerhub Solutions. Available at: https://powerhub.solutions/2025/12/10/riio-et3-the-149-rise-in-uk-transmission-costs-what-it-means-for-businesses/

Privette, A.P., Barros, A. and Cai, X. (2026). Data Centers Water Footprint: The Need for More Transparency. AGU Advances, 7(2). doi:https://doi.org/10.1029/2025av002140.

Robinson, D. (2025). UK prepared to loosen planning rules for massive datacenters. [online] theregister. Available at: https://www.theregister.com/on-prem/2025/01/13/uk-prepared-to-loosen-planning-rules-for-massive-datacenters/833214

Robinson, D. (2026a). Datacenter Builds Could Be Shielded from Judicial Review in UK Planning Reforms. [online] theregister. Available at: https://www.theregister.com/systems/2026/05/20/datacenter-builds-could-be-shielded-from-judicial-review-in-uk-planning-reforms/5243354

Robinson, D. (2026b). UK govt caves on datacentre approval after legal challenge. [online] theregister. Available at: https://www.theregister.com/on-prem/2026/01/23/uk-govt-caves-on-datacenter-approval-after-legal-challenge/5201072

Skidmore, Z. (2025). UK’s National Grid Reports Surge in Data Center Connection Requests. [online] Datacenterdynamics.com. Available at: https://www.datacenterdynamics.com/en/news/uks-national-grid-reports-surge-in-data-center-connection-requests/.

Smart Water Magazine (2026). The water challenge of data centres in England: concentration and peak demand. [online] Smart Water Magazine. Available at: https://smartwatermagazine.com/news/smart-water-magazine/water-challenge-data-centres-england-concentration-and-peak-demand

Starman, D. (2025). 2026 TNUoS Costs Set to increase, but by Less than Feared. [online] Drax.com. Available at: https://energy.drax.com/intelligence/2026-tnuos-to-increase-riio-et3-determinations/.

Sustainability Atlas (2026). Data center energy consumption: tracking global power demand, water use, and carbon intensity. [online] Sustainability Atlas. Available at: https://sustainableatlas.org/post/data-story-data-center-energy-consumption-water-carbon-trends-1624

Tech.UK (2025). News: Ofgem Greenlights £28.1bn of Grid Investment through RIIO-3. [online] Techuk.org. Available at: https://www.techuk.org/resource/ofgem-greenlights-28-1bn-of-grid-investment-through-riio-3.html

Tucker, M. (2025). Data Centres & AI Growth Zones in planning: Change on the Horizon in 2026. [online] Burges Salmon. Available at: https://www.burges-salmon.com/articles/102lxwu/data-centres-ai-growth-zones-in-planning-change-on-the-horizon-in-2026/.

UK Parliament (2026a). MPs to Examine Data Centres Climate Impact as Miliband Says Future Demand ‘inherently Uncertain’ - Committees - UK Parliament. [online] Parliament.uk. Available at: https://committees.parliament.uk/committee/62/environmental-audit-committee/news/212132/mps-to-examine-data-centres-climate-impact-as-miliband-says-future-demand-inherently-uncertain/.

UK Parliament (2026b). Risks and Opportunities to the Sustainability of Data Centres in the UK - Committees - UK Parliament. [online] Parliament.uk. Available at: https://committees.parliament.uk/work/9651/risks-and-opportunities-to-the-sustainability-of-data-centres-in-the-uk/

Vasudevan, H. (2025). Data Centers Are Breaking the Old grid. Let AI Build the New one. [online] Utility Dive. Available at: https://www.utilitydive.com/news/ai-utilities-reliability-cost/805224/

Vieau, F. and Dorociak, V. (2025). Will AI Growth Zones Deliver Local Value or Just National headlines? [online] Techuk.org. Available at: https://www.techuk.org/resource/will-ai-growth-zones-deliver-local-value-or-just-national-headlines.html.

Waltz, E. (2024). Big Tech Backs Small Nuclear, Google and Amazon Invest in Small Modular Reactors to Power Data Centers. [online] IEEE Spectrum. Available at: https://spectrum.ieee.org/nuclear-powered-data-center.

Water Magazine (2026). 84% of Proposed UK Data Centres to Be Built in Water Stressed Areas - Water Magazine. [online] Water Magazine. Available at: https://www.watermagazine.co.uk/2026/04/27/84-of-proposed-uk-data-centres-to-be-built-in-water-stressed-areas/.

Weaver, T. (2026). UK and US data centres now consume around 6% of national electricity. [online] Available at: https://eandt.theiet.org/2026/05/15/uk-and-us-data-centres-now-consume-around-6-national-electricity.

World Energy Council (2019). WEC Energy Trilemma Index Tool. [online] Worldenergy.org. Available at: https://trilemma.worldenergy.org/.

World Energy Council (2021). World Energy Trilemma Index. [online] Available at: https://energyrights.info/sites/default/files/artifacts/media/pdf/world_energy_trilemma_index_world_energy_council.pdf

WRc (2026). Water Efficient Data Centres. [online] WRc. Available at: https://www.wrcgroup.com/uploads/Water%20Efficient/Water%20Efficient%20Data%20Centres.pdf.

Yagüe, A. and Ashtine, M. (2025). The Renewable Route for Data Centre Expansion. [online] The Carbon Trust. Available at: https://www.carbontrust.com/news-and-insights/insights/the-renewable-route-for-data-centre-expansion.

Yañez-Barnuevo, M. (2026). Data Center Power Demands Are Contributing to Higher Energy Bills | Article | EESI. [online] Eesi.org. Available at: https://www.eesi.org/articles/view/data-center-power-demands-are-contributing-to-higher-energy-bills.