Royal Academy of Engineering

AI Growth Lab

Response to the Department for Science, Innovation and Technology's call for evidence

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

It could add value where deployments cut across multiple overlapping or fragmented regulators or governance actors. For example, those at the intersection of health, data protection and product safety, or autonomy, workplace safety and transport rules.

Compared to single-regulator sandboxes, the main advantage is coordination. That could look like a supervised front door for participants, joint trial designs agreed by relevant regulators and more consistent expectations on safety evidence, reporting of results and documentation.

The Academy's response was informed by experts in research and industry, drawn from our Fellowship and Awardee Excellence Community, including leaders in machine learning, robotics and autonomous systems, safety and assurance and computational approaches to science and engineering. Respondents highlighted potential advantages:

- When deployments engage several regulators at once, it can lead to duplicated documentation and long, bespoke approval journeys. A cross-economy Lab could convene relevant regulatory bodies around a joint trial protocol and provide a more coherent view than separate sandboxes would sequentially.
- Across different domains, regulators and operators often ask similar types of performance and safety questions about AI systems. A central Lab could turn them into tested, exportable blueprints, which regulators could adapt, rather than reinventing how to trial safely in each silo. That way, it could focus on better evidence, rather than being purely a deregulation exercise.
- Beyond regulators, local bodies, such as councils, NHS trusts, ports and airports, unions and insurers often determine whether trials proceed. A cross-economy Lab could standardise MOUs, risk-sharing and consent templates with visible support from relevant regulators.
- A cross-economy model is better placed to accredit and build common testbeds, secure data environments and evaluation tools that single-regulator sandboxes may not have the scale or mandate to justify.
- For data-intensive systems spanning domains, cross-economy coordination could offer value because copyright, data provenance and GDPR cut across sectors.

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

Many AI-enabled products sit between categories. This can create uncertainty about approvals, prompt innovators to avoid borderline use-cases, delay projects while they seek legal clarity. A cross-economy Lab risks giving rise to an additional bureaucratic layer on top of existing regulators. If poorly designed, it could slow decisions, add duplicative processes and create confusion about who has authority to approve, enforce or halt a trial, whereas single-regulator sandboxes usually have clearer lines of accountability.

A related risk is blurred accountability when things go wrong. Participants may be unsure whose standards apply, and the public may be unclear which regulator is responsible for enforcement and redress.

There's a risk of incumbents dominating agenda-setting. Lab design should support a diversity of participants, including startups, SMEs, academic innovators, going beyond established market leaders.

Risk appetite may vary across sectoral regulators. The Lab that's trying to reconcile different expectations could end up defaulting either to the most cautious position, offering limited

flexibility, or to an unduly relaxed position that might undermine trust in the wider regulatory regime.

There's risk around generalising from pilots. A successful sandbox test can be misread as proof of safety, security or fitness for deployment. Safety and performance are context-dependent. If the Lab is treated as a substitute for the application-specific assurance work and monitoring that should happen outside the sandbox, it could reach premature conclusions or create regulatory cliff-edges.

Participants will need to share enough evidence for regulators to make well-grounded decisions. Disclosure across a cross-economy Lab, with more parties and broader dissemination, increases the risk of exposing IP, sensitive data or security vulnerabilities. Getting that balance wrong could deter startups, SMEs, deep-tech companies.

It could drift into high-level upstream analysis without translating to practical deployment guidance. Legitimacy will depend on being technical, measurable, closely connected to application contexts.

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.

The experts who informed the Academy's response highlighted the following barriers:

- The barriers causing friction often result from uncertainty and duplication created by how multiple regimes interact. Where AI deployment straddles multiple regimes, for example autonomous vehicle type-approval and authorisation interact with local highway approvals; hospital robots spanning medical device, NHS governance and data protection; school robots combining product safety with children's data rules and safeguarding. This can drive case-by-case multi-regulator approval journeys with duplicated documentation.
- Legacy safety regimes designed for embodied AI systems are largely designed around fixed-function deterministic machines. Core frameworks, such as machinery safety regulations, the Provision and Use of Work Equipment Regulations, the Construction Design and Management regulations, Civil Aviation Authority rules for drones, road-traffic and automated vehicle frameworks, offshore safety-case regimes, and medical-device rules and hospital governance for clinical and assistive robots assume fixed-function systems. This pushes designers into constrained autonomy (i.e. low speeds, fenced areas, rigid modes) and creates a reluctance to update AI components once systems are signed off, delaying improvements and raising assurance costs.
- Some regulators are perceived as requiring deterministic evidence of behaviour in every scenario, which is mathematically unachievable for certain experimental and learning AI systems. This can force scope reduction or fallback to human-in-the-loop approaches that become too slow and expensive.
- Fear of IP litigation deters use of high-value datasets, such as journals, proprietary data and archives. SMEs are particularly affected due to indemnification costs.

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

Our advice is to prioritise sectors where there's clear societal and/or economic need, technology is mature enough to evaluate, deployment faces multi-regulator or evidentiary bottlenecks. Areas which government may wish to consider include:

- Embodied AI across mobility, construction, logistics, healthcare, agriculture, marine and offshore, drones and aviation and similar settings. Embodied AI deployments are more

complex to regulate, with more friction and fragmented regimes. Legacy safety frameworks assume deterministic machines, leading to constrained autonomy, fragmented approval pathways across regimes, and update paralysis once a system is signed off. The Lab could trial bounded operating domains, update and monitoring approaches, integrated pathways across regulators and local gatekeepers.
- Agentic systems, particularly where horizontal barriers constrain deployment.
- Computer vision systems, as they can support stronger assurance and more quantifiable evaluation than other AI systems.
- Healthcare, including clinical decision support and diagnostic imaging, and clinical and assistive robotics. There's high potential but slow translation into routine use due to overlapping governance, such as medical device expectations, NHS governance, data protection and confidentiality, and unclear evidence thresholds. The need for randomised control trials can be another factor.
- Transport autonomy and aviation drones, where deployments face multi-actor approvals and safety assurance challenges, including local approvals and liability and insurance constraints.
- AI for net zero and materials, where chemical safety processes under the Registration, Evaluation, Authorisation and Restriction of Chemicals regime make it difficult to automate the physical testing and scale-up of AI-predicted compounds. A sandbox could support trials in industrial labs with appropriate safety controls.
- Nuclear, where there is a need to enable more efficient adoption of innovation in a tightly controlled, safety-critical environment.
- Financial services, especially tabular models, which are used more freely in the US alongside stronger validation, while the UK is more constrained. Sandboxes could test and signpost engineering-grounded robustness and verification approaches.

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

The experts with relevant experience who informed the Academy's response suggest the Lab will work best if it's engineering-led, outcome-driven and designed to produce reusable evidence. The Academy would like to draw attention to the following lessons:

- Early low or near-zero risk opportunities build confidence and capability, but the aim should be to focus on higher risk, high reward applications.
- Sandboxes should resemble how we certify bridges or engines: clear technical questions, measurable outcomes and disciplined methods. This includes testing assurance techniques, such as robustness validation and verification approaches for machine learning and agreeing up-front what evidence will be considered credible. Pilots should generate reusable evaluation artefacts.
- A sandbox shouldn't be treated as sufficient evidence for deployment, but as the start of a clearly planned next phase. From the outset it should define what additional assurance, monitoring, update-management and incident-handling processes will be needed beyond the trial, who will own this work (e.g. which regulator), a roadmap from evaluation to supervised deployment and, where appropriate, broader adoption. If a pilot succeeds, there should be a concrete route, such as guidance and updated codes of practice, evidence requirements that relevant bodies will recognise, or a pathway into commissioning or procurement. If it doesn't succeed, lessons should be captured and shared in a non-sensitive form.
- The Lab should coordinate multi-regulator trials, but not duplicate regulators' responsibilities. This should be underpinned by clear, published selection criteria, e.g. the risk profile of use-cases, cross-economy relevance, feasibility within the Lab's

resources, the potential to generate reusable learning, alongside published outcome summaries and periodic external review.
- Data provenance and evaluation methods should be public to build trust, while balancing transparency with IP security considerations. A balance can be struck through sharing results and publishing non-sensitive lessons, but avoiding mandatory disclosure of proprietary designs or intrusive monitoring.

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?

We recommend time-limited flexibility, but only where rules are ambiguous, overly rigid in their current interpretation, or impose procedural burdens that are disproportionate for supervised trials. Modifying process friction should be prioritised over weakening core protections.

Candidates for temporary modification within the AI Growth Lab could include:

- Clarifying or relaxing constraints that prevent access to representative data for safety and performance testing, where lawful bases exist and privacy safeguards are in place. For example, the experts consulted for the Academy's response noted that simplistic interpretations of data protection rules can inhibit realistic test datasets, which can limit robustness evaluation.
- Temporary adjustments to documentation and reporting formats where current expectations are inconsistent, so that pilots can trial standardised evidence packs, such as risk cases, Operational Design Domains, incident logs, updated or changed control records, without forcing full certification-grade submissions at the outset.
- Time-limited permissions for bounded trials of embodied AI and autonomous systems in defined operational design domains, where existing frameworks can create grey areas about what is permitted, provided there are clear safety envelopes, stop and rollback mechanisms and agreed monitoring.
- Regulatory process flexibility to enable iterative model updates during a trial without re-triggering full approvals for every minor change, where changes remain within an agreed Operational Design Domain and are governed by change control and post-change monitoring.

Across these, flexibilities must be explicitly scoped, time-limited, tied to measurable evaluation plans and accompanied by predictable oversight, such as regular checks, stable criteria, ability to stop pilots. The purpose should be to test which evidentiary and assurance approaches work in practice and can later be adopted at scale. Relaxing regulation in the abstract should be avoided.

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)?

A sandbox should not create exemptions from core protections that underpin public trust. The following categories can be considered as red lines that should not be eligible for temporary modification or disapplication:

- Fundamental safety protections, particularly where the public, workers, or safety-critical infrastructure could be harmed. Instead, flexibility should sit around how evidence is generated and reviewed.
- Data protection essentials and basic privacy protections. The AI Growth Lab could help participants make fuller use of representative data for testing within the existing data protection framework, for example by clarifying lawful bases, standardising Data Protection Impact Assessments and enabling testing in well-governed environments, such as secure data environments. But there should be no relaxation of core requirements around lawfulness and fairness of processing, purpose limitation, security and the handling of sensitive data and children's data, or of protections for vulnerable groups.
- Any trial should maintain legal equality duties and should be designed to detect and mitigate discriminatory outcomes.
- Sandboxes should not be used to lower cybersecurity requirements or introduce insecure handling of models or data.
- People affected by pilots must retain routes for complaint and redress, and pilots should not undermine existing consumer protections.
- Where systems are trialled in workplaces, duties relating to health and safety and worker protections should remain in force.
- Participation should not require disclosure or waiver of IP protections.
- If users or consumers are involved, they should be only enrolled after providing informed consent, i.e. they should not be experimented on without their knowledge.

Which institutional model for operating the Lab is preferable? Please select one option:

- 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)

What is your reason for selecting this institutional model?

Based on the expert consultation the Academy has conducted, a central government-hosted model with strong regulator involvement was preferred over a single lead-regulator model, largely for capability and coordination reasons.

No single regulator currently has sufficient embedded expertise and mandate to run an AI Growth Lab across sectors. Regulatory contexts and constraints differ substantially. Appointing any one of them as lead risks creating a model optimised for that regulator's operating environment, but ill-fitted for others.

Many barriers are cross-cutting. A central AI Growth Lab can provide a unified entry point and orchestrate multi-regulator pilots that directly address these friction points, rather than forcing innovators to run parallel engagements. A central mechanism could also make it easier to coordinate with adjacent initiatives, such as the Regulatory Innovation Office and the NHS innovator passport mechanism, reducing duplication and fragmentation.

There may be a benefit of a central hub with regulator secondees or embedded teams from key regulators. This would retain domain expertise and statutory ownership with line regulators, while giving the AI Growth Lab the day-to-day capability to design joint trials, build shared templates, and codify reusable patterns that recur across sectors. Crucially, this model would be well placed to build and retain technical skills and capabilities required for an effective AI Growth Lab, including expertise in AI assurance and evaluation methods. Such a

model would also support a culture change across regulators – a barrier that can be as important as formal rules.

This model would also best support shared test and evaluation infrastructure that individual regulator sandboxes would struggle to justify and could better coordinate national regulators with local delivery bodies that often determine whether pilot deployments can scale.

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

The AI Growth Lab should combine time-limited regulatory flexibility in exchange for disciplined process, transparency and proportionate monitoring, to offer a fair deal.

Entry requirements for companies could include:

- Named accountable senior person for the trial
- Evidence of sufficient organisational maturity
- Clear description of the problem the innovation is solving, how it is expected to support economic growth and public value, and a supportive customer
- An initial safety and risk case against baseline practice, including hazard analysis and a clear operational design domain showing where the system will and will not be used
- A revert plan and human-in-the-loop arrangements where appropriate
- When applicable, a short description of the training data
- Ethics preparation where relevant, such as identification of directly affected groups and settings, the proposed consent approach and any required approvals
- An evaluation plan with baseline and target metrics

Monitoring during the trial could include:

- Immediate notification of serious incidents or near-misses against agreed thresholds.
- Regular short reports on agreed KPIs
- Audit and inspection access proportionate to risk
- Requirements should avoid changing reporting regimes mid-trial or requiring re-approval for every minor update

After the trial:

- A public summary of what was trialled and results against the backdrop of the baseline, with a technical annex shared with regulators where needed
- Clear next steps with action owners.
- What would be excessive:
- Continuous real-time data streaming to government beyond incident reporting, especially involving personal or commercially sensitive data
- Mandatory public disclosure of proprietary designs, source code or detailed datasets that compromise IP or cyber resilience.

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? Please select one option:

- Yes
- No
- Maybe
- Don’t know

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

There is a possibility of streamlining, but only with safeguards. Two complementary points were emphasised through the Academy's consultation:

- A successful pilot is not, by itself, proof of safety or fitness for universal deployment. Safety, security and performance are context-dependent, and there is a risk that sandbox participation is treated as a substitute for the downstream assurance work that must still be done.
- Evidence from sandboxes can justify faster permanent changes where results are independently validated, risks and mitigations are well understood, a diverse set of participants have been tested, and impact assessments show clear economic and social benefit. Streamlining should accelerate technical adjustments and guidance updates without weakening public oversight.

The AI Growth Lab can be used to produce an agreed evidence blueprint that line regulators and certification bodies can recognise.

Where legislative change is required, time-limited, targeted modifications should be used with appropriate Parliamentary scrutiny, and only after the conditions above are met.

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

Please select one option:

- Yes
- No
- Maybe
- Don't know

If you answered 'yes' or 'maybe', which technologies would benefit the most?

Several areas were identified for consideration through the Academy's consultation:

- Quantum communications and security, as the technology is advancing faster than many assurance and deployment pathways. A supervised environment could validate performance and security claims and clarify expectations.
- Engineering biology, where governance spans safety, ethics, biosecurity, environmental regulation and product regulation, with complex interactions that can slow responsible experimentation and scale.
- Decentralised technologies, where these raise questions spanning financial regulation, consumer protection and legal accountability that could benefit from time-limited supervised trials.
- Enabling technologies for AI deployment, such as synthetic data generation and privacy-enhancing technologies, as these could be trialled to demonstrate safe ways to test and validate systems while protecting privacy, potentially easing data-access bottlenecks.

Thank you for taking the time to complete the survey. We really appreciate your time. Is there any other feedback or evidence that you wish to share?

Please select one option:

- Yes
- No

If you answered 'yes', please set out your additional feedback or evidence.

The AI Growth Lab will most successfully drive growth if it tests promising innovations from diverse companies of all sizes, which requires targeted communication that clearly articulates the purpose, as this isn't immediately obvious from the title.