BCS' response to the AI Growth Lab consultation

A Government open evidence call from the Department for Science, Innovation and Technology (DSIT)

January 2026

Compiled by Claire Penketh, BCS Senior Policy and Public Affairs Manager

Table of Contents

BCS
The Chartered Institute for IT
3 Newbridge House,
Newbridge Square,
Swindon SN1 1BY

Introduction

The Department for Science, Innovation and Technology (DSIT) sought evidence on a pioneering cross-economy sandbox that would oversee the deployment of AI-enabled products and services that current regulation hinders.

The AI Growth Lab would support growth and responsible AI innovation by making targeted regulatory modifications under robust safeguards and with careful monitoring.

The call for evidence aimed to collect essential information about people's feedback on the proposed AI Growth Lab to help inform government policy development.

BCS answered through a UK growth lens. If the policy intent differs (e.g., primarily governance research vs. economic acceleration), priorities and metrics should be explicit. We favoured transparency (public registers and reports) while avoiding disclosures that could enable test gaming or pose security risks.

Four BCS Fellows contributed their views to this consultation: Emma McGuigan, Independent Advisor, Board Director & Technologist and BCS' AI Lead; Nigel Barnes, Co-Founder, Technology Advisor; Professor James Davenport, Professor of Information technology at University of Bath, and Adrian Hopgood, Emeritus Professor of Intelligent Systems at the University of Portsmouth and member of the BCS AI Specialist Group.

Q1. The AI growth lab would offer a supervised and time-limited space to modify or disapply certain regulatory requirements. To what extent would an AI Growth Lab make it easier to develop or adopt AI?

- Much easier
- Somewhat easier ☑
- I don't think there would be an effect
- Don't know

Q2. What advantages do you see in establishing a cross-economy AI Growth Lab, particularly in comparison with single regulator sandboxes?

Expert members from BCS, the Chartered Institute for IT - the professional membership body for technology, believe a cross-economy AI Growth Lab could provide greater clarity on what is allowed and an opportunity for cross-disciplinary thinking, which could improve consistent, joined-up guidance across regulations that routinely cut across sectors (e.g., data protection, financial services). This could reduce the risk of vested interests taking over sectors, cut down conflicting interpretations and promote more 'blue skies' research environments.

Many real-world AI use cases span multiple regulators (for example, the Financial Conduct Authority, Ofcom for communications, or ICO for data). A common entry point, shared triage, and coordinated supervisory input would therefore accelerate testing and reduce compliance friction. This could be particularly helpful for SMEs, who may not fully grasp the cross-sector and cross-jurisdiction implications of their AI product/service.

From a UKPLC perspective, there's an advantage in the scale-up community because they're working within the constraints without having to understand them all themselves.

It could also build coherence in testing methods (risk assessment, bias testing, security controls) and help smaller firms by embedding compliance-by-design patterns they can reuse. Public transparency of pilots can raise trust and avoid 'one-off' sandboxes proliferating with divergent rules. If the question also covers firms that operate across economies, then a unified Lab aligned with, for instance, the EU AI Act requirements, will ensure that UK products remain competitive internationally. In addition, a cross-economy Lab could avoid the issue of sandboxes that differ significantly, thereby reducing the burden on firms of multiple sandboxes that potentially create inefficiencies and discourage participation.

Finally, a cross-economy model is better able to engage hyperscalers and shared infrastructure providers (computational model access, evaluation tooling) than multiple

Q3. What disadvantages do you see in establishing a cross-economy AI Growth Lab, particularly in comparison with single regulator sandboxes

Scope creep and complexity: a genuinely cross-economy remit risks becoming so broad that it is never 'done', as the sandbox can get too big, with uneven sector coverage and frequent re-prioritisations as AI regulation evolves. This can dilute agility—the very quality sandboxes are intended to provide.

It could also be slow and bureaucratic, and has the potential to be elitist. It also needs to align with market research on where the greatest growth opportunity lies for UK businesses to be truly effective.

There could also be a real risk due to governance delays and inconsistent policy signals from different sectors—for instance, competing rulebooks that differ across jurisdictions or evolve (e.g., risk-tiering approaches). Without a tight mandate, this creates rework and uncertainty, undermining trust and adoption.

Operational load: coordinating multiple regulators, onboarding diverse participants, and provisioning shared computational/ model access is complicated. If not well-scoped, capacity constraints could crowd out SMEs or lead to long queues.

Governance risk: More actors increase the chance of inconsistent advice. Strong triage, issue ownership, and transparent decision logs are essential. In short, breadth brings

value, but only with crisp, well-thought-out scope boundaries, clear success criteria, and time-boxing.

Q4. What, if any, specific regulatory barriers (particularly provisions of law) are there that should be addressed through the AI Growth Lab? If there are, why are these barriers to innovation? Please provide evidence where possible.

Our experts saw few regulatory barriers per se, but identified areas of uncertainty that could deter innovation or market access. The priority is clarity and consistent interpretation, rather than broadly 'turning off' protections.

For instance, the EU AI Act Articles 10 and 15 would likely have to be addressed within the AI Growth Lab to ensure market deployment for goods and services across the EU.

The overall consensus among our experts was that Labs should enable testing for bias, data poisoning, cybersecurity compliance, and model exfiltration risks, among others.

Chief among these is the application of data protection law to modern AI development and use (especially the processing of personal data in model training, fine-tuning, evaluation, and red-teaming). Organisations are often unsure about lawful bases, minimisation, and transparency duties when models are probabilistic and multi-purpose.

For instance, in healthcare, health datasets are tightly controlled to the extent that innovation could be stifled. There are good reasons, of course, as even anonymised data can be traced back to individuals. A twofold solution that could work: (1) a technical check on the quality of datasets (e.g. they should be representative of the target population and depersonalised) and (2) a set of clear, consistent, and rapid protocols for accessing them.

In several regulated domains, explainability/justification requirements limit the use of opaque models—for example, in financial advice, where explanations must meet professional standards—creating hesitancy when model internals are not interpretable.

The Lab could address these by: (1) providing authoritative, regulator-backed patterns (e.g., privacy-preserving training/evaluation, acceptable use of synthetic or de-identified data, model cards and decision rationales); (2) offering safe-harbour style supervision for time-limited tests; and (3) issuing coordinated guidance notes where multiple regulators share competence.

Q5. Which sectors or AI applications should the AI Growth Lab prioritise?

Prioritise where UK growth and productivity impact are highest and where cross-regulatory issues are most acute. Selection criteria should include: (i) clear UK export or scaling potential; (ii) measurable productivity uplift; (iii) significant cross-regulator dependence (e.g., data protection + sector rules); (iv) public value and trust benefits; and (v) ability to reuse methods/tooling across sectors.

Examples include: enterprise AI for regulated services (financial services, healthcare operations, but not safety-critical treatment), AI for secure communications and media integrity (Ofcom/ICO), and public-sector service optimisation (where central government is both rule-setter and user). We do not recommend importing external risk taxonomies wholesale; prioritisation should reflect UK PLC's opportunities and capabilities, with transparent criteria and open calls.

Prioritisation should also follow the EU definitions of high-risk systems: HR tools, educational assessments, and applications that affect fundamental rights.

Q6. What could be potential impacts of participating in the AI Growth Lab on your company/organisation? n/a

Q7. Several regulatory and advisory sandboxes have operated in the UK and around the world, for example, the FCA's Innovate Sandbox, the Bank of England / FCA Digital Securities Sandbox, the MHRA's AI Airlock, and the ICO's Data Protection Sandbox. Have you participated in such an initiative?

- Yes
- NO ☑

Q8. What lessons from past sandboxes should inform the design of the AI Growth Lab?

Design for clarity and agility. The two recurring lessons for sandbox effectiveness are: (1) be explicit about scope—what kinds of systems, data, and decisions the Lab covers—and (2) be explicit about what it is solving for (e.g., resolving legal uncertainty; creating shared testing methods; enabling time-limited derogations, etc.). Access to regulators is the primary value—authoritative, joined-up guidance and rapid issue resolution, over blanket rule suspensions. Publish standard artefacts (risk assessments, evaluation checklists, assurance templates) that participants can reuse after the pilot.

Lean governance helps with a single front door, named regulatory owners for issues, transparent decision logs, and time-boxed pilots with success metrics. Finally, plan for infrastructure: many pilots will require hyperscaler-hosted computational models. Early partnerships avoid bottlenecks and make it feasible for SMEs to participate.

Q9. What types of regulation (particularly legislative provisions), if any, should be eligible for temporary modification or disapplication within the Lab? Could you give specific examples and why these should be eligible?

Our experts cautioned against the broad disapplication of core protections. The greater need is for supervised interpretation and narrowly tailored flexibilities. Where temporary modifications are considered, they should be procedural and reversible, for example: (i) time-limited relaxation of certain notification/reporting timelines; (ii) supervised use of de-identified/synthetic data for experimentation where equivalent real-world data would otherwise be required; (iii) allowance to trial explainability approaches (post-hoc and system-level) where perfect feature-level interpretability is not yet feasible.

Any such flexibility should be granted case-by-case, with explicit guardrails (purpose limitation, data minimisation, independent evaluation, and sunset dates). The emphasis should be on gaining clarity from regulators rather than 'switching off' data protection.

Although the UK is not an EU member state, it would be wise to consider how a UK AI Growth Lab fits within Article 57 of the EU/AI Act, which references explicitly criteria for AI regulatory sandboxes.

Q10. We propose that certain types of rules and obligations, such as those relating to human rights, consumer rights and redress mechanisms, and workers' protection and intellectual property rights, could never be modified or dis-applied during a pilot. What types of regulation (particularly legislative provisions) should not be eligible for temporary modification or disapplication within the Lab (e.g. to maintain public trust)?

Rules protecting human safety and life, fundamental rights, consumer protections and redress, and workers' protections should not be disapplied. Safety-critical thresholds (e.g., clinical safety constraints, radiation dose limits, aviation safety standards) are essential to public trust and must remain in force during any pilot.

Our experts would avoid designating intellectual property rights as 'never modifiable' in blanket terms. IP in the AI context (training data, model outputs, derivative works) is unsettled and, if rigidly fixed, could create a legislative quagmire within a sandbox. Where IP issues arise, handle them via clear participation agreements and targeted guidance rather than statutory disapplication.

As far as sensitive data, such as health care data, is concerned, personal data privacy rules need to be respected to retain trust - albeit, as per our response in Q9, there could be scope for some innovative thinking to avoid stifling innovation. It's about getting clarity from the regulator on what's possible and what's not, rather than per se turning off data protection. In short: preserve safety, rights, and redress; be pragmatic and contract-led on IP.

Q11. What oversight do you think is needed for the Lab?

- Parliamentary scrutiny when modifying or disapplying regulations within the Lab

- A Statutory Oversight Committee made up of sectoral regulators and independent experts ✓
- Public transparency and reporting ✓
- None
- Don't know
- Other (please specify)

Q12. How would this oversight work most effectively?

Create a lean, expert Oversight Committee with a single front door, public transparency, and clear escalation paths. BCS's recent YouGov survey highlighted that public trust is essential to the adoption of technology, including AI. Over 2,200 UK adults responded to the survey in August 2025 (https://www.bcs.org/media/mrlduktp/ai-ethics-uk-report.pdf), and results revealed that public trust in technology professionals is conditional on accountability and regulation. The survey found strong public support for formal oversight of IT professionals, with a large majority agreeing that those involved with public-facing systems should be on a public register and follow a Code of Conduct. Our survey showed responsibility for upholding high ethical and technical standards in IT is seen as a shared duty: half of respondents (50%) think an independent professional body such as BCS should be responsible, while significant proportions also point to the UK Government at (46%), employers, at 42% with with 41% backing a dedicated digital/ AI regulator.

Our experts believe a combination of formal regulatory oversight and paid independent staff would work. It would need both: a carefully constructed, publicly available, published framework and a register of pilots with scope, safeguards, and high-level findings to meet transparency goals—while avoiding disclosure of test results that would enable 'gaming' of tests or reveal security-sensitive details.

Operationally: establish standard entry criteria and triage; assign a lead regulator per pilot; keep auditable decision logs; and require outcome reporting against predefined metrics (safety, fairness, robustness, data protection, user impact). Escalate to Parliament only when proposing permanent legislative change beyond the sandbox's remit. This balances accountability with the agility needed for timely experimentation.

Q13. What criteria should determine which organisations or projects are eligible to participate in the Lab? Please select all that apply:

- You have an innovative product you want to bring to market
- Your innovation is intended for the UK market, or you are a UK based firm ✓
- Your innovation will benefit consumers ✓
- Your innovation is directly connected to AI ✓
- There is a regulatory barrier (legislation) which the AI Growth Lab would help overcome ✓

- There is a significant regulatory compliance resource otherwise needed to test the relevant product, which the AI Growth Lab would help avoid ☑
- Other (please specify)

Q14. Which institutional model for operating the Lab is preferable?If so, what are they and how might they be overcome?

- AI Growth Lab run by central government, with the support of sectoral regulators ☑
- AI Growth Lab run by a lead-regulator
- Don't know
- Other (please specify)

Q15. What is your reason for selecting this institutional model?

IThe central government can convene across regulators, align with public-sector use cases, and avoid duplication with existing sector sandboxes. Embedding the Lab within government (with strong regulator support) also allows costs to be integrated into programme budgets rather than treated as an external overhead. This improves sustainability and strengthens accountability for public value.

In practice, the government is best placed to broker partnerships with hyperscalers for computational /model access and evaluation tooling—capabilities that smaller firms cannot afford individually. A central model with regulator secondments preserves domain expertise while providing a single front door and consistent operating procedures. We do not recommend an independent institute model for the core Lab, as it risks opacity into costs and skills gaps relative to regulators' evolving AI capabilities.

In addition, there was some discussion amongst our experts around whether an independent, trusted, open neutral body, such as the National Physical Laboratory or UKRI, could be of value, should it be decided the central government model wasn't appropriate.

Q16. What supervision, monitoring and controls should there be on companies taking part in the Lab?

Requirements: (i) clear pilot charters (scope, hypotheses, safeguards); (ii) documented risk assessments and DPIAs where personal data is involved; (iii) security controls proportionate to threat (incl. red-teaming for model misuse); (iv) controls on model training/fine-tuning within shared environments to prevent manipulation that could degrade other participants' results; (v) periodic reporting on outcomes against pre-set metrics; and (vi) orderly exit plans with data/model artefact handling.

Regulators should supervise via named case officers; capacity management should prevent a few firms from monopolising resources. It was felt there should be a limit on requests from individual firms to prevent the monopolisation of resources. Without such controls, dominant players could overwhelm the lab, excluding smaller innovators. Sensible caps on concurrent projects and fair scheduling mechanisms would ensure equitable access. Monitoring should also verify compliance with agreed protocols, maintain integrity, and prevent misuse.

Where third-party foundation models are used, ensure supplier assurance (e.g., model provenance, safety policies, incident response) and contractual guardrails compatible with UK law and guidance.

Q17. Do you think a successful pilot in the AI Growth Lab would justify streamlined powers for making changes permanent, as opposed to following existing legislative processes which would take considerably longer?

- Yes
- No ☑
- Maybe
- Don't know

Q18. If you answered 'yes' or 'maybe' to question 22, what is the most effective way to achieve streamlined powers to make permanent legislative changes?

Q19. Would there be value in extending the AI Growth Lab to other high-potential technologies?

- Yes
- No
- Maybe ☑
- Don't know

Q18. If you answered 'yes' or 'maybe' to question 24, which technologies would benefit the most?

Quantum

Q19. Additional feedback?

Two overarching points:

- Define the sandbox: clarity on scope (systems, data, decision types), provisioning (model/computational access, evaluation tooling), and goals (resolve specific legal uncertainties vs. test guardrails) is essential. Without this, many questions become ambiguous.

- Infrastructure and partnerships: meaningful pilots will require hyperscaler-hosted compute and access to state-of-the-art models. Early partnership agreements are critical to affordability and SME inclusion.

Who we are

BCS is the UK's Chartered Institute for Information Technology. The purpose of BCS as defined by its Royal Charter is to promote and advance the education and practice of computing for the benefit of the public.

We bring together industry, academics, practitioners, and government to share knowledge, promote new thinking, inform the design of new curricula, shape public policy and inform the public.

As the professional membership and accreditation body for Information Technology we serve over 60,000 members including practitioners, businesses, academics, and students, in the UK and internationally.

We also accredit the computing degree courses in over ninety universities around the UK. As a leading information technology qualification body, we offer a range of widely recognised professional and end-user qualifications.