BRITISH FRONTIER AI

British Frontier AI and Industrial Acceleration

A sovereign capability, built by a privately led laboratory.

43 min read · October 2026

Policy decisions at a glance

DECISION	PROPOSED SETTLEMENT
Primary mission	An independent British ability to develop, train and operate AI within a few months of the global frontier.
Second mission	Accelerate high-growth British industries where national capabilities and scientific assets create a plausible advantage.
Capability target	By year five: general capability within a few months of the closed frontier on independent measures, and leadership in selected scientific domains.
Market position	Research, train and license models. No mass-market consumer chatbot.
Institutions	A privately led British Frontier AI Laboratory and a separate, publicly controlled sovereign compute operator.
Research structure	A core frontier programme and four scientific divisions: life sciences, defence (in a ring-fenced subsidiary), space and materials.
Model policy	Proprietary control of strategic frontier assets. Older and smaller weights released selectively where this helps adoption without weakening sovereignty.
Ownership	Founders, employees and private investors hold the majority. The state co-invests as a minority on identical terms and appoints one non-executive director.
Sovereign controls	A protective special share, contractual continuity rights over designated models and the National Security and Investment Act. None depends on the size of the state's stake.
Founding finance	Up to £3bn of state equity over five years, invested only alongside private capital at the same time and on the same terms. A further £2bn in the accelerated scenario only.
Five-year illustration	£50bn: £30bn public and £20bn private. An organising scenario, not a calculated minimum.
Delivery principle	Public compute, sovereign contracts and continuity rights underwrite the sovereign minimum. Private capital leads the laboratory's equity from the founding round.

A decision that cannot be deferred. Ownership and sovereign rights must be designed together. This proposal secures sovereignty through protective rights and contracts, not ownership. The laboratory stays under private control, which preserves its autonomy, its appeal to talent and investors and its classification outside the public sector.

A distinction to keep throughout. Access to a competent model, the ability to train the next generation independently and standing at the leading frontier are three different objectives. The programme must measure each separately.

Executive summary

Britain should establish a permanent, British-controlled frontier AI capability, delivered by a privately led laboratory and backed by national compute. The aim is not a domestic ChatGPT or the maximum value of one company. **It is to retain strategic agency in a foundational technology and to speed up scientific and industrial innovation where Britain already has strengths.**

The Council for Science and Technology's advice to the Minister for AI, dated July 2026 and published in October, defines AI sovereignty as agency across the AI value chain. It recommends targeted domestic development alongside continued access to international models [1]. The Alan Turing Institute has described a medium-sized sovereign model and a portfolio approach [2]. This proposal goes further: a national laboratory that trains successive model generations within a few months of the closed frontier and leads in selected scientific domains.

There are two institutions. A publicly controlled compute operator secures sites, accelerators, networking, storage and power. It also hosts designated sovereign models as a last resort. The British Frontier AI Laboratory recruits researchers, trains models, owns its intellectual property and licenses models to government and industry.

The laboratory has a core research group and four specialist divisions sharing one technical stack: life sciences, defence, space and materials. Defence sits in a ring-fenced subsidiary so its security requirements do not burden the rest of the company. NHS use is a demand case within life sciences.

The laboratory is privately led. A private lead investor and the founders hold the majority from the outset. The state invests in each round at the same time and on the same terms as private investors. It holds a minority stake (about 20% at founding, falling to about 15% by year five in the illustration) and appoints one non-executive director. Founders and employees get substantial vesting equity and real operating freedom. Liquidity comes from secondary sales and a later listing. A foreign takeover of the strategic core is prohibited.

Sovereignty does not depend on the size of the state's stake. It rests on three layers. A **special share** gives protective vetoes over change of control, disposal of designated assets, relocation of core capability and concentrated foreign ownership, with no power to direct strategy. **Contracts** give government a lasting licence to run designated models on UK-controlled infrastructure, backed by tested escrow. **Statute**, chiefly the National Security and Investment Act 2021, is the backstop. The design aims to keep the laboratory outside the public sector. That cannot be guaranteed, so a provisional view from the Office for National Statistics (ONS) should be sought before the founding round [24, 25].

The five-year illustration is £50bn: £18bn of public compute and operations, £9bn of sovereign research contracts, up to £3bn of state equity and £20bn of private equity. Private capital is a launch condition for the laboratory but not for the sovereign floor. If private rounds stall, state equity stops too. Compute, contracts and continuity rights carry on, and Parliament decides what happens next. Over ten years, including hardware refresh, the public commitment is roughly £55–70bn.

Much of what frontier leaders spend pays for being first. OpenAI's reported research and development spending rose from $7.81bn in 2024 to $19.18bn in 2025, separate from the cost of serving customers [5]. Most of its 2024 research compute appears to have gone on experiments and unreleased models rather than final training runs [6]. The compute needed to reach a given capability is falling fast [8], and the best open-weight models have trailed the closed frontier by about four months since January 2026 [7]. A capable follower can stay close for a fraction of the leaders' spending. Leading outright in general capability would cost far more, so this proposal aims to lead in selected scientific domains instead.

The case does not rest on a tax windfall. **The value is strategic independence, an option on future breakthroughs and higher productivity and intellectual-property creation in British industry.** Extra gross value added is not tax revenue, and it must be measured net of innovation that would have happened anyway. The cost should be judged as an insurance premium against a defined loss of access, not as an investment that repays the Exchequer.

A six-month foundation phase would cover independent costing, recruitment of founders and a private lead investor, a review of existing British laboratories, legal design of the sovereign controls, a provisional classification view, grid and site options and a staged mandate. Government should not commit to a vehicle or architecture before testing whether an existing British team could deliver faster.

1. The strategic problem

1.1 Frontier AI is not just another software market

Frontier models increasingly underpin research, engineering, programming, analysis and autonomous work. Whoever controls their development influences access, price, licensing, release and research direction. In defence, healthcare and critical infrastructure, dependence on another country's suppliers is a strategic weakness. Foreign suppliers are not unreliable, but access by contract is not the same as independent capability.

A country can own data centres yet lack the models, know-how, training software and teams to build new systems. It can own model weights yet depend on foreign clouds, export-controlled chips and a few key specialists. Sovereignty must be defined by what a country can actually do, not by where a company is incorporated.

The Council for Science and Technology rejects full self-sufficiency in favour of selected capabilities and resilient access [1]. This proposal accepts selective dependency but gives more weight to domestic frontier development as a permanent institution.

1.2 Britain's scientific position

The UK helped create DeepMind and has world-class universities, computational science, biotechnology and engineering. That shows a strong base, not that success can be reproduced on demand. In April 2026 David Silver's London laboratory Ineffable Intelligence raised $1.1bn at a $5.1bn valuation, led by US investors with the government's Sovereign AI Unit and the British Business Bank taking part [12, 13]. British-led teams can attract global capital. The same example shows how quickly their ownership becomes international.

The harder problem is retention, capital and scale. A laboratory near the frontier must fund failed experiments, hardware refresh and repeated model generations. National support addresses that coordination problem while keeping entrepreneurial management.

1.3 Why buying foreign models is not enough

Buying foreign models is the sensible default for most uses where sovereignty is not at stake. Forcing British firms to use a domestic model would undermine the industrial mission. The programme should compete with international systems, not hide from them.

But procurement alone does not build training expertise, control of model weights, the right to modify systems or access during a dispute or crisis. Nor does it build the foundations for frontier scientific models. These are what the programme buys.

1.4 Open-weight progress changes the starting point

Leading open-weight models are a strong base for experimentation, adaptation and training. Since January 2026 the best of them have trailed the closed frontier by about four months on Epoch AI's capability index [7]. Stanford's 2026 AI Index shows a similar picture [3]. Separate analysis puts the fall in compute needed to reach a given capability at between 16 and 60 times a year over 2023–2025, with wide uncertainty [8].

This is the strongest argument that the programme is affordable. It comes with two cautions. Followers benefit from leaders' published research and open releases, and catch-up slows if those dry up. And adapting an open model is not proof of the capacity to build the next one. The programme should use open models heavily at first while building its own training capability, and report the two separately.

2. Two missions, measured separately

2.1 Mission one: permanent frontier capability

The first mission is for Britain to develop, train, update, deploy and operate frontier AI for critical national purposes independently. That needs legal rights, working technology, trained people and installed compute. Owning shares in a valuable company does not meet it.

Five measures apply:

- unrestricted domestic operation of designated strategic models
- recurring training of independently originated advanced systems
- a secure, UK-controlled compute stack

- retention of critical knowledge and deployment artefacts
- general capability within a few months of the closed frontier on an independent index, and leadership in at least one scientific domain on prospective tests.

Resilience under duress is a sixth test. Regular exercises should run designated models without foreign vendors, the laboratory's management or overseas control planes, covering patching, security, retraining and recovery as well as inference. Full independence from foreign chips is not credible or proposed. The aim is operational control and a plan for disruption (section 8.4).

2.2 Mission two: industrial acceleration

The second mission is to combine British scientific and industrial strengths with powerful AI to speed up discovery, design, manufacturing and exports. Most of the gains should go to downstream British firms and research bodies, not the laboratory.

That means broad licensing, deployable weights where appropriate, standard APIs, technical support and clear enterprise terms. It does not mean subsidising every application company or imposing one supplier. Evaluation must separate the laboratory's own sales from the value its customers create, and measure patents, products, trials, designs, exports and productivity net of what would have happened anyway.

2.3 What the programme will not do

It will not build a consumer chatbot, require UK firms to buy its models or set up a separate venture fund. It will not guarantee investor returns or the company's solvency. It will not assume access to clinical or defence data because the state is a shareholder and customer.

3. Two institutions

3.1 Sovereign Compute Operator

The first institution is a publicly controlled, ring-fenced compute programme, built where possible on existing national capability such as the AI Research Resource and Isambard-AI [10, 11]. It controls data centre sites, networking, accelerators, storage, cooling, power and critical systems administration.

It runs designated national models, provides capacity to the laboratory on published terms, maintains secure environments, plans replacement cycles and hosts sovereign models as a last resort. It may contract out construction and routine operations but must keep enough rights and expertise that no foreign supplier can switch off essential systems.

It does not set architecture, appoint researchers or run its own research. It is judged on availability, throughput, cyber resilience, cost per useful computation, commissioning speed, utilisation and power. As a public body its spending scores in the public accounts.

3.2 British Frontier AI Laboratory

The second institution is a UK-incorporated, privately controlled research company. Management controls research, hiring, partnerships, licensing and spending. A core frontier-model team and four divisions share one platform, with defence in a ring-fenced subsidiary. The company owns and licenses

its models, may make significant profits and may raise capital within the investor rules in Appendix C.

Private control is deliberate. It keeps the laboratory free of public pay scales, grant-by-grant approval and ministerial choice of architecture, so leaders can build teams fast and drop failing approaches. The state's interests are protected by the special share, the sovereign contracts and its board seat.

3.3 Existing vehicles and programmes

A new company is not an end in itself. Formation should test three routes: a new company built around an exceptional team and private lead, a co-investment and sovereign-rights deal with an existing UK laboratory or a new company combining a team with an existing research or infrastructure operation. An existing laboratory may bring know-how and a working stack, but its investors or contracts may make the sovereign terms unworkable. The test is capability per public pound, speed to scale and enforceable rights.

The programme must also fit with the existing £500m Sovereign AI Unit, which already invests in UK AI companies, provides compute on Isambard-AI and backed Ineffable Intelligence [13–15]. The two should not compete for the same teams. The state's stake in the laboratory should be held by an arm's-length investor on commercial terms, either the Sovereign AI Unit or the British Business Bank. This also supports the classification design in section 6.

A shortlist should be assessed confidentially, with competitive challenge from credible consortia and an auditable public-value assessment.

4. Research strategy and divisions

4.1 Core frontier programme

The core programme covers reasoning, mathematics, coding, multimodal understanding, planning, scientific representation, training efficiency and model operating systems. It should stay architecture-neutral, leaving founders free to pursue whatever approaches the field rewards.

Early work should use open-weight models and established libraries. In parallel the laboratory must build its own datasets, pipelines, distributed training stack, evaluations and original pre-training, recording which models derive from third-party weights and which are its own.

Open-weight models need a **provenance policy**. Many of the strongest come from jurisdictions whose interests may diverge from the UK's. Third-party weights should be security-assessed and checked for licence restrictions. They stay out of defence and other designated deployments unless approved.

The key output is not a flagship release. It is the ability to produce and improve successive models at competitive cost, with the state able to run specified versions independently.

4.2 Life sciences

Britain's strongest potential advantage in biomedical AI is the NHS. Few countries combine near-universal coverage, a population of almost 70 million, longitudinal records, leading genomic research and a large pharmaceutical industry. The value lies in linking diagnoses, treatments, imaging, pathology, genomics and long-term outcomes. The US has larger datasets but a fragmented system.

The Foresight project shows both the opportunity and the risk. It trained a generative model on de-identified records of about 57 million people in England inside the NHS England secure data environment. NHS England paused it in June 2025 after GP representatives objected that COVID-19 research data had been used without their knowledge [18, 19]. Access to population-scale data can be lost quickly if consent, governance and professional engagement are not settled first.

With UK Biobank's 500,000 whole-genome sequences, Genomics England and the NHS Genomic Medicine Service, Britain can build a globally competitive biomedical AI capability [17]. The division would develop models for molecular representation, protein engineering, genomics, chemical synthesis and disease prediction, supporting drug discovery, trial design and personalised medicine. The goal is for Britain to be a leading location for AI-enabled drug discovery, not only to improve NHS services.

This needs negotiated access, not an assumption that NHS data sits in one place. Health is devolved and records are held by many bodies. The programme must work through secure data environments and approved custodians, with independent oversight and clinicians and patients involved from the start. Intellectual property in a new medicine should normally belong to the inventor or research consortium, with the laboratory paid for its models and contracted work.

4.3 Defence

This division will build secure models for intelligence fusion, mission planning, simulation, engineering, logistics and autonomous platforms, within legal command responsibility and classified security arrangements. Decisions involving force stay under human authority.

The Ministry of Defence's June 2026 AI policy describes wider access to frontier models across classifications, with modular procurement and assurance [20]. The division should supply components the Ministry can integrate and test in realistic exercises.

Defence work sits in a **ring-fenced subsidiary** with its own vetting, information barriers, site security and export-control compliance. This keeps defence requirements from slowing other divisions, allows tighter government rights over the subsidiary than the parent and lets investors with defence exclusions invest in the parent. The subsidiary licenses the core stack from the parent.

4.4 Space

This division will work on Earth observation, onboard inference, spacecraft operations, autonomous mission planning and constellation coordination, under the constraints of radiation, delay, limited power and intermittent links. The UK space sector reports £18.6bn of income and 55,550 direct jobs [21]. The aim is to help British firms win global satellite markets, working with manufacturers, operators, analytics firms and universities. Mission and engineering tests matter more than language benchmarks.

4.5 Materials and advanced engineering

This division will cover chemistry, condensed matter, materials discovery, process engineering, composites, semiconductor design and industrial robotics, linking machine learning to simulation, experiment and manufacturing. The 2026 semiconductor study identifies 703 UK companies in the sector, with opportunities in photonics, packaging, power electronics and design. UKRI's National Materials Innovation Programme links discovery to adoption [22, 23].

The priority is closing the loop from computational proposal to measured result: robot-assisted experiments, inverse design, uncertainty-aware prediction and manufacturable process recipes. A material that cannot be made economically at scale is not an industrial success.

4.6 Allocating resources

Divisions should not get equal fixed budgets. The core team and shared training systems will absorb most early resource. Specialist spending should follow reviewed opportunity, data and test facilities, strategic importance and commercial uptake. Each division has a scientific director and a small industry advisory group that shapes evaluation but does not set the agenda or claim the intellectual property.

4.7 Safety, security and release

A state-backed frontier laboratory must meet the highest safety standards, and its specialist models carry dual-use risks. The laboratory should:

- give the AI Security Institute pre-deployment access for independent capability and safety testing
- test life sciences and materials models for biological and chemical misuse before release, with access controls and staged release where needed
- publish a release policy setting out which evaluations govern the release of weights
- protect its weights in proportion to their value, manage insider risk and commission regular independent red-teaming.

Safety leadership is itself a sovereign asset. It strengthens the UK's voice on international standards and gives government confidence to use the models in sensitive settings.

5. Commercial model and intellectual property

The laboratory sells to organisations, not consumers. Revenue comes from hosted APIs, enterprise inference, on-premises and sovereign-cloud licences, specialist model subscriptions, paid co-development and sovereign contracts.

Customer data and inventions do not transfer to the laboratory by default. Standard terms should separate the laboratory's background models, the customer's data, fine-tunes and new downstream intellectual property. Customers normally keep their data and outputs. Jointly funded research can agree its own terms in advance.

Strategic weights and the training stack stay proprietary, subject to the state's rights. Releasing smaller, older or research models can help adoption and recruitment. Release is a commercial and strategic judgement, governed by the safety tests in section 4.7.

Licensing must be predictable and non-discriminatory within security and export rules, and the laboratory must not use public compute to dominate downstream British markets. An independent review should compare its terms with international suppliers periodically. British firms remain free to choose better foreign models.

5.1 Pricing sovereign contracts

Sovereign contracts should be open-book and priced on an economically significant basis: allowable costs plus a regulated profit rate, modelled on the Single Source Contract Regulations used for non-competitive defence work [32]. This gives the taxpayer transparency and the laboratory a defined margin on its anchor revenue. Commercial use of public compute is charged separately on published terms, so nothing is counted twice.

5.2 What the investment case rests on

Sovereign contracts are valuable anchor revenue, but at a regulated margin they will not carry the valuations in Appendix A. As an illustration only: at 20 times forward revenue, a £60bn valuation implies about £3bn of annual revenue. Year-five sovereign contracts would supply £2.25bn of that at a modest margin. The valuation therefore rests mainly on commercial licensing, especially specialist scientific models where the laboratory can lead. The ten-year model (Appendix E3) must show sovereign and commercial revenue separately, with the margin on each.

6. Ownership and sovereign control

6.1 The economic settlement

The laboratory is a privately controlled company. A private lead investor and the founders hold the majority from the founding round. Founders and employees receive substantial vesting equity and later investors subscribe for ordinary shares. Percentages should not be fixed in policy before the team, lead investor and valuation are settled.

The state invests in each round **at the same time, at the same price and in the same instrument** as private investors. It always holds a minority of both economic interest and votes: 20% after founding and about 15% by year five in Appendix A4. Investing on identical terms alongside a larger private investor is what makes this a commercial investment rather than a subsidy (sections 6.5 and 6.6). The state takes no preference or rights beyond those of private investors in the same round.

Founder and employee incentives are set by the company and its lead investor, not by the state accepting worse terms. Founder equity vests against contribution and retention, and senior awards refresh as the company grows. The offer must make sense to top-tier talent.

6.2 Exit and liquidity

Liquidity comes from approved secondary sales, a later listing and, once the company is cash-generative, dividends or buybacks. Investors can sell economic interests without transferring the sovereign core. A listing is compatible with the special share if designed properly. A foreign takeover of the strategic core is prohibited. Frontier laboratories realise value through listings rather than trade sales, so this costs investors little.

The state does not guarantee investor returns or buy back shares. If the company fails, private money is lost. The continuity and compute arrangements preserve sovereign capability without propping up the company. Secondary sale rules should allow real wealth creation through orderly windows, transparent pricing and security checks on buyers. Overly tight approval rules would damage recruitment.

6.3 The state's board seat

The state appoints **one non-executive director** with the same duties as any other director, and has no other appointment rights. Government's requirements as a customer are set through the sovereign contracts, not board instruction. The public strategic board (section 11) oversees the compute operator, the contracts and the state's shareholding policy. It does not direct the laboratory.

6.4 The sovereign settlement

The state's protections work in three layers. None depends on the size of its stake.

The special share gives consent rights over specified actions only: change of control of the company or a critical subsidiary, sale or exclusive licensing of designated sovereign assets, moving designated research or compute abroad, any foreign holding above a set cap, rights that would block independent UK use and switching off the continuity arrangements. It gives no power over strategy, budgets, pricing, appointments or research. UK defence and aerospace companies use protective special shares of this kind and remain in the private sector. Appendix C1 sets out the full schedule.

The contracts carry the operational substance. The laboratory grants the state a lasting licence over designated models and deployment materials. On agreed triggers, government receives working weights, tokeniser and software artefacts, architecture descriptions, evaluation records and documentation, with the compute operator as host of last resort. Source and training artefacts are escrowed and tested regularly. Because the state pays for these rights, they are a customer's rights over defined assets, not an owner's control of the company. They cannot confer rights to third-party data the laboratory does not hold.

Statute is the backstop. The National Security and Investment Act 2021 covers acquisitions in AI, computing, defence and related sectors, including investments in each round [27]. Export controls and vetting apply to the defence subsidiary, where the state can hold tighter rights.

Government must prove in live exercises that it can recover and run designated systems. A special share without a working model and skilled operators is not independence.

6.5 Public sector classification

The ONS places a body in the public sector if government controls its general policy. Indicators include voting majorities, board appointments, special shares, contracts and the degree of government financing [24]. A publicly controlled body that does not cover more than half its costs from sales at economically significant prices is classed as general government. Sole-source supply to government normally fails that test [25]. Government equity that covers expected losses counts as spending unless private investors take a significant share at the same time and on similar terms. Eurostat's 2018 AREVA ruling rejected private participation that was smaller, later and conditional [26].

A majority state stake would almost certainly fail these tests, because early spending goes mostly on the laboratory's own research, which does not count as sales. So the design avoids public control altogether:

- a minority state stake with one non-executive director
- a special share with protective rights only
- every state subscription made alongside a larger private one, at the same time and on the same terms

- sovereign contracts priced on an economically significant basis, with no guarantees or implied bailout.

The ONS judges in the round and nothing is guaranteed. The residual risk is that the special share plus heavy early reliance on government revenue is judged to amount to control. The ONS can give a provisional view on a new entity's business plan [25]. That view should be obtained before the founding round and published.

6.6 Subsidy control and other legal requirements

Co-investment on the same terms as a significant private investor is strong evidence that the state's equity is not a subsidy. Compute pricing, preferential access and any contracts not on market-equivalent terms still need assessment under the Subsidy Control Act 2022 [28]. The founding documents must also deal with procurement, competition, data protection, export control and intellectual property.

7. Financing

7.1 Financing principle

The state underwrites the **sovereign minimum** through public compute, sovereign contracts and continuity rights. None of these depends on private fundraising. The **laboratory's equity** is led privately from the founding round, with the state as minority co-investor. Private capital is therefore a launch condition for the laboratory but not for the sovereign floor.

State equity never substitutes for private capital. If no credible private lead emerges, or later rounds stall, government chooses between three routes. It can contract the sovereign capability from an existing UK laboratory (Appendix E2, Option D). It can narrow ambition to the sovereign floor. Or it can ask Parliament for a publicly owned laboratory with statutory freedoms over pay and procurement, as the Advanced Research and Invention Agency has. That last route would sit in the public sector and be budgeted in full.

The central assumption is up to £3bn of state equity over five years. A further £2bn applies only in the accelerated scenario, on the same basis, and is outside the £50bn central case.

7.2 Five-year central scenario

FUNDING SOURCE AND USE	FIVE-YEAR ILLUSTRATION	STATUS
Public compute, facilities and operations	£18bn	Public infrastructure
Sovereign capability contracts	£9bn	Public purchase of defined capabilities
State equity in the laboratory	£3bn	Minority co-investment alongside private rounds only
Private equity over successive rounds	£20bn	Private lead from founding. Later rounds not committed
Total	£50bn	Illustrative scenario

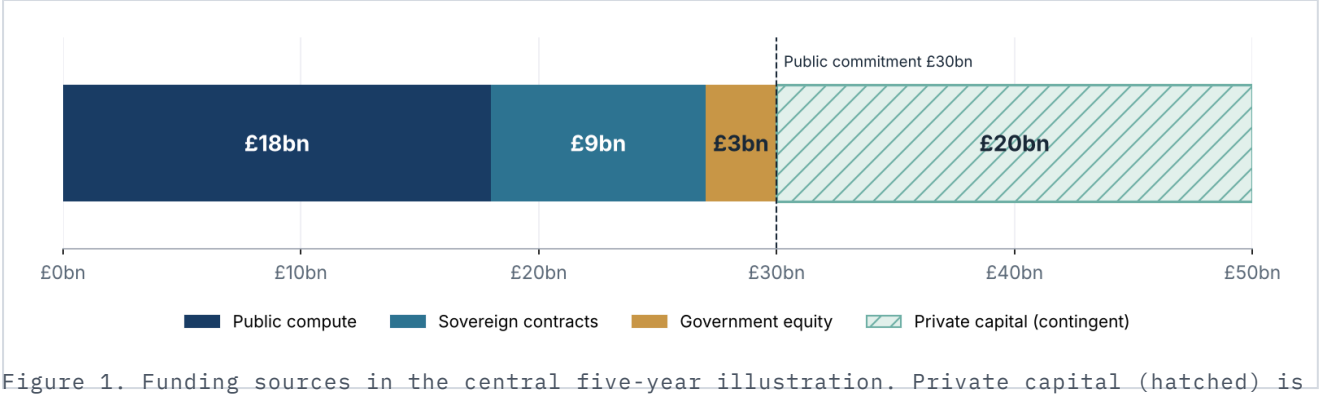


Figure 1. Funding sources in the central five-year illustration. Private capital (hatched) is contingent. State equity is invested only alongside it.

The table doubles as a high-level use of funds. In company accounts, though, equity is financing and contract revenue is operating income. The integrated model must turn both into actual hardware, staff, data and facilities costs.

The public commitment is £30bn, not £50bn. The £20bn private target depends on markets and technical progress. A shortfall does not remove the state's duty to maintain the floor it has bought, but it may limit the commercial expansion and frontier ambition.

7.3 Fiscal treatment

Subject to the ONS view, the elements would be recorded as follows:

- **Compute operator:** public capital and operating spending, scoring in public sector net investment and against borrowing and debt, offset by the assets it owns.
- **Sovereign contracts:** government purchases, scoring as public spending.
- **State equity:** a financial transaction if made on the co-investment terms above, broadly neutral for public sector net financial liabilities. If treated instead as a capital transfer, it scores as spending.
- **Laboratory:** if classed as private sector, none of its spending scores, including spending funded by private investors.

7.4 Avoiding a hidden subsidy

The compute operator must record the economic cost of capacity used for sovereign contracts. Each contract must specify the models, rights and knowledge transfer the state receives. Commercial work is paid for by the laboratory or its customers on published terms. The laboratory must not be paid twice for the same public capacity. Sovereign purchases must be open-book and benchmarked regularly (section 5.1).

7.5 Ten-year commitment

The five-year envelope sits within a ten-year plan covering hardware refresh, facilities, training campaigns, licensing revenue and changing national needs. Years six to ten would need roughly £5–8bn a year of public spending on compute refresh, operations and sovereign contracts. That takes the ten-year public commitment to about **£55–70bn**, to be validated in the ten-year model (Appendix E3).

This is an insurance premium. Even the high case in Appendix E1 does not recover the ongoing cost through tax. The programme is justified only if government values independent capability and its industrial benefits at roughly £5bn a year or more. The business case must make that judgement explicit by defining the loss of access it insures against.

An independent review at year three should decide whether to scale, maintain, partner or restructure. A year-five review recommends the next settlement. The state commits to a capability, not to indefinite revenue for one company regardless of performance.

8. Compute, energy and infrastructure

8.1 The technical challenge

UK public compute includes Isambard-AI in Bristol, with 5,448 NVIDIA GH200 superchips, and further systems in Cambridge. The Compute Roadmap expands the AI Research Resource, a base to build on but not a guarantee of dedicated frontier capacity [10, 11]. The laboratory needs large, tightly coupled training clusters with fast interconnect, storage and checkpointing, secure environments and room to run many experiments at once. Procurement should be performance-based and measured in sustained useful throughput, not accelerator counts.

8.2 Interim capacity and build-out

Research should start before permanent facilities are ready, using existing UK resources and secured contracted capacity as a bridge. The operator should buy in modules with phased acceptance rather than five years of hardware at once, and stress-test utilisation, supply risk, obsolescence and replacement costs.

8.3 Electricity is on the critical path

Data centre sites are limited by grid connections, generation, cooling, land, fibre and planning. The government's AI Growth Zones guidance names connection delays as a major obstacle [29]. Buying accelerators before securing a connection risks billions in idle assets. Energy planning should run alongside national grid and generation planning from the start, and site selection should score connection readiness, resilience, water use, security, build time, environmental impact and room to expand. For scale, the 120,000-accelerator example in Appendix B3 implies about 300MW of facility load.

8.4 Hardware sovereignty means control, not autarky

The UK will keep relying on global suppliers for advanced chips. The 2026 AI Hardware Plan identifies scope for domestic contribution in photonics, edge inference and advanced design [30]. The operator should use trusted suppliers, resilient control planes, auditable firmware and contingency procurement. UK-owned data centres do not remove chip dependency.

The most serious threat is not losing a cloud or software vendor but an exporting government restricting the next generation of accelerators. Installed hardware would keep running, but capability would fall behind over two or three refresh cycles. The programme needs a plan covering how long designated

capability lasts without refresh, whether to hold a strategic hardware reserve, which alternative suppliers are realistic and what allies could provide.

9. Founders and talent

A laboratory at this level is a concentration of exceptional people with exceptional resources. The founding team should include an entrepreneurial chief scientist, a research chief executive and experienced distributed-training leaders, recruited together and free to hire internationally on private-sector terms.

The offer combines vesting equity, real secondary liquidity, long-term research funding, competitive pay, top-tier compute and the chance to lead an institution of historic importance. Researcher autonomy and scientific standards are core inputs, not perks.

Government should also back research collaboration, university secondments and a pipeline of systems and scientific-computing engineers. Without strong platform engineers, scientists spend their time debugging infrastructure. The divisions should recruit both AI specialists and domain scientists, and work with universities in ways that protect academic independence and lawful data use.

10. Delivery

Phase 0: authorisation and founding (months 0–6)

Appoint a small independent formation board with expertise in frontier AI, scientific computing, finance, competition and national security. It should define the sovereign rights and special share and seek a provisional ONS view and a subsidy assessment. It should open confidential talks with founding teams, lead investors and existing laboratories. It should also agree who manages the state's stake, reserve immediate compute and commission independent costing. Parliament should see the public envelope, the ten-year estimate and the classification view before any irreversible commitment.

A multi-year mandate must be available during recruitment so a new team is not funding its first experiments through annual approvals. The founding round goes ahead only once the team, a private lead on agreed terms and the sovereign documents are in place.

Phase 1: build and train initial models (months 6–18)

Recruit research leadership and the training-systems team. Adapt approved open-weight models, build the training stack and evaluations, set up AI Security Institute testing and start original pre-training. Order compute in phases, only where power and deployment milestones are credible. Run a few genuine industrial pilots aimed at learning, not showcase demonstrations.

Phase 2: independent training (months 18–36)

Train progressively more ambitious original models, show reproducible development and secure deployment, bring the first UK-owned cluster into service and launch enterprise and specialist pilots. Later private rounds expand where results support valuations.

Phase 3: frontier-scale expansion (years 3–5)

Add UK-controlled compute, run repeated training generations, grow the divisions, build licensing revenue and hold formal continuity exercises. The year-five review sets the next commitment on technical standing, resilience and industrial evidence.

Phase 4: maintain and renew (years 6–10)

Continue research, replace hardware, renew energy and infrastructure and develop new model generations. The programme must survive changes of partner or leadership without losing the sovereign core, and adapt if new techniques make today's architectures obsolete.

11. Governance and accountability

11.1 Ownership in government

The Council for Science and Technology recommends a single government lead for AI sovereignty [1]. The Department for Science, Innovation and Technology should own the programme through one senior responsible owner, with the Ministry of Defence co-sponsoring defence contracts. Accounting officer responsibility must be explicit for the compute operator and the sovereign contracts. The National Audit Office should have access to the operator, to contract records (written into the contracts) and to the management of the state's stake. The laboratory is audited as a private company.

11.2 Boards

A small public strategic board oversees funding, compute priorities, sovereign requirements and shareholding policy. It does not review research decisions. The laboratory's board, including the state's non-executive director, is responsible for research, people, private capital, accounts and licensing.

11.3 Performance and oversight

Reporting should separate technical capability, sovereignty, industrial adoption and financial stewardship. Bigger models or more accelerators are not success in themselves. Science should be judged by third-party evaluation, an independent capability index, prospective domain experiments, cost, reliability and security. Industry should be judged by uptake, downstream intellectual property and commercial results.

Major commitments need stage gates (Appendix D), but scientists must be able to change direction quickly within an agreed envelope. Payments tied to a set architecture or benchmark invite gaming. Oversight should include red-team security testing, national-security reviews, external scientific evaluation, open-book contract accounts, a hardware asset register, talent retention reporting and annual public accounts with classified annexes where needed.

12. Economic and fiscal assessment

12.1 What the taxpayer buys

The state buys three separate things: independent capability and resilience, an option on future advances and the chance of faster private innovation, with a return through its shareholding. These should not be merged into one figure. Compute is a depreciating physical asset. Sovereign contracts buy rights and knowledge. Equity is a financial asset. Each is accounted for differently.

12.2 The counterfactual

The comparison is not with an economy without AI. British firms can buy international AI services without a national laboratory. The programme adds value only where domestic models give extra access, better scientific performance, lower cost, security or continuity. Spillovers from domestic expertise add further value. Appraisal should compare the options in Appendix E2 on cost, dependency, industrial output and foregone alternatives, including the chance that global capabilities become cheaper to license.

12.3 Measuring returns

£1 of extra gross value added is not £1 of tax. The right method estimates attributable UK GVA, adjusts for displacement and leakage, applies a defensible fiscal conversion, discounts the result and sets it against costs, renewal and opportunity cost. Appendix E1 shows that the case rests on resilience and option value rather than fiscal recovery. Hence the insurance framing in section 7.5.

12.4 Research budget priorities

Government plans £86bn of public research and development spending over the 2026–2030 spending review period, including £38.6bn for UK Research and Innovation [9, 31]. That is scale, not spare cash. Much of it funds institutions the laboratory will rely on. At its peak the programme would cost a large share of UKRI's annual budget. The submission must separate reallocated money, existing AI and compute commitments and genuinely new spending. Basic science should not be raided to fund model training, since that would weaken the ecosystem frontier advances come from.

13. Decisions required

Government is asked to take five decisions in sequence:

- **Adopt the two missions and the capability target**, and authorise the foundation phase.
- **Adopt the ownership model:** a privately led laboratory with the state as minority co-investor on identical terms. The state appoints one non-executive director. Sovereignty is secured by special share, contracts and statute.
- **Approve the founding envelope** once a founding team, private lead, independent compute estimate and provisional ONS view are in place.
- **Establish the sovereign settlement:** special share, continuity licence, escrow and any statutory instruments, with subsidy analysis.

- **Approve the phased compute plan** and its grid-connection critical path.

Government should publish a short annual statement covering what Britain can run independently, what the laboratory has built, its standing against the frontier, public money committed, private capital raised and evidence of industrial gains.

A policy that outlasts changes of government needs a legally durable core and financing that does not depend on ministerial enthusiasm for each training run. Durability is not immunity from scrutiny. The laboratory must keep earning confidence through results.

Technical Appendix A: Integrated financing model

A1. Status

All amounts are illustrative nominal pounds over five years. They are not quotations, commitments, budgets or valuation forecasts. They show the public and private split, the sources and uses of funds and the contingent obligations. A due-diligence model must replace them with real staffing, architecture, facility, energy, tax and subsidy figures before approval.

The central total is £50bn: £30bn public and £20bn private. The public share is £18bn for compute and operations, £9bn for sovereign contracts and £3bn of equity. The equity is a financial asset bought alongside private investors, not payment for a service. Private equity funds research and growth not already paid for by government contracts.

A2. Illustrative annual cash profile (£bn)

PROGRAMME SOURCE	Y1	Y2	Y3	Y4	Y5	TOTAL
Public compute and operations	2.00	3.00	4.00	5.00	4.00	18.00
Public sovereign contracts	1.00	1.50	2.00	2.25	2.25	9.00
State equity (co-invested)	1.00	0.75	0.50	0.50	0.25	3.00
Total public	4.00	5.25	6.50	7.75	6.50	30.00
Private equity target	1.50	2.50	4.00	5.50	6.50	20.00
Combined	5.50	7.75	10.50	13.25	13.00	50.00

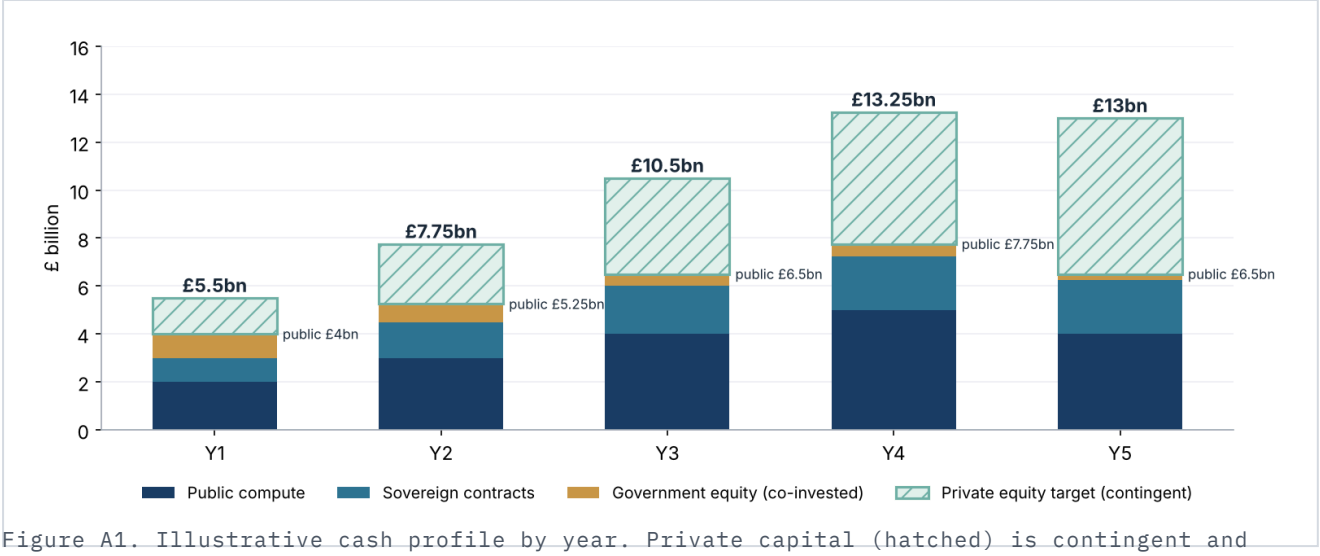


Figure A1. Illustrative cash profile by year. Private capital (hatched) is contingent and state equity is invested only alongside it. Not a forecast.

Equity is shown in the year of each round in Appendix A4, so the tables match. Private amounts are targets, not a schedule. The sovereign plan must be deliverable at the contracted public level even if private rounds slip. But a £27bn floor of compute and contracts should not be presented as carrying the same ambition as the full scenario.

A3. Funding scenarios

SCENARIO	PUBLIC COMPUTE	SOVEREIGN CONTRACTS	STATE EQUITY	PRIVATE RAISED	FIVE-YEAR TOTAL
Sovereign floor	£12bn	£6bn	£0.5–1.5bn	£1.5–6bn	£20–25.5bn
Central	£18bn	£9bn	£3bn	£20bn	£50bn
Accelerated	£28bn	£14bn	£5bn	£28bn	£75bn

State equity is always limited to co-investment alongside private capital actually raised, so it is lower in the floor scenario. The floor may keep useful sovereign capability but should not be described as near-frontier. The central and accelerated scenarios are hypotheses for cost validation, not proof that money alone buys capability. The extra £2bn equity authority applies only in the accelerated scenario.

A4. Illustrative cap table

The table shows mechanics, not an agreed valuation. The state subscribes at the same price and in the same instrument as private investors in every round, and private investors take the larger share each time. At founding, the team and option pool hold the pre-money equity in a 62.5:37.5 split.

ROUND	PRE-MONEY	PRIVATE / STATE CASH	POST-MONEY	STATE	FOUNDERS	EMPLOYEES	PRIVATE
Founding (Y1)	£2.5bn	£1.5bn / £1.0bn	£5.0bn	20.0%	31.2%	18.8%	30.0%
Series A (Y2)	£8.0bn	£2.5bn / £0.75bn	£11.25bn	20.9%	22.2%	13.3%	43.6%
Series B (Y3)	£15.0bn	£4.0bn / £0.5bn	£19.5bn	18.6%	17.1%	10.3%	54.0%
Series C (Y4)	£28.0bn	£5.5bn / £0.5bn	£34.0bn	16.8%	14.1%	8.4%	60.7%
Series D (Y5)	£53.25bn	£6.5bn / £0.25bn	£60.0bn	15.3%	12.5%	7.5%	64.7%

Private rounds total £20bn and state subscriptions £3bn. The valuations are illustrative and do not predict £60bn (section 5.2 explains what that would need). Option refreshes, preferences, debt and secondary sales would change the figures. The state's sovereign rights do not depend on its percentage.

A5. Accounting discipline

The model should keep separate ledgers for the compute operator and the laboratory, eliminating transactions between them in consolidated reporting so no pound is counted twice. It must also separate cash committed from cash spent, equity from grants and procurement, primary from secondary sales, cash raised from paper valuations and running costs from capital assets. Public accounts should disclose future contract commitments, the present value of obligations and the cost of maintaining sovereign models if the company stops trading.

Technical Appendix B: Compute and energy costing

B1. Bottom-up procurement

Costing should start from the chief scientist's approved experiments: accelerator hours, cluster size, useful throughput, utilisation, interconnect, storage, dependencies and redundancy. These become phased purchases, leases and replacements, with a reserve for failed experiments and changed architectures. Peak floating-point figures are a weak measure. Procurement should test real workloads at full-cluster scale.

B2. Illustrative power costs

At constant load, annual energy in terawatt-hours equals IT load (MW) × power usage effectiveness × 8,760 ÷ 1,000,000. The table assumes an effectiveness of 1.25. Prices should be checked against the Department for Energy Security and Net Zero's industrial price statistics, including network and policy costs and any AI Growth Zone terms.

IT LOAD	ANNUAL ENERGY	AT £100/MWH	AT £150/MWH	AT £200/MWH
50 MW	0.55 TWh	£55m	£82m	£110m
150 MW	1.64 TWh	£164m	£246m	£329m
300 MW	3.29 TWh	£329m	£493m	£657m

Power is material but small next to chips, networking, storage, maintenance, refresh and facilities. Real tariffs, grid charges and utilisation are needed before a funding decision.

B3. Illustrative hardware capital

Hardware cost is installed accelerators multiplied by the all-in cost per accelerator, with facilities, operations and renewal added separately. For example, 60,000 accelerators at £80,000 each cost £4.8bn for one generation. 120,000 at £100,000 cost £12bn and draw about 300MW. These are simplified assumptions, not quotes, and device counts do not compare across generations. Supplier bids should be independently benchmarked, with acceptance tests, guaranteed power, phased site readiness and replacement rights. Lock-in and single-source risks must be shown to the investment committee.

B4. Operational sovereignty test

A quarterly drill should confirm designated models can be loaded, run, secured, restored and updated entirely on British-controlled systems. An annual exercise should simulate losing a major cloud or software vendor and review the refresh-denial plan (section 8.4). Research capability needs its own test of training and reproduction, since passing a deployment drill does not prove the UK can build new models.

Technical Appendix C: Sovereign rights heads of terms

C1. Special share reserved matters

The state's prior consent is required for the following and nothing else:

- sale or exclusive licensing of designated sovereign intellectual property and models
- change of control of the company or a critical subsidiary, including the defence subsidiary
- moving designated research, compute or operations outside the UK
- any single foreign holding above a cap set during legal design
- granting rights that would block independent UK use of designated models
- ending or disabling the continuity arrangements.

Consent is not needed for hiring, pay, architecture, pricing, publication, budgets, board appointments or routine contracts. The share carries no right to direct general policy. The list of designated assets must be precise, kept up to date and narrow enough to be commercially workable. Decisions should follow published principles and fixed response times.

C2. Model continuity licence

The state receives a perpetual, irrevocable licence, fully paid or royalty-defined, to run specified sovereign models for public purposes on UK infrastructure, with rights to patch them. Third-party assets remain on their original terms. Escrow covers agreed weights, code, infrastructure recipes, model cards, documentation and deployment artefacts, plus training configurations where feasible, all checked regularly. The compute operator is the host of last resort.

Triggers include insolvency, material breach of continuity obligations, failure to meet essential public-service needs and a prohibited change of control. They must be clear enough for lenders, investors and courts. National-security access must be separable from ordinary licensing disputes.

C3. Board and investor rights

The state appoints one non-executive director and has no other appointment rights. Founders and employees have real voting participation and vesting shares. The shareholders' agreement provides secondary windows, security-based transfer conditions and a listing path. Investors get standard information and governance rights but no guaranteed returns and no veto over the sovereign settlement.

An **investor eligibility policy** covers primary rounds as well as secondary sales. It defines eligible jurisdictions and investor types, the treatment of sovereign and state-linked funds and how holdings count towards the foreign cap, with National Security and Investment Act notification where required. The policy affects both the pool of capital and its price, so it must be settled before the founding round.

C4. Company failure

The company could fail despite useful research. Insolvency planning must stop critical artefacts getting trapped offshore or in disputed licences. A continuity trust, pre-agreed receiver access, escrow and a team-transfer plan all need specialist legal testing. Any automatic insolvency right must be checked against insolvency law, security interests, licences and other creditors' rights.

C5. Classification, competition and data access

The ONS should assess the shareholding, board rights, special share and contracts together, with a provisional view before the founding round (section 6.5). Subsidy and competition advice must cover compute pricing, preferential access, contracts and licensing. Clinical data stays with approved custodians in secure environments. Classified defence material sits under separate controls in the defence subsidiary.

Technical Appendix D: Performance gates

WHEN	TECHNICAL	SOVEREIGNTY	COMMERCIAL AND INDUSTRIAL	DECISION
Before founding round	Credible team and 24–36 month research plan	Special share and contracts agreed. Provisional ONS view	Private lead on agreed terms. Independent value assessment	Form, or negotiate with an existing laboratory
Month 12	Training stack running. Reproducible baseline models	Secure compute and artefact custody	First paid or validated pilots	Release next tranche or fix gaps
Month 24	Original training campaign. Independent tests	Reproducible UK deployment under domestic control	External results from specialist divisions	Expand, reconfigure or pause hardware
Month 36	Several model generations. Competitive cost and quality	Continuity exercise, retraining capability, refresh-denial plan	Repeat licensing customers. Early productivity evidence	Year-three review: scale, maintain, partner or restructure
Month 60	General capability within a few months of the closed frontier on an independent index. Leadership in at least one scientific domain	UK-based operation, refresh and research continuity	Verified additional industrial output and revenue	Approve next five-year settlement

No single public benchmark can be the target. General capability should be tracked on an independent index. Scientific models should be tested on held-out tasks, physical experiments and customer outcomes. Independent experts should judge whether the laboratory is building frontier-relevant models or just fine-tuning public ones.

Escalation triggers include compute cost inflation, grid delays, export restrictions on accelerators, loss of key talent, failed private rounds, supplier concentration, adverse classification and repeated missed milestones. If private investment falls short, government should publish the gap between the sovereign floor and the full ambition rather than quietly filling it with public money.

Technical Appendix E: Economic assessment and options

E1. Fiscal sensitivity

The table shows what different year-ten GVA outcomes would mean. It is not a forecast. Attributable receipts equal GVA uplift × the share additional to the programme × an assumed fiscal conversion. The result is *annual* receipts in year ten, not a ten-year total.

CASE	YEAR-TEN GVA UPLIFT	ATTRIBUTABLE TO PROGRAMME	FISCAL CONVERSION	ANNUAL RECEIPTS
Low	£5bn	25%	20%	£0.25bn
Central	£15bn	50%	25%	£1.88bn
High	£30bn	60%	30%	£5.40bn

Against an ongoing public cost of £5–8bn a year after year five (section 7.5), even the high case does not pay for itself through tax. A full appraisal would model the whole benefit path, displacement, where profits arise, research costs, renewal, alternative uses of funds, discounting and foreign model availability. Equity value must be assessed separately and not added to output benefits.

E2. Options to appraise

Option A: buy international models. Simplest and often cheapest in the short term, but no independent training capability, limited strategic access and few research spillovers. The do-minimum comparator.

Option B: national mid-sized model. Less funding and infrastructure, with more control for selected public functions. May be enough for many security tasks but does not reach the frontier.

Option C: British frontier laboratory plus sovereign compute (recommended). As set out in this paper: privately led, with the state as minority co-investor and sovereignty secured by special share, contracts and statute. Higher cost and risk, larger spillovers and a direct route to independent model generations.

Option D: an existing British laboratory. Possibly faster access to world-class founders and working systems, if enforceable sovereign rights can be agreed with existing investors at a reasonable price. Also the main fallback if Option C cannot secure a private lead.

Option E: anchor customer without state equity. The state buys capability and provides compute but takes no stake. Mistral, a privately led European laboratory now valued at $24bn, shows such a company can attract large private capital [16]. Lower public cost but weaker alignment and no share of upside. The special share would still need at least a nominal holding.

Option F: allied pooling. Joint funding of compute or specialist divisions with like-minded partners. Shares cost and risk but dilutes control. Could complement Option C in specific domains.

The foundation phase should compare Options C, D and E against the same sovereign obligations, timelines and independent cost estimates.

E3. Ten-year financial model

The model should cover public and private investment, headcount, compute, depreciation, replacement, energy, cash burn, tax, equity returns and downside losses. It must show **sovereign and commercial revenue separately, with the margin on each**, and test whether commercial revenue can support the valuations in Appendix A4. It should use at least three hardware and research architectures, drawing on trend data such as Epoch AI's [4]. Scenarios should include a faster-moving frontier, slower catch-up and restricted accelerator exports.

It must also estimate the minimum annual public cost of keeping sovereign capability if the company fails or stops training. That figure is the programme's critical fiscal liability and cannot be replaced by assumptions about future valuations.

Sources

These sources support the factual baseline, not the illustrative financing or economic scenarios. Some company figures are reported second-hand. Global AI spending figures are not directly comparable with this programme without adjusting for accounting, exchange rates and consumer operations.

[1] Council for Science and Technology (2026), *AI sovereignty: preserving UK agency in an AI-enabled world*, letter dated 30 July 2026, published 19 October 2026. https://www.gov.uk/government/publications/ai-sovereignty-preserving-uk-agency-in-an-ai-enabled-world/ai-sovereignty-preserving-uk-agency-in-an-ai-enabled-world-web-version

[2] Centre for Emerging Technology and Security, Alan Turing Institute (2026), *Securing the UK's AI Future: A Sovereignty Vision for 2030*, 2 September 2026. https://cetas.turing.ac.uk/publications/securing-uks-ai-future-sovereignty-vision-2030

[3] Stanford Institute for Human-Centered AI (2026), *The 2026 AI Index Report: Technical Performance*. https://hai.stanford.edu/ai-index/2026-ai-index-report/technical-performance

[4] Epoch AI, *Trends in AI* (accessed 25 October 2026). https://epoch.ai/trends

[5] Ars Technica (2026), report of OpenAI financial statements 2024–25, June 2026. https://arstechnica.com/ai/2026/06/leaked-financial-docs-show-openai-is-losing-billions-of-dollars-a-year/

[6] Epoch AI, *Most of OpenAI's 2024 compute went to experiments* (data insight). https://epoch.ai/data-insights/openai-compute-spend

[7] Epoch AI (2026), *Open models lag state-of-the-art closed models by 4 months*, 29 May 2026. https://epoch.ai/data-insights/open-closed-eci-gap

[8] MIRI Technical Governance Team, *Catch-up algorithmic progress might actually be 60× per year*. https://techgov.intelligence.org/blog/catch-up-algorithmic-progress-might-actually-be-60x-per-year

[9] Department for Science, Innovation and Technology (2026), *Main Estimate memorandum 2026 to 2027*. https://www.gov.uk/government/publications/dsit-main-estimate-memorandum-2026-to-2027/dsit-main-estimate-memorandum-2026-to-2027

[10] Department for Science, Innovation and Technology (2025), *UK Compute Roadmap*. https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

[11] Department for Science, Innovation and Technology (2026), *AI Research Resource: advanced supercomputers for the UK*. https://www.gov.uk/government/publications/ai-research-resource/airr-advanced-supercomputers-for-the-uk

[12] University College London (2026), *UCL researchers lead two of Europe's largest-ever AI funding rounds*, 28 April 2026. https://www.ucl.ac.uk/news/2026/apr/ucl-researchers-lead-two-europes-largest-ever-ai-funding-rounds

[13] British Business Bank (2026), *British Business Bank and Sovereign AI invest into AI superintelligence company Ineffable Intelligence*. https://www.british-business-bank.co.uk/news-and-events/news/british-business-bank-and-sovereign-ai-invest-ai-superintelligence-company-ineffable-intelligence

[14] Department for Science, Innovation and Technology (2026), *First backing through UK's Sovereign AI*, 16 April 2026. https://www.gov.uk/government/news/ai-firms-pioneering-drug-discovery-cheaper-supercomputing-and-more-get-first-backing-through-uks-sovereign-ai

[15] University of Bristol (2026), *Isambard-AI supercomputer powers £500m UK Sovereign AI fund*, 20 April 2026. https://www.bristol.ac.uk/research/centres/bristol-supercomputing/articles/2026/isambard-ai-supercomputer-powers-500m-uk-sovereign-ai-fund.html

[16] Reuters (2026), *French AI company Mistral hits $24 billion valuation in funding round*, 8 September 2026. https://www.reuters.com/world/europe/french-ai-company-mistral-hits-24-billion-valuation-funding-round-2026-09-08/

[17] UK Biobank (2025), *Whole genome sequencing data*. https://community.ukbiobank.ac.uk/hc/en-gb/articles/15468060583709-Whole-genome-sequencing-data

[18] NHS England, *Secure data environments*. https://digital.nhs.uk/data-and-information/research-powered-by-data/support-and-resources/background/secure-data-environments

[19] Digital Health (2025), *AI project to predict health outcomes paused over GP data concern*, 4 June 2025. https://www.digitalhealth.net/2025/06/ai-project-to-predict-health-outcomes-paused-over-gp-data-concern

[20] Ministry of Defence (2026), *Putting artificial intelligence at the heart of UK Defence*, 10 June 2026. https://www.gov.uk/government/publications/putting-artificial-intelligence-ai-at-the-heart-of-uk-defence/putting-artificial-intelligence-ai-at-the-heart-of-uk-defence

[21] UK Space Agency (2026), *Annual report and accounts 2025–2026*. https://www.gov.uk/government/publications/uk-space-agency-annual-report-and-accounts-2025-2026/uk-space-agency-annual-report-2025-2026

[22] Department for Science, Innovation and Technology (2026), *Semiconductor sector study 2026*. https://www.gov.uk/government/publications/semiconductor-sector-study-2026/semiconductor-sector-study-2026

[23] UK Research and Innovation (2026), *National Materials Innovation Programme*. https://www.ukri.org/what-we-do/browse-our-areas-of-investment-and-support/national-materials-innovation-programme/

[24] Office for National Statistics, *UK economic statistics sector and transaction classifications: the classification process*. https://www.ons.gov.uk/methodology/classificationsandstandards/economicstatisticsclassifications/uk economicstatisticssectorandtransactionclassificationstheclassificationprocess

[25] Office for National Statistics (2025), *Defining the boundary between the general government sector and public non-financial corporations in economic statistics*, 28 November 2025. https://www.ons.gov.uk/economy/governmentpublicsectorandtaxes/publicsectorfinance/articles/definingtheboundarybetweenthegeneralgovernmentsectorandpublicnonfinancialcorporationsineconomicstatistics/2025-11-28

[26] Eurostat (2018), *Advice on the statistical treatment of capital injections in AREVA*. https://ec.europa.eu/eurostat/documents/1015035/8683865/Advice-2018-FR-Nat-accounts-treatment-of-capital-injections-in-AREVA.pdf/8e668574-6bfc-467a-81aa-680a536be121

[27] HM Government, *National Security and Investment Act: details of the 17 types of notifiable acquisitions*. https://www.gov.uk/government/publications/national-security-and-investment-act-guidance-on-notifiable-acquisitions/national-security-and-investment-act-guidance-on-notifiable-acquisitions

[28] Department for Business and Trade, *Subsidy control principles assessment guide*. https://www.gov.uk/government/publications/subsidy-control-principles-assessment-guides/subsidy-control-principles-assessment-guide-for-subsidies-or-schemes-of-interest-ssois-and-subsidies-or-schemes-of-particular-interest-ssopis

[29] Department for Science, Innovation and Technology, *Delivering AI Growth Zones*. https://www.gov.uk/government/publications/delivering-ai-growth-zones/delivering-ai-growth-zones

[30] Department for Science, Innovation and Technology (2026), *UK AI Hardware Plan*, 8 June 2026. https://www.gov.uk/government/publications/uk-ai-hardware-plan/uk-ai-hardware-plan

[31] UK Research and Innovation (2026), *Strategy 2026 to 2031*. https://www.ukri.org/publications/uk-research-and-innovation-strategy-2026-to-2031/ukri-strategy-2026-to-2031/

[32] Single Source Regulations Office, *Baseline profit rate recommendations*. https://ssro.gov.uk/category/baseline-profit-rate-recommendations