Yottar

REPORT

INVISIBLE CONSTRAINTS

The Economic Impact of Limited Grid Capacity Visibility



November 2025

Yottar

CONTENTS

EXECUTIVE SUMMARY 3

CONTEXT 4

METHODOLOGY 5

COSTS 5

- DELAYED AND DETERRED INVESTMENT 6
- NETWORK REINFORCEMENT INEFFICIENCY 7
- SYSTEM BALANCING AND CONSTRAINT COSTS 9
- ADDITIONAL DEVELOPER COSTS 10
- TOTAL COSTS 11

ENABLING GRID CAPACITY VISIBILITY 12

RECOMMENDATIONS 13

CONCLUSIONS 14

EXECUTIVE SUMMARY

The UK's ability to capture growth from the AI boom and the electrification supercycle is limited by access to electrical power. Beyond the physical shortage of capacity, a major constraint is weak visibility of where usable capacity exists at node and asset level. This causes projects to cluster at constrained points while spare capacity elsewhere is missed.

This paper estimates the annual economic cost of poor grid capacity visibility across Great Britain for connections from 11 kV to transmission, covering data centres, EV hubs, utility solar, grid-scale storage, electrified heat, commercial and industrial real estate, and public estates. We quantify four cost categories: 1) delayed and deterred investment, where private sector projects start later or not at all because usable capacity is unclear; 2) network reinforcement inefficiency, where avoidable or oversized works arise from sub-optimal siting; 3) system balancing and constraint costs, where curtailment and losses rise when assets connect at constrained nodes; and 4) additional developer costs, where application churn, studies, design rework and land negotiations add avoidable spend that is passed through to customers.

The central estimate of the total annual cost of poor capacity visibility is £2.71 billion, within a range of £1.25 billion to £4.69 billion. Of the four cost categories assessed, delay and deterred investment and network reinforcement inefficiency account for most of the total, at £1.20 billion and £1.08 billion respectively in the central case. System balancing and constraint costs add £133 million, and additional developer costs contribute £302 million. These costs are ultimately borne by consumers through higher bills and through lost or delayed growth.

The drivers are straightforward. Developers cannot see up-to-date, machine-readable information on capacity, contracted queue positions, timed reinforcements or feasible ramp and curtailment options at node level. As a result, siting and sequencing decisions are weaker than they could be, avoidable or oversized reinforcement is built, curtailment costs and losses rise, and application churn, redesign and unproductive land discussions increase.

The solution is to urgently enable access to datasets that can enable accurate grid capacity visibility. We recommend that the Government and Ofgem urgently close the remaining data gaps needed for accurate grid assessments. Specifically, we recommend:

1. A clear commitment from DESNZ and Ofgem to identify and close all data gaps needed for accurate grid assessments within 12 months.
2. A priority lane for enabling access to critical gaps within the next 3 months.
3. Accuracy assessments, including clear quality metrics (for example, completeness, timeliness, and error rates), published methodologies, and periodic independent audits so users can trust the data and identify issues quickly.
4. Interoperability and accessibility, data needs to be interoperable between datasets and so should include common unique identifiers of any assets referenced in a dataset.
5. Network data incentives: amend network data incentives to focus on incentivising networks to enable access to underlying data, rather than building digital products themselves.

Improving capacity visibility does not replace the need for timely network investment, planning reform or queue changes. However, it will reduce delay, cut avoidable build and operating costs, and help the UK realise more of the economic gains from AI and electrification every year.

1. CONTEXT

There are two megatrends that will significantly impact the economic success of countries this century: the AI boom and the electrification supercycle.

Data centres are critical to unlocking the economic benefits of AI, providing the computing infrastructure needed to develop, train, and deploy advanced models at scale. Independent analysis finds the UK data centre sector could add £44 billion in additional GVA by 2035, creating about 40,200 direct operational jobs and 18,200 construction jobs, with £9.7 billion extra in tax receipts if growth is enabled[1].

Within the electrification supercycle, the EV transition could lift the UK automotive sector's annual GVA by £16.1 billion by 2035 in a high-uptake scenario, supporting around 167,000 additional jobs across the economy[2]. On the supply side, solar is assessed to contribute around £3.5 billion GVA and 28,700 jobs a year by 2035, while grid-scale battery storage contributes about £1.6 billion GVA and 13,900 jobs a year by 2035[3]. For buildings, electrified heat delivers macro benefits in the near term, with modelling showing a £6.8 billion uplift to UK GDP in 2030 and about 138,400 full-time jobs when residential heat is decarbonised in line with an accelerated pathway[4].

Beyond these two megatrends, wider commercial and industrial development is projected to grow from ~£281-287bn in 2024 to ~£300-306bn by 2028.

However, the ability of the UK to realise these economic growth opportunities is limited by access to electrical power, which is now the binding constraint across all asset classes. Grid connection offers are, on average, about 5 - 5.5 years later than developers request. Reinforcement and equipment lead times compound this, with typical lead times reported at ~15 months for 33 kV transformers, ~2 years for 132 kV, and longer for EHV, alongside doubled switchgear lead times in parts of the supply chain[5].

However, it is not just the availability of capacity that is the issue. *Visibility* of where there is more or less capacity is critical. Available capacity varies by asset and location across the networks. Substations and cables all have different ratings, operational limits, operational load flows, planned upgrades and contracted capacity that changes over time. Yet developers cannot currently see a clear, up-to-date picture of this. Public electricity network heatmaps do not provide the context required by developers, rarely show sufficient detail to enable confident investment decisions, and do not allow for data to be combined with other data relevant to asset development. The result is that spare capacity at some nodes is not

[1] TechUK
[2] CBI
[3] Solar Energy UK
[4] Cambridge Econometrics
[5] Ofgem

discovered or used, while applications cluster at already constrained points. The objective of this paper is to set out the economic cost of the lack of greater grid capacity visibility and to propose recommendations for addressing these data gaps.

2. METHODOLOGY

We estimate the costs associated with poor visibility of grid capacity. The scope is connections in Great Britain from HV networks (11 kV) through to transmission, covering data centres, EV hubs, solar, grid-scale storage, electrified heat, commercial and industrial real estate, and public estates.

We quantify four cost categories using unit models by segment, then aggregate to £ per year for low, central and high cases.

Inputs come from observed pipelines and cost norms. For each channel we apply an affected-share and a visibility-attributable share so that only the portion caused by data gaps is counted. The framework is structured to avoid double counting across channels.

3. COSTS

This section reports results for the four factors and their low, central and high totals, with supporting tables:

1. Delay and deterred investment. Projects operate later or divert because usable capacity is unclear.
2. Network reinforcement inefficiency. Avoidable or oversized works arise from sub-optimal siting.
3. System balancing and constraint costs. Redispatch, curtailment and losses increase when assets land at constrained nodes.
4. Additional developer costs. Application churn, studies, rework and land negotiations that are passed through to customers.

Only the share attributable to visibility gaps is counted, not physics or planning. Totals are structured to avoid double counting between factors.

3.1 DELAYED & DETERRED INVESTMENT

Poor visibility of usable grid capacity delays and diverts investment. It affects the assets we assess here: data centres; EV charging hubs; utility solar and grid-scale storage; electrified heat in buildings; and large Commercial & Industrial (C&I) and real-estate projects, including public-estate schemes. We value the economic impact in two parts: operational GVA delayed (value that would have been produced if the asset were operating sooner) and construction GVA deterred (one-off domestic value that does not occur when investment does not occur). We attribute only the share caused by visibility gaps rather than physical limits, planning, or queue policy.

Delayed and deterred investment (£m per year)

Cost component	Low	Central	High
A. Baseline delay across the whole pipeline (short, widespread timing drag from coarse data)	150	300	450
B. Severe delay for the materially impacted subset (late node/voltage switches, missed ramp options)	240	480	720
C. Deterred operational GVA (projects that divert/park because a viable path to power isn't visible)	90	180	270
D. Deterred construction GVA (one-off UK build value that doesn't happen when investment moves/pauses)	120	240	360
Total	**600**	**1,200**	**1,800**

We value the economic impact of weak capacity visibility as the sum of delayed operational GVA and deterred value.

First, we recognise that all projects experience some friction because current maps and statements don't show usable headroom, contracted pipeline or dated reinforcements at node level. This creates a baseline delay (Row A) that applies to the entire pipeline: we estimate the monthly operating value (using a return-on-cost proxy) and multiply by a short delay period, then apply a modest "visibility share" to reflect that only the data gap (not physics or planning) is counted.

Second, a material subset of projects suffers severe delay (Row B). These are the cases where lack of node-level clarity leads to late node/voltage switches, missed ramped connections, or multiple failed applications before a viable point is found. Here we use longer delay months and a higher visibility attribution, consistent with what developers report in practice (especially for large loads).

Third, some investment is deterred or relocated because developers cannot evidence a credible, timely path to power at the outset. We split this deter effect into operational GVA (Row C), which does not materialise in the UK while the project is elsewhere or paused, and construction GVA (Row D), a one-off loss of domestic build activity when construction shifts abroad or is shelved. For construction we apply a conservative value-added ratio (around 30 percent of capex) to the deterred portion of the pipeline and again apply the visibility share so we only attribute the data-gap component.

The low/central/high scenarios vary three levers in combination: the delay months (shorter vs longer), the size of the severely impacted subset (smaller vs larger share of the pipeline), and the deterred volume (e.g. fewer vs more large campuses and power-hungry schemes relocating), along with conservative shifts in the construction value-added ratio.

3.2 NETWORK REINFORCEMENT INEFFICIENCY

This factor captures the avoidable network build that arises because developers cannot see, with sufficient detail, which connection nodes could offer capacity with fewer upgrades. Without clear, node-level visibility of usable capacity, projects often proceed at a point of connection that requires upgrades which could have been avoided by choosing a different node or better sequencing the connection to align with scheduled works. The cost shows up in two places. First, a socialised portion falls on all customers' bills where wider reinforcement is funded through network charges. Second, a developer-borne portion is paid directly by the connecting party for sole-use or extension assets and for reinforcement triggered at the connection voltage; these costs then flow through to consumers via higher prices for the goods and services those assets provide. In this section we count only the inefficient build itself (the avoidable reinforcement spend), attributing the part caused by lack of visibility.

Network reinforcement inefficiency (£m per year)

Cost component (inefficient build only)	Low	Central	High
A1. Distribution – socialised on bills	210	420	640
A2. Distribution – developer-borne	130	260	400
B1. Transmission – socialised on bills	140	280	430
B2. Transmission – developer-borne	60	120	180
Total	**540**	**1080**	**1650**

To assess avoidable network build due to a lack of visibility we start from two separate spend bases and keep them distinct by voltage level and by who pays.

First, for the socialised component where the costs of reinforcement are added to energy bills, we take representative Great Britain annual envelopes for capacity-driven reinforcement inside the price controls: distribution load-related reinforcement (ED2) and transmission demand-driven wider works (ET2/NOA delivery). These are costs that, if incurred, are recovered through network charges and therefore land on bills.

Second, for the developer-funded component, we construct a parallel base for connection/extension/sole-use works and any reinforcement at the connection voltage that is charged to the connectee; these are not in ED2/ET2 and are paid directly by developers.

On each base we apply a visibility-attributable inefficiency rate that represents schemes which proceed at a point of connection that require reinforcement which could have been avoided, downsized, or shifted had capacity been visible at the siting/phase-planning stage. We use different ranges for distribution and transmission, reflecting that siting decisions tend to drive a larger share of avoidable build at distribution, and a somewhat lower but still material share at transmission. The low/central/high scenarios simply vary these rates within defensible bands; multiplying those rates by the corresponding spend bases yields the figures in the table for the four rows (distribution/transmission × socialised/developer-borne). Summing row pairs gives the socialised and developer totals, and summing all four rows gives the grand total.

3.3 SYSTEM BALANCING & CONSTRAINT COSTS

Poor siting driven by weak grid visibility does not only add to build costs. It also raises the day-to-day cost of running the electricity system. When new demand or generation is sited sub-optimally, NESO has to intervene to keep flows and voltages within limits. Most of these actions are taken through the Balancing Mechanism, where NESO accepts bids to turn down and offers to turn up to manage constraints on lines and boundaries. Constraint management often means paying one plant to reduce output while paying another elsewhere to increase output. Some actions are taken outside the BM through contracted services that are brought into settlement (non-BM actions). NESO may also procure ancillary services for voltage or stability that become more expensive if power is in the wrong place. In parallel, sub-optimal siting increases electrical losses when power travels further or through already loaded asset

System balancing and constraint costs (£m/year)

Cost component	Low	Central	High
Redispatch & constraints (share of NESO balancing costs)	18	87.5	255.2
Losses (share of GB T&D losses valued at wholesale)	7.7	45.5	164.8
Total	**25.7**	**133**	**419.9**

The approach for this factor anchors on observed system costs. Annual spend on redispatch and constraints, system curtailment payments, and the cost of transmission and distribution losses are taken as the starting point. Only the portion plausibly driven by poor siting is attributed to visibility gaps, using conservative shares. Results are cross-checked with a small set of worked examples at known constrained boundaries. These exemplars illustrate how new demand or generation in the wrong place lifts redispatch, curtailment and losses, and they provide a reality check on the top-down attribution.

Low, Central and High scenarios come from varying the size of the cost bases and the attribution shares. The intent is to show an auditable range that reflects uncertainty without overstating impacts.

3.4 ADDITIONAL DEVELOPER COSTS

Poor visibility does not only slow projects or raise system costs. It also creates additional avoidable costs for developers that are ultimately passed through to customers. When developers don't have good visibility of grid capacity, they submit multiple applications, waste time negotiating with landowners, and often have to redesign schemes. These costs sit on developer P&Ls but are paid for in the end by consumers.

Additional developer costs (£m per year)

Cost component	Low	Central	High
Applications & fees	40	100	180
Studies & surveys (internal + external)	20	78.8	212.5
Design & engineering rework	13.5	75	280
Land discussions & negotiations	8	48	144
Total	**81.5**	**301.8**	**816.5**

Estimates anchor on evidenced application activity and fee levels. UKPN received more than 70,000 applications in 2023. Using an average blended Application & Design fee of around £1,000, and scaling to all DNOs yields an indicative 250,000 applications and about £250 million a year in grid application fees.

For studies and surveys, design and engineering rework, and land discussions, the calculation links simple activity ratios to the application base. A conservative proportion of applications is assumed to trigger extra studies, a subset to progress to offers that then require redesign after unfavourable terms, and a further share to kick off early land engagement before capacity clarity. Each activity uses modest unit costs that blend internal time with light external spend, and only the portion plausibly caused by poor visibility is attributed.

3.5. TOTALS ACROSS ALL FOUR FACTORS

The figures below add up the visibility-attributable costs from the four factors set out above.

Total economic cost of poor grid-capacity visibility (GB, £m per year)

Cost component	Low	Central	High
1) Delay and deterred investment	600	1,200.00	1,800.00
2) Network reinforcement inefficiency	540	1,080.00	1,650.00
3) System balancing and constraint costs	25.7	133	419.9
4) Additional developer costs	81.5	301.8	816.5
Total (all factors)	**1,247.20**	**2,714.80**	**4,686.40**

4. ENABLING GRID CAPACITY VISIBILITY

Both the government and Ofgem have been supportive of enabling grid capacity visibility. Two approaches have been explored to enable visibility.

The first has been that Ofgem has required the networks to publish heatmaps of network capacity. However, these have proven to be insufficient solutions to the challenge as they represent over simplifications of network assessments. Heatmaps generalise over areas, hide node-level details, and are often static or lagged. They cannot express the possibility of ramped connection profiles, curtailment risk, or dated reinforcement effects in a way that supports decisions. More generally, heatmaps represent an attempt to deliver centralized digital innovation by regulatory direction - an approach that was arguably unlikely ever to be successful, and one that has deterred innovators and investors from developing better solutions.

The second approach has been to require that the networks publish network data via open data portals on the details of their networks. Significant progress on open and shared data has been made. This progress has enabled significantly greater network visibility. However, critical gaps remain. These gaps mean that despite all the data that has been published, it is still not possible to accurately assess grid capacity.

There are processes to close some of these gaps that will further improve capacity visibility, however these efforts are insufficiently resourced and progress is very slow.

The costs highlighted in this paper could be avoided if there were a concerted effort to accelerate the process of enabling wider access to the data required to enable grid capacity visibility.

5. RECOMMENDATIONS

To enable timely publication of the electricity network data required to enable accurate, up-to-date grid capacity assessments, our recommendations are as follows:

01 **A clear commitment from DESNZ and Ofgem to identify and close all data gaps needed for accurate grid assessments within 12 months.**

02 **A priority lane for enabling access to critical data gaps within the next 3 months.**

03 **Accuracy assessments, including clear quality metrics (for example, completeness, timeliness, and error rates), published methodologies, and periodic independent audits so users can trust the data and identify issues quickly.**

04 **Interoperability and accessibility, data needs to be interoperable between datasets and so should include common unique identifiers of any assets referenced in a dataset.**

05 **Network data incentives: amend network data incentives to focus on incentivising networks to enable access to underlying data, rather than building digital products.**

6. CONCLUSION

The evidence in this paper shows that poor visibility of usable grid capacity carries a material, recurring cost to the UK economy. Across the four categories assessed, the total cost is between about £1.25 billion and £4.69 billion a year, with a central estimate of £2.71 billion. Most of this burden arises from avoidable delay and deterred investment and from inefficient reinforcement decisions, with a smaller but still meaningful contribution from higher system balancing costs and developer rework. These costs ultimately fall on consumers through higher bills and foregone growth.

Improving visibility does not create capacity and does not replace the need for timely network investment, planning reform or queue management. It does, however, let developers and networks use existing and planned capacity more efficiently and lower the cost of operating the system. The solution is primarily a data task: making accurate, machine-readable network information available at asset and node level, with clear quality metrics and interoperability.

Delivering the recommendations in this paper would give developers and networks a shared view of the system within months, not years. That would accelerate viable projects, reduce avoidable reinforcement, lower constraint costs, and cut wasted developer spend. In the context of the UK's AI and electrification ambitions, enabling grid capacity visibility is a cheap and no regrets action that supports every other electricity system reform and pays back annually

.

Yottar

ABOUT YOTTAR

Yottar is an energy-tech startup building a digital twin of the grid to enable accurate grid capacity assessments. Yottar works with asset developers and landlords to support optimised site development and prioritisation according to grid capacity. To find out more please get in touch: hello@yottar.tech.