# Hard to Compute

## Why Startups Need More Power

November 2025

STARTUP
C*ALITION

# Authors

**Martha Dacombe**
Emerging Technologies, Policy Lead

**Vira Veselukha**
ClimateTech, Policy Lead

**Vinous Ali**
Deputy Executive Director

# About Startup Coalition

Startup Coalition is an independent advocacy group that serves as the policy voice for the UK's technology-led startups and scale-ups. We were founded in 2010 by Mike Butcher, Editor-at-Large of TechCrunch, and Jeff Lynn, Chairman and Co-Founder of online investment platform Seedrs. We fight for a policy environment that enables early-stage British tech companies to grow, scale, and compete globally.

Our network includes over 4000 startups, scale-ups, and investors. We have been instrumental in building proactive coalitions of businesses and investors on issues integral to the health of the UK's startup ecosystem. Our work has seen many successes, from the establishment of the Future Fund to the creation of the Scale-Up Visa.

# Introduction

Compute is the lifeblood of the digital economy. It is the infrastructure that powers artificial intelligence, enables machine learning breakthroughs, processes vast datasets, and drives the digital transformation that is reshaping every sector from healthcare to finance. As we enter into the next Industrial Revolution, access to computational power has become synonymous with economic competitiveness, scientific leadership, and national security. Without sufficient computational capacity, even the most brilliant innovations remain locked in laboratories, and the most promising startups hit insurmountable walls.

The Government's *Compute Roadmap* can be sorted into three broad buckets. Firstly, building out AI growth zones; secondly, bolstering and expanding the AI Research Resource (AIRR), and finally, evolving procurement to support the development and diversification of the stack in the UK – crucially ensuring that UK startups are at the heart of this process.

There is no doubt that the Roadmap demonstrates genuine ambition; however, executing it will require laser-focus from the Government and them reaching out into spaces they don't currently occupy.

**We may be biased, but we firmly believe that in this case UK startups will not only be net beneficiaries of getting this right, but will actually help enable and accelerate execution of the plan.**

This report therefore is in two parts:

**PART ONE:** Examines the need for compute and offers thoughts on how a blended model with elements reserved for secure and sovereign compute would ensure the UK brings compute on stream quickly to meet growing demand.

**PART TWO:** Outlines how at every part of the stack, and in related areas, startups can support Government to meet the objectives laid out in the *Compute Roadmap*, as well as simply being beneficiaries. From chip design to reimagining data centres, startups are innovating to meet future demand in a way that is efficient, sustainable and sovereign (where it needs to be).

# The History of Compute Policy Approaches

**March 2014** — The official launch of ARCHER supercomputer, housed at the University of Edinburgh's Advanced Computing Facility.

**November 2021** — ARCHER2 launches post £79m invested - 11x more powerful than ARCHER, and places UK amongst world's compute leaders.

**June 2022** — Worlds first exascale machine (1000 pflops compared to ARCHER2'S 28/s) online. The Government commissioned an Independent Review on the Future of Compute. The review warned that the UK was falling behind without a long term plan. It also recommended building exascale supercomputer, dedicated AI compute clusters and national coordination body.

**August 2022** — The CHIPS and Science Act passed in the USA, authorising $280bn in new funding to boost domestic research and manufacturing of semiconductors in the USA.

**March 2023** — Government publishes Future of Compute Review Independent Report, and earmarks £100m for the new AI Research Resource (AIRR).

**October 2023** — Plans for first UK exascale supercomputer announced.

**November 2023** — AIRR upgraded to £300m. Enables Dawn, a Cambridge supercomputer alongside the "Isambard" AI system in Bristol, together forming the AIRR. Could mean a 30x increase in UK public compute capacity.

**August 2024** — UK Government shelves £1.3bn of tech and AI projects - scraps plans for first exascale supercomputer in Edinburgh and additional AIRR funding.

**January 2025** — AI Opportunities Action Plan published, the Government accepts all 50 recommendations from Matt Clifford. Government announces expanding AIRR capacity, by at least 20x by 2030, as part of the AI Opportunities Action Plan. AI Growth Zones announced, these intend to help accelerate the build out of data centres.

**February 2025** — The Government opens bidding for AI Growth Zones for local authorities.

**June 2025** — Chancellor Rachel Reeves confirms £750 million investment for a national supercomputer hosted at the University of Edinburgh.

**July 2025** — UK Compute Roadmap published - sets out a 10-point plan including investment of up to £2 billion to deliver a diverse, joined-up compute ecosystem. AIRR capacity to increase from 21 AI ExaFLOPS in 2025 to 420 AI ExaFLOPS by 2030.

**September 2025** — The Government designates the North East of England as an AI Growth Zone.

**November 2025** — AI Growth Zone in North Wales announced, along with reforms on planning, energy access and the provision of support from the Government to accelerate AI Growth Zone development.

In 2023, under the last Conservative Government, the *Independent Review on the Future of Compute* was published. This report, led by Zoubin Grahramani FRS, established compute as a **critical capability** requiring urgent government intervention. The Review linked the success of compute build-out as integral to success in other policy areas, such as economic growth, net zero targets and overall competitiveness of the British digital economy.

In January 2025, the then new Labour Government published the *AI Opportunities Action Plan,* written by Matt Clifford. The plan elevated compute beyond a technical resource to a matter of economic security, arguing that access to compute would form a key pillar of the UK's economic resilience. This represented a subtle but important shift in emphasis from the Independent Review, and pushed the Government to take seriously the procurement of sovereign, and domestic compute resource as essential to the competitiveness of the British economy.

In the *AI Opportunities Action Plan,* developing sovereign compute capabilities was designated as a specific aim. Indeed, in the Plan, sovereign AI is described as 'owned/and or allocated by the public sector, [and] will enable the UK to quickly and independently allocate compute to national priorities'.[1] This emphasis on a sovereign capability is important, as it compels the Government to seriously invest in compute to power access to this critical resource for startups, academics and researchers building the future. However, the plan also outlines how the UK will also need privately owned compute to meet the grand ambitions of the Roadmap to expand overall capacity in the UK.

It should not be ignored that the UK is already in a strong position when it comes to the development of AI. The UK's AI ecosystem represents a genuine success story, there are over 3,700 AI companies in Britain, and the overall sector was valued at a staggering £72.3bn in 2024[2]. The UK is Europe's undisputed leader, and this is a lead that other countries in Europe and further afield envy. However, not a week passes without announcements of new investment, regulatory loosening, or new initiatives to make 'x country' a world leader. The Government should be conscious of this and not rest easy.

[1] AI Opportunities Action Plan
https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan
[2] UK Government https://www.business.gov.uk/campaign/grow-your-business-in-the-uk/artificial-intelligence/

## What the Flop is Compute?

*The Independent Review of the Future of Compute* in 2023 defined compute as an umbrella term, including but not limited to, high performance computing (HPC), high throughput computing (HTC), large scale computing (LSC) and supercomputing. Cloud computing does not fall entirely under this umbrella. Cloud compute offers scalable access to existing commercial hardware, but not necessarily the sovereignty or security guarantees required for national AI infrastructure. Sovereign compute implies physical control, location certainty, and prioritised access in times of disruption. The Review also outlined that a common way to measure computational power is in 'flops' – floating point operations per second: a floating point operation is roughly equivalent to a single arithmetic calculation involving two numbers.

Compute falls into three categories, of firstly: public compute, owned or operated by the British Government This can be considered genuinely 'sovereign'. The second form is private compute, which is owned or operated by private firms, but still within the UK. This could be a privately owned data centre for example. And, finally, there is overseas or foreign compute, which is located out of the UK, and is operated by foreign providers, but access can be rented out or used by British companies if they wished to.

| Type | Classical HPC | AI Compute | Emerging and Novel Uses |
|---|---|---|---|
| Uses | Scientific research<br><br>Engineering simulations<br><br>Climate modelling<br><br>Data-intensive academia, material science<br><br>Relies heavily on machine learning but not AI-specific acceleration | Training, running, and production of AI models<br><br>Requires GPU clusters, TPU pods, and hyperscale data centres<br><br>Training and inference have different compute needs | Quantum processors<br><br>Neuromorphic chips<br><br>Experimental accelerators |
| Infrastructure | Large-scale compute with extensive memory and storage<br><br>Cheaper processing power suffices (no need for AI acceleration)<br><br>Needs to be stable and continuous | Specialised compute resources tuned for training vs inference<br><br>Needs infrastructure for sustained deployment of trained models<br><br>Geographic distribution is critical | Modular data centre space<br><br>Plug-and-play access to diverse compute<br><br>Facilities must be retrofittable<br><br>Infrastructure built for experimentation |
| Challenges | Oversubscribed resources<br><br>Risk of being sidelined in the AI boom<br><br>Concerns about access + continued investment in non-AI workloads | Founders burn through cloud credits quickly (especially on inference)<br><br>Tight startup budgets hinder access<br><br>Regional gaps in availability and accessibility | Lack of infrastructure tailored to emerging tech needs<br><br>Shortage of quantum and emerging talent<br><br>Funding mechanisms are often not fit-for-purpose or too narrowly defined |
| Needs | A national strategy that expands access to foundational compute via grants or cloud credits<br><br>Sustained investment and predictability in hardware refreshes | Affordable, accessible compute for both training and inference<br><br>Programmes for credits or subsidised access<br><br>Distributed access across regions | Flexible, modular infrastructure, cross-domain expertise, and open-ended funding streams<br><br>A national springboard for experimentation and testbeds |

## The Compute Roadmap – What Does It Outline?

**On the issue of domestic research capacity**. The Government aims to 20x AI Research Resource (AIRR) capacity by 2030, from 21 AI exaFLOPS to 420 AI exaFLOPS. This expansion would ensure the AIRR enables the training of multiple AI models a year, provides an up-to-date research capability and is explicitly designed to keep up with the expected exponential growth in AI computational demands**.**

To achieve this, £2 billion has been committed to building modern public compute systems through 2030, backed by two new AI supercomputers that will form the backbone of the AI Research Resource. To fulfill this strategy, the Government is clearly getting going, and currently has tenders on their website for the AIRR, including one for supercomputing hardware[3], and another for AI cloud compute[4]. On top of this, the Roadmap outlines a plan to develop AIRRPORT, a single access portal to AIRR, ending the fragmented maze that currently confronts users seeking computational resources.

Furthermore, the Sovereign AI Unit and AI Security Institute will receive dedicated compute access, ensuring national priorities have the computational muscle they need. The Government also announced a new network of National Supercomputing Centres (NSCs), which aim to contribute to a federated and strategic experience of the UKs compute capabilities, and the Edinburgh Parallel Computing Centres (EPCC) will be the UK's first National Supercomputing Centre as part of this cluster of Government capacity.

**The infrastructure promises are comprehensive.** AI Growth Zones (AIGZ) are central to this, designated as areas for developing compute capacity and power throughout the UK, like data centres and supporting compute facilities, in which investment and jobs flow from. The latest announcement of a new AIGZ in North Wales, for example, will create an estimated 3,000+ jobs.[5]

**Building sovereign AI and leveraging existing compute capabilities are outlined as key priorities in the Roadmap.** The Roadmap recognises that the UK is advanced in areas such as chip design, and outlines how the UK needs a more effective pull through mechanism of procurement to enable 'the most promising British technologies to scale and succeed'.[6] Indeed, the Roadmap targets compute stack development through reformed procurement systems. Action 9 commits to supporting British companies in developing sustainable compute technologies, leveraging AIRR and AIGZs for access and scaling opportunities, working alongside initiatives like ARIA's scaling compute programme and the National Semiconductor Centre. Crucially, Action 10 changes procurement eligibility criteria to prevent novel British compute technologies from losing out to incumbents, while establishing compute as a priority for the Sovereign AI Unit to ensure domestic companies deliver critical AI infrastructure.

The Roadmap is the most comprehensive compute strategy the British Government has ever attempted. The question now is not whether the ambition is sufficient, but whether it can be executed at the pace innovation demands, and critically, whether it will actually reach the startup ecosystem that drives economic growth. The Roadmap focuses heavily on research infrastructure and cites national capacity,

[3] UK Government Tender Notice https://www.find-tender.service.gov.uk/Notice/062921-2025
[4]UK Government Tender Notice https://www.find-tender.service.gov.uk/Notice/062921-2025
[5] Department for Science, Innovation and Technology https://www.gov.uk/government/news/ai-growth-zones-to-create-thousands-of-jobs-and-unlock-up-to-100-billion-in-investment-as-new-site-confirmed-for-north-wales
[6] UK Compute Roadmap – https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

but the pathway from these flagship investments to accessible, affordable compute for British startups and scale-ups remains unclear.

# The Infrastructure Challenge

While the UK excels in talent, research, and operating environment, we rank a dismal 17th globally in AI infrastructure. According to the Social Market Foundation's analysis, "infrastructure is the Achilles heel of the UK's AI ambitions."[7] The UK faces unprecedented economic and technological risks if its compute capacity development fails to match international competitors and domestic demand. The Tony Blair Institute for Global Change has warned that compute infrastructure disparities between nations represent the new digital divide, emphasising that "compute is not just a source of scientific and economic progress, but the new benchmark of global power economically and geopolitically."[8] This creates a frustrating dynamic where we seed innovative companies full of exceptional British talent, yet there is no dynamic compute ecosystem for them to easily access when they want to scale or innovate.

Our own economic modelling produced independently by policy and research consultancy Public First estimates that based on current trends, the UK will require an additional 414 data centres (5.75 GW) to fully meet growing demand for compute, digital services, and AI workflows. This chimes with the Government's own projections which estimated that the UK will need at least 6GW of AI-capable data centre capacity by 2030.

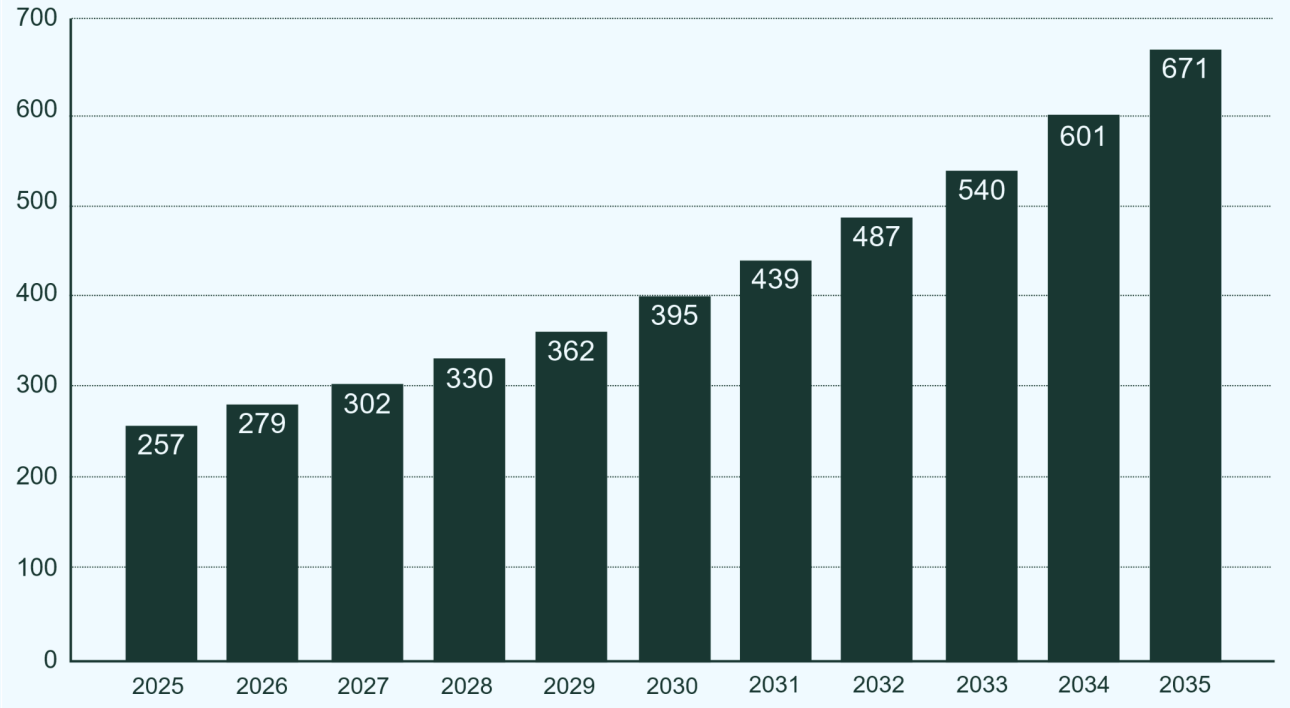


*Projected Number of Data Centres (2025 - 2035)*[9]

However, we are currently only on track to deliver 63% of this required capacity, highlighting a significant infrastructure gap. Without accelerating investment and deployment, the UK risks limiting its digital growth and competitiveness in AI and emerging technologies.

[7] Social Market Foundation: How to power AI: Boosting compute capacity for UK AI https://www.smf.co.uk/publications/data-centres-and-uk-ai-ambition/
[8] Tony Blair Institute for Global Change Report: State of Compute Access: How to Bridge the New Digital Divide https://institute.global/insights/tech-and-digitalisation/state-of-compute-access-how-to-bridge-the-new-digital-divide
[9] Public First Economic Modelling

On the issue of building large scale infrastructure, recent high-profile investments announced during President Trump's second state visit provide cause for measured optimism regarding these challenges. Multiple deals included NVIDIA's commitment to sell 120,000 GPUS, totalling £11 billion to the UK [10], representing a significant vote of confidence in Britain's technological future, along with CoreWeave's collaboration with DataVita to establish a new data centre in Lanarkshire.[11] Further, in September 2025, AWS and the Government confirmed an £8 billion investment for the construction, operation and maintenance of data centres[12], and Google has agreed to invest £5 billion over the next two years to expand an existing data centre in Hertfordshire.[13]

However, one major constraint remains: the availability of sufficient energy for data centres. These facilities are extremely power-intensive, and the UK's aging power grid has struggled to keep pace with surging demand; approximately two-fifths of existing data centres have already encountered power availability issues as a result.[14]

This strain has led to long delays in hooking up new facilities, turning access to electricity into the single biggest limitation on data centre growth. In the London area alone, about 400 GW of pending data centre connections are queued for grid access, far more than current infrastructure can support.[15] Regulators estimate most of those projects may never materialise under present conditions.[16] Without rapid investment in grid capacity and sustainable energy, these bottlenecks could slow the rollout of AI hardware despite new investments, a concern the Government is beginning to address by designating data centres as Critical National Infrastructure and pursuing grid connection reforms.

While these investments signal growing confidence in the UK's AI potential, private capital alone cannot close the infrastructure gap. The Government must translate its *Compute Roadmap* commitments into rapid, coordinated action, ensuring sustained access to affordable compute resources, combined with material support for startups who are building the solutions to the bottlenecks which undermine our compute ambitions. If the Government does not take this opportunity, it will risk watching homegrown innovation migrate to jurisdictions with superior infrastructure.

**Why does this matter to startups?**

Already, AI startups face a compute cliff as they grow. In the early seed stage, access is straightforward and affordable, hyperscalers like AWS, Google Cloud, and Azure typically offer generous credits to attract promising young companies. These programs effectively subsidise experimentation and initial development. But this changes dramatically at scale.

---

[10]The Guardian https://www.theguardian.com/technology/2025/sep/17/jensen-huang-nvidia-uk-ai-superpower-500m-nscale
[11] Data Vita Blog https://blog.datavita.co.uk/blog/datavita-and-coreweave-to-launch-ai-infrastructure-in-scotland-powered-by-renewable-energy-as-part-of-ps1-5bn-investment
[12] UK Government Press Release https://www.gov.uk/government/news/chancellor-announces-8-billion-amazon-web-services-investment-as-she-vows-to-make-every-part-of-britain-better-off
[13] Google Press Corner https://www.googlecloudpresscorner.com/2025-09-16-Google-Opens-Waltham-Cross-Data-Centre-as-Part-of-Two-year-GBP5-Billion-Investment-in-the-UK-to-Help-Power-its-AI-Economy
[14] Aggreko https://www.aggreko.com/en-gb/news/2022/noeur-news/power-struggle-for-data-centres
[15] https://www.theregister.com/2025/02/10/london_has_400_gw_of/
[16]The Register https://www.theregister.com/2025/02/10/london_has_400_gw_of/

As startups mature, compute costs can skyrocket from manageable to existential. The credits dry up, usage multiplies exponentially, and before they know it infrastructure spending is eating into a higher proportion of their runway. For many AI startups, this transition from subsidised compute to market rates creates a make-or-break moment. At the same time, the UK's public compute capacity, as organised through the AIRR, is often difficult for startups to navigate and get reliable access to, as the resource has often operated specifically for academic users.

This shortfall hits startups particularly hard. While academic users often require massive computational resources for research, industry users, especially startups, face a different problem. As UKRI found, industry users require less computational power on average than their academic counterparts, but cost is paramount for these users which is compounded by their reliance on cloud-based computing solutions."[17]

This brings us to **PART TWO - where startups can support.**

[17] UKRIhttps://engagementhub.ukri.org/ukri-infrastructure/shaping-the-future-of-uk-large-scale-compute/

# Procurement Reform and Leveraging Existing Capabilities to Build British Compute Capacity

The UK Government's Roadmap states that in an 'ideal end state, we envision an AI Growth Zone powered by a full UK designed compute stack – from chip to system to software – demonstrating what sovereign capability truly looks like, and opening up new export and growth opportunities for UK firms'[18]. The report outlines a route of creating a 'Compute Bridge', as a pull through mechanism for British compute technologies. This will operate firstly, via R&D and commercial routes, through for example UKRI and strategic programmes like ARIA and the Semiconductor Centre. Secondly, the 'Bridge' consists of support for companies through incubation, testing and integration in the National Supercomputing Centre, along with the procurement of novel computing technologies. Finally, it consists of proper scaling up and implementation achieved through scale deployments, where there are demonstrated deployments at scale in AI Growth Zones, and products are made ready for mass adoption.

The Roadmap prioritises achieving Sovereign AI by developing a compute stack built on strategic British technologies. This aims to establish the UK's technological independence in the global compute race: reducing reliance on foreign technologies while positioning the UK as an essential supplier of critical technologies that other nations depend on.

For this to take place, the Roadmap outlines

- **Action 9:** for the Government to 'support British companies to develop sustainable and secure compute technologies, using AIRR and AI Growth Zones to provide access, testbeds and opportunities to scale', and
- **Action 10:** to 'establish compute as a priority area for the Sovereign AI Unit, supporting research into new computing paradigms and helping British companies grow into global leaders'.

Both these actions aim to ensure that the domestic capabilities, in areas such as hardware, for example, receive the support needed to reach scale and in turn reduce reliance on others for AI infrastructure. Further, in the Roadmap, the Government outlines how it intends to allocate a portion of AIRR capacity specifically for high potential technologies, including the likes of chips, architectures and software stacks. We welcome these recommendations as essential commitments to ensure the success and scalability of British startups building in the compute stack.

## Where to Go Next

While these commitments to action are the correct ones, we believe that the Government can go further. Startups and scaleups in the UK are right now offering the answers to the key issues and bottlenecks in leveraging sovereign capabilities, and the Government should use its position to invest in these technologies with absolute urgency, as these companies cannot be left waiting around.

[18] UK Government Compute Roadmap
https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

It is worth noting that there has been a genuine shift in the Government's approaches to procurement, not only within legislation like the Procurement Act 2023, but also in the policy direction set out in the *AI Opportunities Action Plan*, and of course, the *Compute Roadmap*. Specifically, the Roadmap outlined that in order to develop British compute technologies, the Government will 'improve the eligibility of novel compute technologies for procurements' and that this will ensure 'that innovative and emerging technologies are not discounted from procurement exercises because they do not meet criteria set to evaluate performance against incumbent solutions'.[19]

This could be brought to life through the use of **Advanced Market Commitments (AMCs).** These are binding commitments to purchase innovative products or services that don't yet exist commercially. In short they create a viable market to propel innovations that may otherwise take years to get to market.

Guaranteed demand, so to speak, can ensure the Government gets its hands on the latest innovations whilst also acting as a clear demand signal to other potential buyers and investors while de-risking that investment at the same time. Fundamentally, future contracts can be leveraged by startups to unlock significant private investment, giving them the runway to build.

Regular procurement processes do none of the above and would fail UK startups miserably while also ill-serving a Government that has ambitions of leading the revolution from the front.

**Therefore we propose:**

1. **To deploy AMCs for strategic technologies.**
   Target early-stage innovations critical to building sovereign compute, this can be informed by the critical technologies outlined in the Industrial Strategy.
2. **Link AMC milestones to national compute needs.**
   Structure AMC contracts around specific Government performance targets, such as improving compute energy efficiency or reducing latency across AIRR-connected systems.
3. **Bridge the R&D-to-market gap**.
   Use AMCs as a complementing scheme to UKRI and Innovate UK grant programmes, providing commercial pull-through for technologies proven in pilot phases but not yet viable for traditional procurement.

This approach would supercharge the domestic compute ecosystem, build more advanced technologies than those that exist in the market currently, and ensure that British innovations find their first customers in British Government systems. The Government must follow its own Industrial Strategy, in prioritising frontier technologies with what they describe as holding 'the greatest growth potential'[20].

[19] UK Compute Roadmap – https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

[20] UK Government Industrial Strategy https://assets.publishing.service.gov.uk/media/68595e56db8e139f95652dc6/industrial_strategy_policy_paper.pdf

## Speed and Competitive Necessity

The UK's response must match the pace and ambition of startups. Traditional procurement cycles measured in years cannot compete with Silicon Valley speed, where product generations turn over in 18-month cycles. This speed mismatch has concrete consequences. For example, chip supply chains require 1-2 year lead times for orders, meaning a department that procrastinates on procurement doesn't just delay, it guarantees obsolescence.

Moving to a new mode will need to be supported by new ways of thinking. Building on the *AI Opportunities Action Plan*'s idea of AI leads to drive mission-focussed work, we believe a Compute Mission Director in DSIT can spearhead action on AMC rollout.

Cheque-size as well as speed is also critical. The demand signals and de-risking that markout AMCs only work if the guaranteed contracts are large enough. Government must not salami-slice this funding across multiple initiatives but instead bet big on credible winners.

Finally, getting this right will require significant technical oversight and expertise, as demonstrated in organisations like ARIA, to validate credible delivery paths and distinguish genuine innovation from hype. Therefore, AMC milestones must link directly to national compute needs and performance targets, ensuring the Government purchases address specific strategic requirements rather than simply subsidising firms.

Nation states can successfully support the development and manufacturing of chips, but this demands an exceptional degree of commitment to make it happen. ARIA has proven successful as a market-making and market-shaping force by acting with speed, taking risks appropriate to the technological moment we face. The Government must learn from this model and apply it to compute procurement specifically.

Through AMCs, the Government can strengthen existing British capabilities in the compute stack and build a sovereign base of strategically important technologies. Britain can establish strategic interdependence in the supply chain of AI infrastructure, not by hoping startups survive the valley of death, but by guaranteeing the demand that makes their success inevitable.

**Therefore we propose:**

1. **Appoint a Compute Mission Director**.
    a. This role would have ministerial backing, budget authority, and explicit mandate to pick winners and move at commercial speed.
2. **Ensure that AMCs have heft.**
    a. This could potentially be £100m+ within a single AMC run rather than fragmenting investment across multiple smaller initiatives.
3. **Establish technical oversight capability modelled on ARIA**.
    a. This would hold the ability to validate delivery paths, assess chip architectures, and structure milestones around performance rather than process.
4. **Issue clear ministerial direction.**
    a. This direction should instruct commercial and legal teams to enable innovative procurement with high risk tolerance, explicitly prioritising domestic capability development.

5. **Link AMC milestones to *Compute Roadmap* targets**.
    a. Ensuring purchases address specific national compute capacity needs and performance requirements.
6. **Accelerate procurement timelines.**
    a. This would match chip supply chain realities, as decisions must happen within months, not years, to avoid purchasing obsolete technology

### Case Study: Vaire:

**Vaire Computing is a UK-headquartered deep-tech semiconductor company developing near-zero energy processors for the AI era.**

Founded in 2021, the company is pioneering adiabatic reversible computing (ARC), a breakthrough approach that recycles, rather than dissipates, energy during computation. This allows performance to continue scaling even as power budgets remain fixed, addressing the most pressing constraint in AI infrastructure: energy efficiency.

Vaire's test chip validates that the core technology behind near-zero energy chips works - demonstrating **1-4 - 1.8x energy recovery in standard CMOS**, with no need to rewrite existing software. Vaire's roadmap for further chip tapeouts will focus on delivering order-of-magnitude improvements in energy efficiency in transformer workloads.

**Vaire sells its technology in two ways:**

- Complete chip systems for companies running AI applications (like telecoms, data centres, and defence), designed to work easily with existing open-source software.
- IP licences for chip manufacturers who want to build Vaire's technology into their own products.
- Strategically, Vaire is an energy-sovereign alternative to the incumbent GPU monopoly. Its technology offers Governments and hyperscalers a path to expand AI capacity without exceeding power and carbon constraints - a critical enabler for the UK's and Europe's long-term sovereign compute ambitions. Since demonstrating energy recovery with its first test chip, Vaire has started work with partners across the compute and energy ecosystem and has letters of intent from multiple enterprise customers.

## AIRR: Changing Access to Compute

AIRR can best be understood as the centre of public compute in the UK, and in the *Compute Roadmap*, the Government outlined that it intended to increase the size of this resource by 20x.

The Government has supercomputers such as Isambard-AI, hosted at the University of Bristol, and Dawn in Cambridge, hosted at the University of Cambridge – these supercomputers are main constituent parts of the AI Research Resource. Furthermore, to meet the Government's ambitions of expanding the AIRR's capacity, the Comprehensive Spending Review confirmed funding of up to £750 million for a new supercomputer at the University of Edinburgh, which will far exceed the capacity of ARCHER2, the current supercomputer hosted by the university.[21]

In the *Compute Roadmap*, the Government also indicated that the AIRR will also include complementary cloud compute capabilities, in order to deliver a blended model that industry, academia and research can access, and that this will include a new network of National Supercomputing Centres (NSCs).

In order to achieve these ambitious goals for AIRR, the Government committed itself to three objectives to kickstart the delivery. First, they will launch a programme to identify opportunities for public-private partnerships that could coordinate the delivery and operation of new supercomputing capacity; second, they will commence procurement of cloud compute; and finally, they will deliver AIRRPORT, which will provide a single front door to the AIRR service.

## How Does AIRR Currently Grant Access?

The UK's AI Research Resource (AIRR) provides compute access through two routes, both ostensibly open to industry but with structural limitations that disadvantage startups:

**Gateway Route** offers first-time users up to 10,000 GPU hours over 3 months for testing and benchmarking algorithms. **Innovator Route** provides 50,000-150,000 GPU hours over 6 months for larger projects involving research teams and partnerships. While both routes technically permit "industrial research" and "experimental development," several features effectively exclude most startups:

1. Fundamental research is essentially restricted to research organisations only. This means industry applicants can only access AIRR for applied activities like feasibility studies or product development, not basic AI research, limiting the type of foundational work startups might pursue.
2. The allocation sizes are mismatched to commercial needs. Even the larger Innovator route caps at 150,000 GPU hours over 6 months. For context, training frontier models or developing commercial AI products often requires orders of magnitude more compute, making AIRR suitable primarily for academic-scale experiments rather than the compute-intensive work that would help British startups compete on the international stage.

In practice, AIRR functions as a programme primarily focused on serving academic research that happens to permit some industry participation, rather than a genuine pathway for startups to access the compute they need to build competitive AI products.

[21] The University of Edinburgh https://www.ed.ac.uk/news/university-set-to-host-ps750m-national-supercomputer

## AIRRPORT

The introduction of AIRRPORT offers the opportunity to change how industry, academia and researchers interact with the AIRR. The Government outlines its aims in this to create a seamless experience for users, enabling them to submit jobs, move workloads, access tools and manage data across AIRR systems. However, we remain deeply skeptical that a portal alone will solve the fundamental access issues that startups face.

The problem runs deeper than user interfaces. The entire allocation framework remains designed for academic researchers with long-term projects, not startups that need rapid, flexible access to test ideas and iterate at speed. Simply, the Government must make access as easy as possible for the startup community, a community that will drive economic returns on these investments and will often support the critical missions this Government has set out to achieve.

Clearly, the selection process for access must be transparent, merit-based, and aligned with broader policy objectives - and yes, these can be as broad as economic growth from which all else derives. AIRR should operate more like a competitive grant program, with clear criteria, regular application cycles, and published success metrics.

**We propose the creation of a dedicated access programme for startups and scaleups within the AIRR, with simplified onboarding and compute credits redeemable across supercomputers and approved providers for startups. This programme should be overseen by a specialist director, to provide the clarity of mission and leadership AIRR needs.**

This goes further than what the *Compute Roadmap* suggests – that a *'*portion of AIRR capacity will be reserved for emerging, high potential technologies – including novel chips, architectures, and software stacks'.[22] Reserving capacity will not be enough. The process itself will need to be simplified and offer more support. Researchers have the scaffolding of their institutions around them - startups don't. We must find a way to bridge that shortfall to ensure access is competitive. Again, we can look to the *AI Opportunities Action Plan* for a solution:

> "Strategically allocate sovereign compute by appointing mission-focused "AIRR programme directors" with significant autonomy. These could be modelled after the Defense Advanced Research Projects Agency (DARPA) or the Advanced Research and Invention Agency (ARIA) to quickly and independently provide large amounts of compute to high-potential projects of national importance, operating in a way that is strategic and mission driven. Allocation is an essential part of any compute strategy: spreading large amounts of compute thinly will have little impact. We will have to make choices about when to subsidise compute and when to provide it at cost, recognising that this could form part of an attractive offer to entrepreneurs and researchers deciding where to base themselves."[23]

We accept that done well this could likely result in issues around subsidy control; however, this presents an opportunity to capitalise on the flexibility we now have outside of the European Union. Current rules

[22] UK Compute Roadmap – https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

[23] UK AI Opportunities Action Plan https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan

would mean that if a startup had for example received an Innovate UK grant of £150,000 last year and now wanted to access the AIRR it would be limited to approximately £150,000 of GPU hours before subsidy control kicked in. This would limit the effectiveness of the AIRR for some of the UK's most promising AI startups. Government should be thinking carefully about how to create a new bespoke route to ensure AI startups aren't caught when attempting to action the UK Government's stated objectives of keeping the UK ahead in the AI race.

**Therefore we propose:**

1. **Create a dedicated startups access programme in AIRR.**
    a. This would simplify onboarding procedures and issue compute credits that were redeemable across both supercomputers and approved cloud providers.
2. **Create a competitive grant programme in AIRR.**
    a. This programme would have clear, published criteria for access, transparent and merit based selection processes informed by the Industrial Strategy set by the Government.
3. **Appoint mission focused AIRR programme directors.**
    a. Implement the *AI Opportunities Action Plan* recommendation by appointing AIRR "programme directors" with ARIA-style autonomy to: rapidly allocate compute to strategic, high-growth projects; exercise significant independent decision-making authority; and make strategic choices about when to subsidise versus provide compute at cost.
4. **Tackle head-on issues of subsidy control by creating bespoke new pathways.**
    a. The way in which subsidy control rules would kick in may mean that startups that have been supported by the Government in the last 2 years, for example through a grant, may be locked out of the AIRR. This runs counter to stated Government policy objectives. Therefore, DSIT should direct the CMA to ensure there is a streamlined route for AI startups that essentially acts as a block carve-out.

# Changing the Energy Environment

With an estimated 499 existing UK data centres, the UK is already the world's second-largest market and there seems no sign of stopping, with dozens more under construction.[24] All of these data centres require power. Data centres are estimated to use around 2.5% of the UK's electricity capacity today, with this demand projected to rise to as much as 6% by 2030[25]. Critically, access to affordable and reliable energy is a key enabler for new compute investment. Further, if the UK is to hit its ambitious climate commitments, this energy must be green.

To solve this, the Government has rightly identified the interconnectedness of compute and energy. The AI Energy Council was launched after the *Compute Roadmap* was published, with AI infrastructure planning explicitly tied to power access and sustainability. In November 2025, the Government announced that access to discounted energy would be a core offer of AI Growth Zones. Likewise, public investments in the AIRR supercomputers and node centres presume that new capacity will be backed by cleaner generation. Indeed, increasing sovereign compute capacity should not just accommodate the climate goals, but should accelerate them, as a key driver of demand for electricity.

## Grid Connections

In the same announcement, the Government stated that "timely connections to the electricity grid for AI data centres is the single biggest blocker for establishing AI Growth Zones".[26] This is critical: reinforcing the network can take from five to fifteen years, far slower than the pace of AI demand.[27] Recognising the need for faster connections across the economy, Ofgem has moved to a "first-ready, first-connected" regime: NESO now prioritises strategic clusters and aim to clear long queues of applications for data centre connections.[28] Data centres were designated 'Critical National Infrastructure' in 2024, securing access to fast-track consent.[29]

In practice, several UK-wide compute corridors are now emerging. Hyperscalers like AWS are committing £8 billion to UK data campuses over 1.3 GW in London alone, and AI Growth Zones will combine grid upgrades with planning support to distribute capacity outside the South East.[30] However, many parts of Britain see long waits for new connections or face firm limits on local output.

AI Growth Zones will receive special treatment for grid connections:

[24]Statistica https://www.statista.com/statistics/1228433/data-centers-worldwide-by-country/#:~:text=Leading%20countries%20by%20number%20of,located%20in%20the%20United%20Kingdom.
[25]Data Centre Review https://datacentrereview.com/2024/11/understanding-renewable-energy-use-in-data-centres/
[26] UK Government https://www.gov.uk/government/publications/delivering-ai-growth-zones/delivering-ai-growth-zones#section-two
[27]UK Government https://www.gov.uk/government/news/clean-energy-projects-prioritised-for-grid-connections
[28] Neso Energy https://www.neso.energy/news/our-new-approach-long-term-connections-reform
[29]UK Government https://www.gov.uk/government/news/data-centres-to-be-given-massive-boost-and-protections-from-cyber-criminals-and-it-blackouts
[30]UK Government https://www.gov.uk/government/news/chancellor-announces-8-billion-amazon-web-services-investment-as-she-vows-to-make-every-part-of-britain-better-off

- Two new connection mechanisms will be created: a reallocation mechanism that gives priority access to freed-up capacity, and a reservation mechanism that allows capacity to be held at a specific physical connection point for a strategically important project.
- The Government will work with Ofgem to enable viable options for AI Growth Zone developers to build their own high voltage grid infrastructure.
- AI Growth Zones will receive guaranteed support through the new Connections Accelerator Service.

Whilst these bespoke solutions are very positive, NESO's Centralised Strategic Network Plan (CSNP) does not yet explicitly focus on demand (and therefore economic growth), undermining the approach to prioritising grid connectivity outside of these specific zones. This has to change.

**Therefore we propose:**

1. **NESO's CSNP must adopt a new "demand" principle, accelerate its economic growth potential.**

## Abundant Green Supply

As demand scales, scrutiny of the environmental footprint of compute will only grow. Data centres use very large amounts of energy, so meeting net-zero goals will require abundant low-carbon energy supply through new renewables, energy storage, and nuclear capacity. Official policy already emphasises this link. For example, the *Compute Roadmap* calls for exploring "sustainable solutions such as renewables, advanced nuclear, and innovative grid solutions" to power AI.[31]

Practical examples of firms engaging with this challenge are already emerging. Major tech campuses are pairing with carbon-free supply, for instance Google's new UK data centre is tied to long-term renewable Power Purchase Agreements (PPAs) and battery storage to approach 95% carbon-free power.[32] We would like to see the Government exercising its unique convening power to better join up hyperscaler demand with energy producer supply, maximising the economic value of both.

Further, the Government's Flexibility Roadmap[33] highlights the critical role of controllable demand and distributed energy resources in balancing supply, integrating more renewables, and reducing curtailment. It identifies large energy-intensive sites, including data centres, as key participants in flexibility markets, capable of providing demand-response, frequency response, and other ancillary services to support the grid. However, current regulatory frameworks limit this potential. Backup generators at data centres are classified under the Medium Combustion Plant Directive (MCPD) and specified-generator regulations. While units used solely for on-site emergencies are exempt from certain permitting requirements, they are explicitly barred from providing balancing services or participating in demand-side-response schemes.[34] This misalignment prevents data centres from fully contributing to grid flexibility, despite the

[31]UK Government https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap
[32]New Energy World https://knowledge.energyinst.org/new-energy-world/article?id=139863#:~:text=Google%20is%20also%20investing%20in,research%20in%20science%20and%20healthcare.
[33] UK Government https://assets.publishing.service.gov.uk/media/68874ddeb0e1dfe5b5f0e431/clean-flexibility-roadmap.pdf
[34]UK Government https://www.gov.uk/guidance/specified-generator-when-you-need-a-permit

Government's Roadmap signalling their strategic importance. Regulatory reforms could unlock this potential by allowing data centres of all sizes to participate in flexibility markets, turning on-site generation and controllable load into active assets that help integrate renewable energy, reduce curtailment, and improve system efficiency.

On the advanced nuclear side, in November 2025, the Government announced the first three small modular nuclear reactors (SMRs), to be built by Rolls Royce, in Wylfa, North Wales, which is geographically adjacent to the AI Growth Zone which straddles the Menai Strait.[35] This is an excellent example of practical symbiosis between low carbon energy and compute build out, and should be emulated across the UK.

**Therefore we propose:**

1. **Across every element of Clean Power 2030, data centres must be viewed as a catalyst, in terms of reliable, high, and sustained demand for green electrons.**

## Abundant Cheap Supply

But for the UK to be the home of compute, green energy is insufficient. It must be cheap. Today, the UK's electricity prices are among the highest globally, roughly four times those in the United States and well above the median of comparable countries.[36] This price differential is not just an abstract economic statistic; it has very real consequences for energy-intensive industries, particularly technology and digital infrastructure. Running a data centre in the UK, for example, costs roughly four times as much as operating a similar facility in the US purely because of the cost of electricity.[37]

Since data centres form the backbone of AI development, cloud computing, and digital services, these elevated energy costs directly undermine the UK's competitiveness in the industries that will define the next wave of global growth. With this in mind, the Government announced in November 2025 that they would "develop a targeted pricing support mechanism to recycle grid constraint cost savings to projects that deliver measurable system benefits" such as data centres. Where data centres in AI Growth Zones facilitate these savings, the Government has announced that they will receive a discount on electricity costs. For a 500MW data centre this will be up to:

- £24/MWh in Scotland
- £16/MWh in Cumbria
- £14/MWh in the North East

The global race for technological leadership is accelerating, and these measures are positive. But unless the UK can bring its energy costs in line with international rivals, it will increasingly become a difficult,

[35]BBC News https://www.bbc.co.uk/news/articles/c051y3d7myzo
[36]Social Market Foundation https://www.smf.co.uk/wp-content/uploads/2025/01/How-to-power-AI-Feb-2025.pdf#:~:text=growth%2C%20particularly%20in%20certain%20parts,data%20centres%20to%20operate%20elsewhere
[37]House of Commons Library https://commonslibrary.parliament.uk/research-briefings/cbp-10315/#:~:text=Industry%20stakeholders%20and%20commentators%20have,for%20reforms%20to%20electricity%20pricing

and ultimately unattractive, place to build and scale the technologies of the future, including data centres situated outside of AI Growth Zones.

The introduction of zonal pricing could have incentivised the building of data centres close to excess supply, reducing curtailment payments, which are projected to top £3.7bn by 2030.[38] In the absence of zonal, however, there are still measures that the Government could take, and this starts with cutting the cost of levies on energy bills for business and households.

To address costs for businesses, successive Governments have applied targeted levy discounts to energy intensive firms. In 2024, the then Conservative Government unveiled a new "British Industry Supercharger", which exempts 500 of the UK's most energy-intensive businesses from renewable energy levies, and a discount on electricity network charges of 60%. This was increased to 90% in October by the Labour Government.[39] Alongside this Labour also introduced a new British Industrial Competitiveness Scheme alongside its Industrial Strategy, which exempts eligible businesses from paying for the Renewables Obligation, Feed-In Tariffs, or Capacity Market charges. This will cut industrial power costs by roughly £35-40 per MWh for about 7,000 firms in manufacturing sectors such as automotive, aerospace, and chemicals. This piecemeal approach, however, does not fix the underlying issue of high power, and merely spreads the cost across other energy consumers, including households and small businesses.

For data centres, energy bill levies disincentivise the introduction of local constraint markets, which could lead to cheaper energy in times of excess supply. In the near term, it is critical for the Government to consult on energy bill levies, to investigate their impact and possible avenues for more equitable, and efficient placement.

**Therefore we propose:**

1. **The Government must consult on reforming energy bill levies.**

## Green Data Centres

The most effective way to reconcile the growing demand for compute with the UK's climate goals lies in innovation within the data centre ecosystem itself. While policy and infrastructure reform are critical, the technologies running inside data centres and the ways those centres are designed, cooled, and integrated into their communities offer the most promising routes to sustainability.

Startups and innovators are already developing hardware and systems that drastically reduce energy consumption, recover waste heat, eliminate water use, and enable data centres to act as flexible grid assets rather than static power loads. The UK has a genuine opportunity to lead in green data centre technologies. This emerging ecosystem represents not just a response to environmental pressures, but a chance to build a globally competitive industry rooted in climate-aligned innovation.

[38] Octopus Energy https://octopus.energy/blog/stop-wasting-wind-power/;
[39] UK Government https://www.gov.uk/government/news/british-businesses-to-save-over-400m-a-year-as-government-slashes-electricity-costs

Here are a few startups which are providing the solutions for the current bottlenecks plaguing building compute.

## Case Study: Deep Green

Traditional data centres waste huge amounts of energy as heat, often vented into the atmosphere while local communities struggle with heating costs. **Deep Green has flipped this problem into a solution.**

Their modular, high-density data centres are designed to reuse heat energy by channeling it into community facilities such as swimming pools and district heating schemes. They are also waterless and ultra-efficient, avoiding the cooling issues that plague many operators.

By embedding AI and HPC workloads directly into local infrastructure, Deep Green creates two benefits at once: affordable, sustainable heating for communities and efficient, scalable compute capacity for any business that is using AI and HPC computing. This model directly addresses concerns that AI infrastructure drains local resources and alienates communities. Instead, it positions compute as a public good: a source of warmth, resilience and decarbonisation.

**If scaled, Deep Green's approach could redefine the social contract between AI infrastructure and the places it operates.**

## Case Study: Heata

The rapid expansion of cloud computing has created a hidden challenge: almost all the electricity consumed by data centres eventually becomes heat and most of that heat is simply expelled into the air. At the same time, millions of UK households face rising energy bills and an escalating fuel-poverty gap. Heata bridges these two crises with an elegant, decentralised solution.

**Instead of concentrating servers in vast facilities that require expensive, power-hungry cooling, Heata distributes compute units directly into homes.**

Each unit is a compact, secure server that attaches to a household's vented hot water cylinder. As it processes workloads, the waste heat is transferred into the water effectively delivering free or low-cost hot water as a by-product of cloud compute.

This approach tackles several systemic issues at once. Homes save £150 - £400 a year on energy bills, validated through independent testing by the Building Research Establishment. And

the carbon footprint of compute is dramatically reduced: by avoiding mechanical cooling and displacing domestic gas use, each MWh of computing can result in more energy avoided than consumed.

**By sending processing power directly to where heat is needed, Heata turns a structural inefficiency into a social and environmental benefit: lower bills, lower emissions, and a more resilient energy system.**

## Case Study: Yotter

Grid capacity has become a significant barrier to building new data centres in the UK. Most projects fail long before planning because developers cannot see where viable power actually exists leading to months-long delays, expensive DNO applications and land deals that collapse when connections prove impossible.

**Yottar solves this by giving data-centre developers instant visibility of where the grid can support new load.**

Using a constantly updated digital twin of the electricity network, Yottar provides real-time capacity signals, power-flow analysis and congestion modelling, allowing teams to test grid viability for any location in seconds.

This transforms site selection. Instead of spending thousands on speculative applications or waiting months for a DNO response, developers can immediately filter out constrained zones, prioritise strong locations and assess far more sites in less time. Many operators now evaluate 10x more options while dramatically cutting wasted application costs.

**For data-centre deployment, Yottar turns the grid from a bottleneck into a planning advantage, ensuring new facilities are built where power is genuinely available, not where developers hope it might be.**

## Case Study: Zendo Energy

**Zendo helps data centres navigate the growing unpredictability of AI-driven power demand and volatile energy markets that increasingly rely on intermittent renewable sources such as wind and solar.**
The company provides an intelligent Energy OS software platform that provides data, predictive analytics and simulation tools to help operators reduce energy costs and optimise capacity at their facilities. They offer a smart energy procurement service, aligning tariffs to real and forecasted computing loads, which recently helped a London data centre cut energy costs by around 30% while accessing cleaner power.

Zendo is developing advanced forecasting and capacity-planning tools that predict computing loads in order to unlock stranded capacity in facilities that often operate at only 50% utilisation despite severe grid constraints.

**In a market where energy costs can make up to 70% of e end customer bills and UK energy prices remain among the highest in Europe, Zendo offers a practical, scalable way to reduce costs, lower carbon and maximise returns on theUK's compute infrastructure.**

## Case Study: Chainergy

One of the biggest barriers to new AI compute in the UK is simply accessing power. Grid connections can take years, and most data centres face high energy costs and carbon-intensive supply.

**Chainergy turns this constraint on its head by colocating modular, AI-ready data centres directly at biogas plants.**

Instead of exporting electricity into a congested grid, Chainergy uses clean, existing on-site generation to power GPUs immediately, cutting connection delays entirely and reducing emissions by almost 70%.

This approach creates a uniquely sovereign and cost-efficient compute model: deployments can be built in under six months, power costs are significantly lower than traditional cloud providers, and customers such as AI Scale Ups, Enterprises, universities and public-sector organisations can access secure, privately networked GPU capacity without expanding their own estates.

Biogas operators benefit too, earning more for surplus energy while gaining upgraded digital infrastructure and optional heat-recovery systems that displace fossil fuel use on farms or in greenhouses.

**Chainergy demonstrates how next-generation data-centre design can solve multiple system challenges at once, accelerating AI infrastructure, strengthening local economies, and turning rural renewable assets into high-value compute hubs without adding pressure to the grid.**

# The Way Forward

The startup ecosystem stands ready to deliver the compute infrastructure Britain needs for the next industrial revolution. British founders are already building solutions to our biggest bottlenecks. The Government has committed the resources to back them, and is willing to work with a wide array of stakeholders to deliver the positive business environment needed. Now we must match this ambition with execution of the *Compute Roadmap* – moving faster, procuring smarter, building out the compute capacity at scale in the UK.

When it comes to changing the energy and planning environment, the UK cannot afford piecemeal or incremental steps. Data centres must be treated as nationally significant infrastructure, with streamlined approvals, guaranteed access to clean power, and clear integration into grid and nuclear expansion plans. At the same time, innovative models from startups such as Deep Green, Heata, Yotter Tech, Zendo Energy and Chainergy show that the compute sector itself can be part of the solution, making workloads flexible, reusing heat locally, and turning power demand into a grid-stabilising asset. By aligning planning reform, energy decarbonisation, and startup-driven innovation, Britain can build an AI infrastructure that is globally competitive, climate-aligned, and rooted in tangible community benefit.

The opportunity before us is equally stark. British startups are not waiting for permission to innovate. They are already designing the technologies which could solve the bottlenecks which plague building compute capacity. These founders are not asking for handouts – they are asking for customers, for access, and procurement processes which buy with the same urgency with which they build.

Data centres are a critical backbone of the AI-driven economy, but their growth must be managed so that it drives low-carbon transition rather than derailing it. Achieving this requires mission-oriented collaboration between Government and startups that are already providing solutions to concerns posed by data centres. The UK has many advantages; a strong service sector, a growing AI strategy, and an emphasis on clean power, but must act quickly to align policies. Implementing measures such as advanced network planning, flexible grid codes for large loads, site standards for efficiency, and market incentives for flexibility will all help. Capturing data centres' surplus heat and renewables not only cuts emissions but also creates local jobs. If the UK successfully tackles the challenges above through innovative cooling, on-site generation, and smarter grid management the data centre boom can fund and justify even more renewables and smart infrastructure. In short, by planning strategically now, the UK can turn potential grid stress into an opportunity: expanding digital capacity and accelerating clean energy, delivering sustainable growth for consumers and businesses alike.

The technologies outlined in this report, from Vaire's development of ultra-efficient chip architectures to Deep Green's use of distributed data centres that capture and repurpose waste heat for community heating, and Yotter Tech's real-time mapping of grid connection availability, illustrate that British startups are already addressing many of the core bottlenecks in the national compute stack.

The underlying message is simple: back British innovators with the same determination we court foreign investment. Britain's AI future will not be built by announcements. It will be built by the founders designing chips in Cambridge, the engineers optimising cooling systems in Bristol, the teams mapping grid capacity in real-time, and the policymakers willing to procure their solutions before competitors do. The Roadmap has set the destination. Now we must travel at founder speed.

STARTUP
C✳ALITION