Defensive acceleration: the strategic pivot needed for UK biological resilience

Dr Richard Moulange, Oliver Payne,
Dr Paul-Enguerrand Fady
and Dr Cassidy Nelson

January 2026

THE CENTRE FOR
LONG-TERM RESILIENCE

Contents

Executive Summary **3**

Introduction **5**

UK SONAR: An AI-enabled pandemic early warning system **12**

Concept 12

Economic and strategic case 13

Implementation approach 14

UK PROVE: A test-bed to accelerate next-generation biosecurity technologies **15**

Concept 15

Economic and strategic case 16

Implementation approach 17

UK SCAN: Validating legitimate access to biological tools and materials automatically with frontier AI **18**

Concept 18

Economic and strategic case 19

Implementation approach 20

UK TRUST: Unlocking privileged frontier AI access for national security **21**

Concept 21

Economic and strategic case 22

Implementation approach 23

Conclusion **24**

References **25**

Acknowledgements

This report was prepared by The Centre for Long-Term Resilience's Biosecurity Policy Unit, with special acknowledgement to Frankie Di-Nozzi and Aira Gonzales for its design. We particularly wish to thank those experts who reviewed earlier drafts of the report and provided feedback.

Suggested citation: Moulange, R. et al., 2026. 'Defensive acceleration: the strategic pivot needed for UK biological resilience'. *The Centre for Long-Term Resilience*. Available at: https://doi.org/10.71172/x093-veqz

Contact: Dr Richard Moulange, AI Policy Manager, richard@longtermresilience.org

Executive Summary

The UK Strategic Defence Review (SDR) 2025 identifies chemical and biological defence R&D as *the* essential and urgent activity. The increasingly fractured geopolitical landscape and the rapid maturation of artificial intelligence (AI), engineering biology and their intersection represent significant shifts to the threat landscape. While the accelerating advancement of AI offers transformative technological scientific and economic opportunities and benefits to the UK economy, it also lowers barriers for state and non-state actors to develop biological and chemical weapons.

To date, mitigation efforts for AI biological risks have focused on advancing managed access and model safeguards. However, the proliferation of open-weight frontier models—which can be accessed by anyone and cannot be meaningfully safeguarded—renders these containment strategies insufficient.

Relying solely on restrictions is no longer a viable security strategy. Instead, the UK must actively shift the offence–defence balance towards a defensive advantage. The Centre for Long-Term Resilience (CLTR) proposes a strategic pivot to **defensive acceleration** for UK chemical and biological defence. At its heart, defensive acceleration is the prioritisation of technological development in defensive domains so they advance before offensive applications mature. It is distinct from directionless proliferation, focusing innovation efforts on safety-enhancing technologies.

We offer a list of 22 promising interventions and explore four priority options further in concept notes operationalising defensive acceleration:

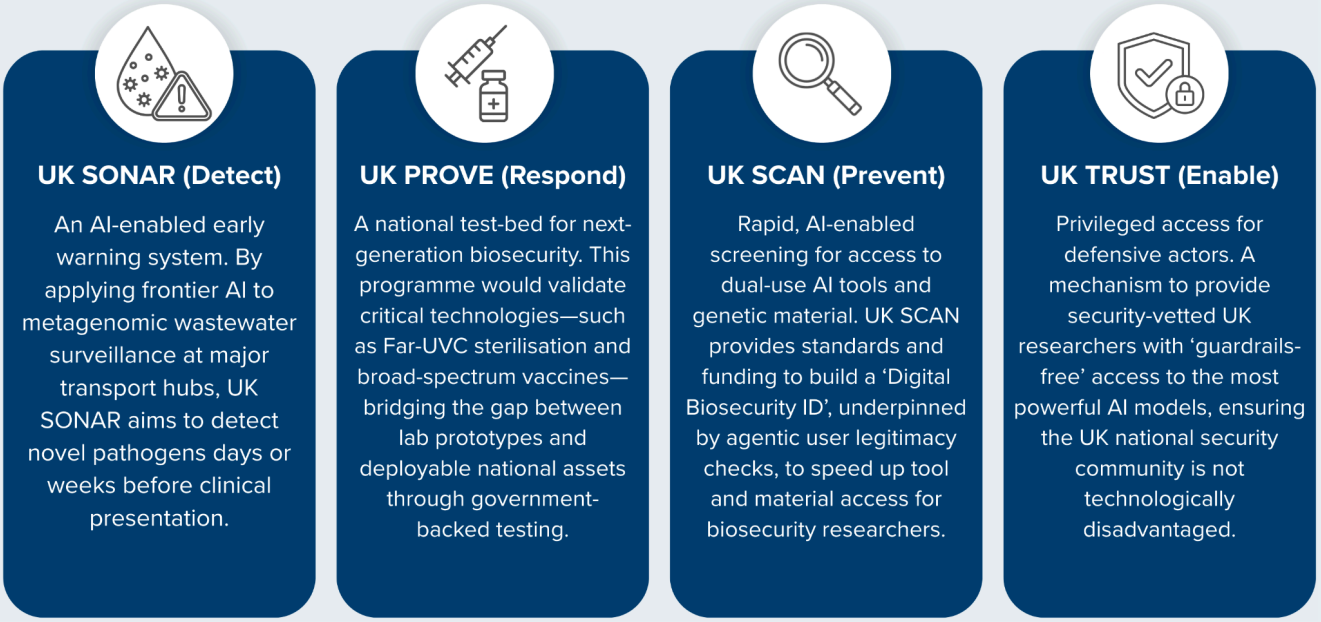

These interventions are designed to address the profound market failures that currently prevent the private sector from delivering necessary national security capabilities. We focus especially on biological defence, but these interventions can apply across the chemical and biological spectrum.

Implementing these four pilots would require an estimated investment of £40–65mn per year, offering significant value for money when measured against the safety and resilience benefits for chemical and biological defence. Beyond immediate security, these programmes leverage the UK's comparative advantages in life sciences and AI safety, creating exportable intellectual property and strengthening our contribution to the AUKUS and NATO alliances.

Introduction

The UK Government identifies AI and engineering biology as two of the six technologies most critical to the UK economy and its national security (Department for Science, Innovation and Technology, 2025). Both technologies have the potential to transform the strategic landscape, including by enabling a wide spectrum of threat actors—both state and non-state—to develop and deploy chemical and biological weapons (CBW). Defence against CBW has been designated the "essential and urgent activity" for the newly established Defence Research and Evaluation organisation in the UK's most recent Strategic Defence Review (2025).

There is growing recognition that the intersection of AI and engineering biology poses particular national security challenges. AI systems could both reduce barriers for different threat actors wishing to develop known CBW and help engineer novel pathogens with even greater catastrophic potential (Rose & Nelson, 2023). The concern that AI systems could "meaningfully assist non-state actors in advancing the development, production, acquisition or use of chemical or biological weapons, as well as their means of delivery" is recognised internationally by 27 nations, including the UK and the U.S. (Department for Science, Innovation and Technology, 2024).

Existing mitigations are insufficient in the short term

There have been two major types of AI-specific mitigation for AI-enabled CBW development to date: *managed access* and *model safeguards*. Managed access considers how to limit advanced AI model access to only those with a legitimate need to complete dual-use activities, such as developing vaccines for pandemic pathogens (Carter & Butchello, 2026). Safeguards are designed to directly prevent users from accessing dangerous information (Frontier Model Forum, 2025). These include both natural-language instructions for developing weapons and specific technical outputs (such as novel sequences of pandemic pathogens).

We consider it almost certain[1] that—over the next two years—managed access and safeguards alone will fail to prevent all but the most unsophisticated threat actors from using advanced AI assistance for CBW development. This is driven particularly by the *open-weight release* of many AI models with frontier or near-frontier capabilities: managed access is impossible for open-weight models (which can be freely shared on the internet) and almost all open-weight safeguards are easy to disable (AI Security Institute, 2025).

Defensive acceleration is necessary, alongside managed access and safeguards

At its heart, *defensive acceleration* is the strategic prioritisation of technological development in defensive domains so they advance before offensive applications mature. It is distinct from directionless technological proliferation, ensuring that safety-enhancing innovations are brought online faster than offensive capabilities.

[1] Here, and elsewhere within this report, we use specific probabilistic language corresponding to the Professional Head of Intelligence Assessment Probability Yardstick (Intelligence Analysis, 2025)

To counter AI-enabled chemical and biological threats, industry and civil society experts, as well as the UK Government, have repeatedly highlighted the need to accelerate biological and chemical defence (OpenAI, 2025; Parthemore & Weber, 2026; Ministry of Defence, 2025). This is especially important—even disregarding any potential AI-enabled threat actor uplift— given the more dangerous geopolitical landscape in 2026, with verified chemical weapons use in Ukraine by the Russian Federation (U.S. Department of State, 2024; Foreign Commonwealth and Development Office, 2024; Ministerie van Defensie, 2025) and growing fears of deliberate biological attack (Ministry of Defence, 2025).

Given that we expect managed access and safeguards to fail to contain future AI-enabled biological and chemical misuse, we assert that defensive acceleration *'def/acc'* is, in fact, a necessary option for UK biological and chemical security.

Indeed, across multiple critical national security activities, def/acc is an attractive option for three reasons:

1. **Economic benefits:** Def/acc focuses on advancing science and technology in key strategic domains. Investments in R&D and new defensive technologies help grow the economy and create new high-skilled jobs while addressing market failures in building societal resilience to security threats.

2. **Politically viable:** Def/acc will unlock more 'win–win' policy options, unlike restricting AI access and information. Creating new companies and deploying new technologies to safeguard national security does not create obvious 'policy losers' and pilot def/acc projects can likely deliver value very quickly.

3. **Strategic advantage:** Exciting def/acc options already exist that could greatly enhance the UK's ability to detect, prevent and respond to chemical and biological threats and demonstrate to international partners that the UK is a world-leader in multiple areas of (AI-enabled) science and technology.

Operationalising defensive acceleration

Previous attempts to define and deliver defensive acceleration have remained very general. For example, Matt Clifford, Chair of UK ARIA and former Adviser on AI to the UK Prime Minister, explicitly summarises it very broadly: "[...] it's about building the infrastructure for the future where our values win" (Clifford, 2025). He highlights excellent areas where a def/acc mindset might be useful, including AI-enabled biodefence, but does not provide specific examples.

This generality makes it harder to identify which strategic priorities are most crucial to address first, and—within a specific domain—which projects are likely to be most amenable to acceleration and provide the most defensive benefits.

We suggest a three-step framework to implement a def/acc agenda:

1. **Strategic and tactical prioritisation:**
    a. Ruthlessly focus on those domains where acceleration is most crucial to safeguard national security. This likely involves making painful trade-offs. It is especially important to consider the scale of potential risks—we recommend starting by prioritising the highest-impact ('Catastrophic') risks within the National Security Risk Assessment.

    b. Within these areas, carefully interrogate and identify which technologies are most likely to offer significant net defensive benefits (after considering how such technological development could advantage attackers). For more information on the *offence–defence balance*, see Box 1.

2. **Differential technological development:** For a given strategic domain (1a) and set of promising defensive technologies (1b), develop a subset of those technologies where state intervention is expected to accelerate development the most, compared to a market-neutral status quo. We identify three such types of acceleration:
    a. *First*: Choose technologies that otherwise would not be researched and commercialised by the market (often because they generate large positive externalities to society in the form of improved safety and security). Focus on capturing market share and monopolising critical components by building supply chain segments domestically.

    b. *Faster*: Choose defence-biased technologies that can be developed and deployed more quickly than equivalent offence-biased ones.

    c. *Further*: Choose technologies that are especially amenable to rapid and scalable deployment by defensive actors. Impressive prototypes that are expensive and slow to access and use do not actually provide defensive uplift or enhance national security.

3. **Real-world uplift by design**: During early-stage technological development, focus relentlessly on identifying users, sharing early ideas and gathering feedback. Make an explicit plan, detailing how the new technology will be provided to and used by key defensive actors. Prioritise resourcing this deployment plan before attempting to develop a new round of technologies.

Box 1: The offence–defence balance

The offence–defence balance describes whether, within a given domain, attackers or defenders hold a structural advantage. In some domains—such as cybersecurity—offence is often considered dominant: attackers need only find one vulnerability, while defenders must protect every possible entry point. In other domains—such as conventional territorial defence—defenders may hold advantages through fortification, knowledge of terrain and shorter supply lines.

The offence–defence balance is not fixed: technological change can shift it in either direction. The development of nuclear weapons dramatically advantaged offence in strategic warfare; encrypted communications shifted some cyber capabilities back towards defence. For def/acc, the offence–defence balance matters because it determines which technologies are worth accelerating. In offence-dominant domains, defensive technologies that can shift the balance towards safety—even marginally—are especially valuable.

Policymakers must consider whether a given technology is 'dual-use': if accelerating development advantages attackers as much as defenders, the net security benefit may be zero. The goal is to identify technologies that disproportionately benefit defenders—'defence-biased' technologies—and prioritise these for acceleration.

Additionally, adversaries could misidentify defensive preparations as offensive activity, potentially triggering escalatory responses. Transparency is essential to signal intent. Def/acc programmes should be designed with openness by default wherever possible (e.g. publishing methodologies, sharing findings with allies and engaging with multilateral mechanisms such as the Biological Weapons Convention). This helps reduce misperception risks whilst building international norms in favour of defensive acceleration.

Chemical and biological defence is a priority defensive acceleration area

The Strategic Defence Review (SDR) 2025 identifies chemical and biological defence R&D as the 'urgent and essential activity' for UK Defence R&D. We strongly concur, especially given the rapid acceleration of LLMs, AI agents and specialised AI-enabled scientific tools, as well as the recognition by industry that frontier AI systems can potentially meaningfully uplift novice actors building chemical and biological weapons (OpenAI, 2025; Anthropic, 2025).

The SDR's focus on CB defence—coupled with the National Security Risk Assessment's pandemic scenario, which estimates up to 840,000 deaths in the UK as a reasonable worst-case—directly addresses Step 1a. Biology is generally an offence-dominant domain (Aveggio et al., 2025), which complicates identifying defence-biased technologies in Step 1b. Aveggio et al. highlight that biological attackers generally have advantages in:

i. **Time for pathogen spread**: Transmissible pathogens can self-replicate and transmit rapidly over large distances, unlike therapeutics and vaccines
ii. **Financial burden**: Much less biological material is required to attack (e.g. a single vial) rather than defend (e.g. billions of vaccine doses)

iii. **Threat surface**: Attackers choose one of more than 1,500 known human-infecting pathogens (not even including future potentially engineered ones) while defenders have to defend against this entire spectrum

iv. **Consequences of failure**: Attackers can keep attacking until they are caught or deterred, but (at least for pandemic threats) defenders cannot afford to let even one attack through.

They highlight that defenders have one advantage:

v. **Access to knowledge and materials**: Defenders have much greater funding, infrastructure and the ability to (mostly) operate in the open, while attackers have to act covertly and avoid detection with fewer resources.

We additionally note a second defender advantage[2]:

vi. **Permanent environmental hardening**: Defenders have the potential to modify their environment to permanently prevent pathogen build-up and spread (e.g. via installing ventilation, always-on sterilisation equipment and physical barriers)

Considering these advantages, we generated 22 interventions that could help meaningfully accelerate chemical and biological defence and other national security priorities in the UK (see Table 1). These are categorised according to whether they primarily support:

- Chemical and biological defence
- Intelligence advantage
- Privileged defensive access
- Cross-cutting benefits

Box 2: AI-enabled interventions are often, but not always, an especially strong option

Many of the def/acc interventions we propose involve using frontier AI systems to enhance existing defensive activities or to more quickly and effectively develop and deploy new ones. While AI is not necessary to accelerate some aspects of chemical and biological defence—and there is a risk that policymakers could inadvertently favour 'shiny' interventions that are only tangentially AI-enabled—there are significant opportunities for AI to enhance security and intelligence technologies and practices.

With respect to chemical and biological defence, AI systems have already revolutionised aspects of modern life sciences research, most notably protein structural prediction and design: the 2024 Nobel Prize in Chemistry was awarded for the development of AI scientific tools like AlphaFold, ProteinMPNN and RFdiffusion. These have transformed protein science: structural identification tasks that would have taken a PhD researcher in the early 2000s years can now be completed with a single line of code.

This strongly suggests that AI systems could markedly improve some technical CB defence activities, especially those related to -omics–based biosurveillance, microbial forensics and attribution, as well as vaccine and therapeutic design.

[2] We are not the first to note the possibility of permanent defender advantage via changes to the built environment. This has been explored extensively in military contexts under 'collective protection', and there are a number of other in-depth investigations (see, for example, Blueprint Biosecurity, 2025)

Table 1: A long-list of defensive acceleration ideas across four categories with HMG's role specified. In bold are four interventions we consider in depth in this report.

Intervention	Category	HMG role
Develop automated vaccine design capabilities	Chemical and biological defence	Support + coordinate
Develop automated therapeutic discovery capabilities	Chemical and biological defence	Support + coordinate
Develop world's-first AI-enabled environmental metagenomics capability	Chemical and biological defence	Lead
Enhance microbial forensics and attribution activities with frontier AI	Chemical and biological defence	Lead
Enhance UK Biothreats Radar and promote international expansion	Chemical and biological defence	Lead
Develop and deploy next-generation AI-enabled customer and sequence screening mechanisms for nucleic acids	Chemical and biological defence	Support + coordinate
Enhance clinical metagenomics (mSCAPE) programme with frontier AI	Chemical and biological defence	Support + coordinate
Develop AI-enabled clinical anomaly detection for high-consequence transmissible pathogen symptoms	Chemical and biological defence	Support + coordinate
Develop new standards and provide funding for rapid testing and validation of next-generation biosecurity technology	Chemical and biological defence	Support + coordinate
Retool existing AI capabilities towards the highest-priority threats	Chemical and biological defence	Support + coordinate
Identify priority AI-enabled automation options for vaccine, diagnostic and therapeutic creation as part of 100 day mission (100DM)	Chemical and biological defence	Support + coordinate
Develop and deploy automated AI scientific tool capability evaluation	Chemical and biological defence	Lead
Pre-release access programmes to frontier AI for national security	Privileged defensive access	Support + coordinate
Randomised controlled trials for defensive actor uplift	Privileged defensive access	Lead
Provide defensive actors with maximally helpful 'rails-free' frontier AI models	Privileged defensive access	Lead
Reserve national compute allocation for defensive actors	Privileged defensive access	Lead
Ensure government officials use the most advanced AI models (closed- and open-weight)	Cross-cutting benefits	Lead
Promote culture of defensive innovation and raise awareness of historical dual-use developments	Cross-cutting benefits	Lead
Develop and deploy AI-enabled tabletop exercises	Cross-cutting benefits	Lead
Develop automated integration capability for all-source threat assessment	Intelligence advantage	Lead
Develop personalised red-teaming capability for intelligence analysts	Intelligence advantage	Lead
Build technical capability for privacy-preserving cross-domain query integration	Intelligence advantage	Lead

Building priority concept notes for CB defensive acceleration

We identified four priority interventions—bolded in Table 1, above—that we feel are especially promising:

1. **UK SONAR:** AI-enabled environmental metagenomics for engineered pathogen biosurveillance
2. **UK PROVE:** Expanded testing and validation mechanisms for next-generation biosecurity technologies
3. **UK SCAN:** Enhanced AI-enabled screening systems to accelerate legitimate access to dual-use AI models and synthetic nucleic acids
4. **UK TRUST:** Pre-release 'Trusted Tester' programmes to share frontier AI capabilities with defensive actors working on critical national security priorities

In this report, we provide concept notes with additional detail on each of these interventions, focusing particularly on the economic and strategic benefits that can provide the UK and sketching a potential implementation approach.

Given the rapid pace of development in AI, engineering biology and their intersection, we believe it is crucial that policy ideation and development move faster. This ensures that targeted, well-scoped governance can help steer these powerful emerging technologies for the public good. As such, we are sharing ideas in early forms to kick-start that process.

We are aware that additional work would be necessary to develop each priority idea further into a formal business case, but early feedback is essential to ensure that policy is 'effective by design'. We provide approximate economic costings but appreciate that more detailed analysis would be required.

We are excited to iterate on these ideas with colleagues in government and industry and help develop and deliver new capabilities that are essential to keeping the UK safe from chemical and biological threats.

UK SONAR: An AI-enabled pandemic early warning system

Concept

UK biosecurity faces a serious challenge in its ability to rapidly and reliably detect novel and high-consequence pathogens at scale, prior to members of the public falling ill and presenting to hospital. This poses important questions for our national security given increasing biological threats, including from state actors[3,4] and AI-enabled developments.[5] While the UK is currently growing its world-leading *clinical* biosurveillance capabilities, notably the UK Health Security Agency (UKHSA) mSCAPE initiative, such programmes are small in scale and leave a serious gap in pre-clinical environmental detection. To be truly resilient to current and future biological threats, accelerating early warning is critical.

The UK has an opportunity for global leadership in biosecurity and AI-enabled biotech. The U.S. Nucleic Acid Observatory (NAO) has been developing and testing pathogen-agnostic metagenomic surveillance technology since 2021, which can detect both known and novel pathogens from wastewater alone. But NAO has not yet deeply integrated new AI technologies that can dramatically improve both speed and performance. The UK could build a new global gold standard for biosurveillance by:

- ➔ Adopting an AI-first approach to metagenomic wastewater sequencing
- ➔ Integrating with NHS systems and existing wastewater surveillance infrastructure

The UK should establish a Surveillance Observatory for Nucleic Acid Recognition (SONAR) using AI-enhanced metagenomic surveillance to detect novel biological threats before they cause widespread harm to the UK. SONAR would monitor environmental samples from major airports and urban centres to create an early-warning system for both known and unknown biological threats. We could identify emerging pathogens—whether natural, accidental or deliberately engineered—days to weeks before conventional surveillance methods. By integrating advanced AI systems from the outset, UK SONAR would use anomaly detection to automatically flag, investigate and escalate suspicious signals for human review.

A strategic pilot (one airport, 10 wastewater sites) would cost £15mn in Year 1. This could expand to six airports (covering all four nations and 51% of all UK air travellers), 111 public wastewater sites (50% of England's population) and military sites. This investment would position the UK as the undisputed global leader in biosurveillance—unique in collecting both clinical and environmental metagenomic health intelligence—and provide critical early warning capability against biological threats that AI may make easier to engineer.

Sequencing costs have fallen 100-fold over the past decade,[6] whilst AI capabilities for analysing genetic sequences have improved dramatically in the last two years. The UK can now build operational AI-enhanced metagenomic surveillance at scale. This can be done in an anonymised fashion exclusively using aggregated data to prevent privacy concerns.

---

[3] National Security Strategy 2025: Security for the British People in a Dangerous World (Cabinet Office, 2025)
[4] 2025 Arms Control Treaty Compliance Report (U.S. Department of State, 2025)
[5] Safety and security risks of generative artificial intelligence to 2025, Annex B (Department for Science, Innovation and Technology, 2025)
[6] DNA Sequencing Costs: Data (U.S. National Human Genome Research Institute, 2022)

Economic and strategic case

Economic and healthcare value: *UK SONAR helps mitigate catastrophic economic consequences from pandemics, catalyses future R&D and improves day-to-day public health.*

1. The Covid-19 pandemic cost the UK upwards of £310 billion. Early pathogen detection helps mitigate severe economic disruption by accelerating decision-making and countermeasure development.
2. Data-driven public health systems save lives. Pathogen-agnostic metagenomic surveillance provides day-to-day benefits to the public and the NHS by:
    a. Measuring other disease prevalence (such as seasonal influenza), improving public health and vaccine campaign efficiency
    b. Gathering data on disease evolution to improve the design of future medical countermeasures and catalyse future scientific R&D

Strategic value: *UK SONAR would demonstrate international leadership, support other UK biosecurity efforts and could be exported to close allies.*

1. Building UK SONAR with an AI-first mindset would position the UK as the undisputed global leader in metagenomic biosurveillance and would:
    a. Strengthen the UK's position with allies and partners by providing novel—and often, unique—biological threat intelligence first
    b. Support the Biothreats Radar and National Biosurveillance Network build-out
2. UK SONAR would be an exportable capability: the UK can offer next-generation AI-enabled metagenomics surveillance to close partners, enhancing our security and generating revenue. AUKUS Pillar 2 might be one such export mechanism.

Comparative advantage: *The UK is best-placed globally to build AI-enabled environmental metagenomic surveillance, given existing strengths in genomics, AI security and attribution.*

1. World-leading pathogen genomics
    a. UKHSA mSCAPE is the world's first operational clinical metagenomics surveillance system, collecting real-time data from 30 NHS sites.
    b. UKHSA monitored 302 wastewater facilities for SARS-CoV-2, covering 74% of England's population and detecting infection 4–5 days earlier than in the clinic.
2. The UK is the world-leader in harnessing AI technical expertise for national security purposes through its AI Security Institute.
3. UK SONAR datastreams could feed directly into the UK Microbial Forensics Consortium and help train new state-of-the-art AI attribution technologies

Definitions	
Metagenomic surveillance	A biosurveillance technique that identifies both known and novel pathogens without specific targeting, usually from wastewater samples
Pathogen-agnostic	The ability to detect a pathogen—known or novel—with limited or no prior genetic information
Microbial forensics / attribution	The use of microbial genetic data and computational analysis to determine the origin and potential source of a biological agent and whether it was engineered

Implementation approach

Governance

UK SONAR should be led by the Defence Science and Technology Laboratory (Dstl) with a Senior Responsible Owner appointed within the Ministry of Defence (MOD), who should promote close collaboration with UKHSA and AISI. UK SONAR is closely related to existing MOD 'Detect' Pillar UK Biological Security Strategy activities, especially the UK Microbial Forensics Consortium (UKMFC). Similar to UKMFC and AISI, there should be an Advisory Board with significant industry expertise (e.g., Oxford Nanopore Technologies, Ginkgo Biosecurity) and non-profit expertise (e.g., U.S. Nucleic Acid Observatory), but programme ownership should remain solely within MOD, not shared across multiple Departments.

Clear backing from senior ministers is essential to enable a small team to move quickly. They should reuse existing authorities from the UKHSA Environmental Monitoring for Health Protection (EMHP) programme for wastewater surveillance, where appropriate, though new authorities may be necessary for airport surveillance. MOD should work quickly to facilitate additional surveillance options for military sites.

Operational model

Note: Budgetary and human resource estimates are approximate.

SONAR should be initialised with 9 FTE dedicated personnel (4 computational + 5 wet-lab) and an initial procurement budget of £5mn, with yearly running costs of £10–30mn. Over 3 years, SONAR should scale up personnel, airport surveillance and wastewater surveillance:

- Start with one major London airport (Heathrow) and scale to six major regional airports nationally to ensure coverage of all four nations (Manchester, Edinburgh, Belfast, Birmingham and Cardiff[7])
- Surveil 10 public wastewater sites (10% coverage of England) in Year 1, growing to 111+ (50%) in Year 3, plus additional MOD sites.
- Funding should draw from existing MOD budgets, but the Integrated Security Fund and the National Security Strategic Investment Fund may also be appropriate.

Time	Personnel (FTE)	Airports (UK coverage)[8]	Public wastewater sites (England coverage)[9]	Total yearly cost (£mn)	Additional milestones
Month 12	14	2 (39%)	26 (20%)	15	– Coordination with UKMFC – AI anomaly detection pilot
Month 24	24	4 (48%)	74 (40%)	23	– Data-sharing across HMG – Semi-automated attribution algorithms pipeline active – Additional MOD site pilot
Month 36	36	6 (54%)	111 (50%)	30	– Capability ready to export to close allies and partners – MOD sites scale-up

[7] Note that Bristol Airport passenger numbers are more than ten times that of Cardiff Airport, so wastewater sequencing at Bristol Airport is very likely important to ensure reliable coverage of whole UK population
[8] UK passenger coverage calculated by author using UK Civil Aviation Authority data (UK CAA, 2025)
[9] Wastewater site coverage calculated by author using EMHP programme database (UKHSA, 2021)

UK PROVE: A test-bed to accelerate next-generation biosecurity technologies

Concept

Emboldened adversaries and rapid advances in AI and engineering biology raise the spectre of future engineered pandemics that require cutting-edge technology to counter. Crucial biosecurity technologies—mRNA vaccines, monoclonal antibodies and rapid test kits—helped turn the tide in the UK's fight against Covid-19. This was possible only because of dedicated peacetime investment in strategic research, testing infrastructure and manufacturing capability. But deliberate, engineered biological threats require that the UK be able to rapidly deploy new, validated countermeasure capabilities at scale.

The science for next-generation biosecurity technologies already exists, but profound market failures prevent large-scale commercialisation. Promising technologies—including AI-designed broad-spectrum viral vaccines, advanced personal protective equipment and bioresilient infrastructure—sit in laboratories awaiting validation systems to bring them to deployment. Manufacturers lack clear demand signals and regulatory certainty, which will leave the UK reliant on outdated countermeasures during the next pandemic. No individual firm is incentivised to fund large-scale randomised controlled trials (RCTs) without guaranteed procurement, and the government cannot procure without validated efficacy data.

The UK should establish a Pandemic Resilience and Operational Verification Ecosystem (UK PROVE): a national testbed for critical biosecurity technologies to bridge the gap between pilot capability and deployable national security asset. UK PROVE would coordinate RCTs of promising technologies, develop new standards and link validated capabilities to Advanced Market Commitments (AMCs) for national procurement. Testing and standardisation enables uptake and scaling, which generates economic benefits and reduces unit costs. By absorbing early-stage regulatory risk, UK PROVE unlocks private sector investment and positions the UK as the standard-setter for NATO biosecurity capabilities.

A strategic programme would cost £20mn per year, focused primarily on trial funding, with coordination activities requiring minimal additional resources. This investment would catalyse domestic innovation in strategically critical technologies, improve UK resilience to biological threats and generate peacetime health benefits.

Definitions	
Broad-spectrum vaccines	Vaccines designed to provide protection against multiple strains or entire families of viruses, rather than single specific pathogens
Bioresilient infrastructure	Building systems and technologies—such as Far-UVC sterilisation and advanced air filtration—that reduce pathogen transmission through environmental controls
Advanced Market Commitment (AMC)	A binding government commitment to purchase a specified quantity of a product once it meets defined standards, used to de-risk private sector development
Publicly Available Specification	A fast-track British Standards Institution document that can later convert into a full British Standard

Economic and strategic case

Economic and healthcare value: *UK PROVE activities enhance UK productivity in peacetime alongside delivering deployable national security capabilities to protect the UK in a crisis.*

1. The UK loses £56bn annually to workplace sickness absence.[10] Bioresilient infrastructure can reduce transmission of seasonal influenza, RSV and other respiratory pathogens. Reducing these losses offers substantial economic return.
2. UK PROVE uses public funds to de-risk technologies, not to design them from scratch. By providing regulatory certainty and demand signals through AMCs, the programme creates conditions for private capital to flow into commercialisation.
3. Domestic supply chains for critical technologies—such as elastomeric respirators—create jobs, boost economic growth and reduce dependence on fragile international supply chains that failed during the Covid-19 pandemic.

Strategic value: *UK PROVE establishes the UK as NATO's testbed for biosecurity technologies and captures global standard-setting authority.*

1. International partners, like the U.S., have larger manufacturing capacity but struggle with rapid regulatory coordination. UK PROVE can serve as an agile testbed for NATO, validating technologies and building critical sub-components domestically that can then scale through allied supply chains.
2. By establishing the first validated standards for next-generation biosecurity infrastructure, the UK can shape specifications for international markets. This drives export demand for UK engineering and biotechnology firms and positions British Standards as the international benchmark.
3. Validated technologies could be exported to close partners through mechanisms such as AUKUS Pillar II, generating revenue and strengthening alliance relationships.

Comparative advantage: *The UK combines a strong bioeconomy with regulatory agility, proven clinical trial capabilities and excellent standards-setting infrastructure.*

1. The UK has existing models (MHRA AI Airlock and AI Growth Lab) to rapidly iterate on AI-enabled medtech and novel vaccine and therapeutic designs. The centralised NHS enables the UK to approve and deploy novel biosecurity technologies faster than fragmented U.S. or EU systems.
2. The UK hosts cutting-edge capabilities in AI-driven biology through industry (e.g., Google DeepMind) and academia (e.g., MRC Laboratory of Molecular Biology, University of Oxford Jenner Institute) and has an existing foundation of broad-spectrum vaccine research.
3. The UK ranks in the top five globally for Phase I and II clinical trials.[11] The Covid-19 RECOVERY trial demonstrated that the UK can validate countermeasures faster and more rigorously than any other nation, enrolling over 47,000 participants and identifying dexamethasone as a life-saving treatment within months.

[10] The cost of working age ill-health and disability that prevents work (DWP, 2025)
[11] The road to recovery for UK industry clinical trials (Association of the British Pharmaceutical Industry, 2024)

Implementation approach

Governance

The Department of Health and Social Care (DHSC) should lead UK PROVE under its Pandemic Preparedness portfolio, but activity must necessarily be cross-cutting given biosecurity is a whole-of-society challenge. The UK Health Security Agency should lead on randomised controlled trials, coordinating closely with the MHRA for medical technology approvals and HSE for workplace safety guidance on bioresilient infrastructure.

Challenge-style public–private partnerships—similar to the Defence and Security Accelerator model—should drive innovation whilst maintaining government accountability. The British Standards Institution should support development of new standards, beginning with Publicly Available Specifications that can later convert to full British Standards as evidence matures.

The programme requires senior ministerial backing to enable rapid coordination across departments. Future funding under UKRI's Tackling Infections strategic theme and the MRC Infections and Immunity Board should prioritise engineered biological threats and broad-spectrum countermeasures. Building Advanced Market Commitments will require close collaboration with HM Treasury, with overall cross-departmental coordination driven by mechanisms such as the Preparedness Funders Forum.

Operational model

Note: Budgetary and human resource estimates are approximate.

UK PROVE should launch with a core team of 8 FTE (programme managers, regulatory experts and clinical trial leads) and an annual budget of £20mn, the majority of which funds trials. The programme should scale personnel and activities whilst maintaining lean central coordination.

Time	Personnel (FTE)	Potential priority activities
Month 12	8	– Publish a Publicly Available Specification for elastomeric respirators – Initiate Far-UVC human safety trial – Launch first DASA-style challenge call for broad-spectrum vaccine technologies
Month 24	10	– Elastomeric respirator PAS converted to draft British Standard – Far-UVC trial results submitted to HSE for guidance update – First Advanced Market Commitment announced for validated technology
Month 36	12	– Updated HSE guidance on Far-UVC deployment published – British Standard for elastomeric respirators finalised – First AUKUS Pillar II export of UK PROVE-validated technology

Funding should draw from existing DHSC Pandemic Preparedness budgets, with potential contributions from the Integrated Security Fund for technologies with clear national security applications. The relatively modest budget reflects UK PROVE's role as a coordinator and de-risker rather than a primary funder of R&D. More ambitious versions of the programme could trial more technologies simultaneously, as well as funding additional technical research.

UK SCAN: Validating legitimate access to biological tools and materials automatically with frontier AI

Concept

Customer screening is the digital perimeter of biosecurity: the process of verifying the identity of individuals ordering dangerous biological materials or accessing dual-use AI tools and confirming they have legitimate reason to do so. Currently, this perimeter is porous due to high costs: manual 'Know Your Customer' (KYC) checks are slow, labour-intensive and require specialist knowledge. Many LLMs—and almost all dual-use AI scientific tools—have no customer screening whatsoever.[12] While the landscape has several unknowns, less than 10% of nucleic acid providers verifiably confirm they screen customers to prevent biological misuse.[13]

Legitimate researchers face increasing barriers to use advanced AI tools for dual-use defensive work. Some frontier AI companies have implemented more stringent 'refusal safeguards' on chatbots and AI agents to prevent them providing weapons development advice. Targeted safeguards are absolutely necessary to prevent biological misuse, but sweeping restrictions—that fail to distinguish between users—could hamper medical countermeasure development. Consistent customer screening enables more nuanced refusal safeguards, accelerating defensive biosecurity activity.

Frontier AI agents can solve this problem cheaply at scale. Recent trials demonstrate that AI agents verify institutional affiliation, publication history, and sanctions status with accuracy comparable to or better than human screeners.[14] The UK should establish an AI-enabled Screening Customers for Access and Nucleic acids (UK SCAN) programme to scale this technology. UK SCAN would develop standards and provide funding to build an AI-enabled verification API that validates legitimate users' access to:

1. **Synthetic nucleic acids:** Automating customer verification for DNA synthesis providers to ensure dangerous genetic material is only shipped to authorised users.
2. **Dual-use AI models:** Verifying users of frontier AI and specialised dual-use AI tools, allowing tool providers to grant 'guardrails-free' access to legitimate scientists.

The private sector requires support to deliver this capability. Firm-level incentives to invest in screening R&D are weak, given security benefits accrue to the public as a whole. A centralised, public 'Digital Biosecurity ID' service is more efficient than a fragmented patchwork of private efforts (which increases compliance burden for individual scientists). Government involvement is necessary where classified information informs screening decisions. A strategic programme would cost £1–5mn per year, focused primarily on challenge funding and standards development. This modest investment reflects UK SCAN's role in seeding private-sector R&D rather than building government-owned infrastructure.

---

[12] Global Risk Index for AI-enabled Biological Tools (Webster et al., 2025)
[13] IBBIS Global DNA Synthesis Map (Ameneiros et al., 2025)
[14] Evaluating AI-Assisted Customer Verification for DNA Synthesis Screening (Acelas, Palya, Flyangolts, Fady & Nelson, forthcoming)

Definitions	
Customer Screening (KYC)	The process of verifying the identity and legitimacy of a customer before providing access to sensitive goods or services. This includes verifying institutional affiliation (e.g., university or company employment) and confirming that the stated research purpose aligns with the user's professional background.
Dual-use AI scientific tools	Specialised AI models trained on biological data (such as protein structures or genetic sequences) that can be used to accelerate beneficial medical research but could also be misused. These are a subset of all AI scientific tools.

Economic and strategic case

Economic value: *UK SCAN eliminates the trade-off between speed and security. By automating verification, UK biotech start-ups can "move fast and be safe", accessing materials and tools without friction and accelerating economic growth.*

1. AI-enabled screening reduces compliance costs by orders of magnitude, removing a major barrier to entry for small synthesis providers and software developers.
2. By defining standards for AI-enabled biosecurity KYC, UK SCAN contributes to the development of a broader AI assurance market, driving the growth of this emerging industry and promoting the adoption of safe and responsible AI.
3. There are significant export opportunities for AI-enabled screening: UK SCAN positions domestic firms to provide verification services to U.S. and EU markets, which are currently considering stricter screening mandates.

Strategic value: *UK SCAN positions the UK to shape emerging markets and demonstrates commitment to allies' security concerns.*

1. Maintaining visibility over dangerous gain-of-function research is a growing concern for the U.S. and other allies. UK leadership in screening demonstrates commitment to address priority security concerns and helps strengthen international relations.
2. Early investment in standards development for AI-enabled customer screening creates a first-mover advantage: British standard-setting can influence other markets, ensuring that British companies can easily operate internationally.

Comparative advantage: *The UK already has strong insurance and verification markets and AI assurance capabilities, as well as existing biosecurity guidance to build upon.*

1. London is the global hub for insurance and verification markets and the UK's established expertise in AI assurance extends naturally to AI–biosecurity verification.
2. Existing policy infrastructure—including gene synthesis screening guidance and the Academic Technology Approval Scheme (ATAS)—provides a foundation to design robust but light-weight automated customer screening systems.

Implementation approach

Governance

UK SCAN should be led by the Department of Science, Innovation and Technology (DSIT), reflecting its focus enabling safe digital infrastructure and existing synthesis screening guidance. Some technical R&D may need to happen internally, but the majority of capability-building should occur in the private sector through challenge funding. The AI Security Institute (AISI) and the National Cyber Security Centre (NCSC) should jointly develop the technical standards for automated verification agents, working closely with the National Protective Security Authority (NPSA) to identify appropriate personnel security standards for tool and nucleic acid access. AISI and NCSC should conduct red-teaming against the most promising proposed screening systems.

While the vast majority of checks can be automated using open-source data (e.g., publications, university and company directories), there will be rare cases where classified information indicates otherwise legitimate users should be considered higher-risk. UK SCAN should coordinate closely with other parts of the UK Intelligence Community (UKIC) to identify such instances and to facilitate secure data-sharing.

Operational model

Note: Budgetary and human resource estimates are approximate.

UK SCAN should launch with a budget of £1–5mn per year and a small dedicated team of 4 FTE within the DSIT AI Strategy & Preparedness Directorate, coordinating closely with DSIT Engineering Biology. It should draw on additional resource from AISI, NCSC and NPSA, where necessary. This low cost is possible because the programme is not building physical infrastructure, but rather software standards and APIs.

Time	Personnel (FTE)	Potential priority activities
Month 12	4	– Launch DASA-style challenge fund to develop AI–KYC agent – Begin developing standards for dual-use AI model access – Coordinate with UKIC to design ‘edge-case’ classified information-sharing protocol
Month 24	6	– Pilot AI-enabled customer screening mechanisms with UK-based DNA synthesis providers and AI–biology labs, benchmarking against human compliance officers – Run internal red-teaming exercise against the strongest AI–KYC agents – Finalise and publish model-access standards
Month 36	6	– Publish results of AI–KYC agent pilot – Deploy API-based ‘Digital Biosecurity ID’ that integrates AI–KYC agent and allows verified researchers to access nucleic acid providers and dual-use AI tools instantly – Continue regular red-teaming activities to assure real-world API deployment

Funding should draw primarily from DSIT budgets, with potential contributions from the Integrated Security Fund. The modest budget reflects UK SCAN’s role as a convener and standard-setter rather than a capability builder. Supplementary funding could support further in-house technical work to build next-generation screening systems directly. The programme could consider developing secure, automated systems that minimally query classified information in a privacy-preserving way to support customer verification.

UK TRUST: Unlocking privileged frontier AI access for national security

Concept

In offence-dominant domains like biology, defenders need every available advantage—but frontier AI currently reaches attackers and defenders at the same time. In practice, defenders in government often receive access *slower* than potential adversaries, due to vetting bottlenecks, liability concerns and cumbersome procurement processes. This matters because early access to frontier AI capabilities could meaningfully enhance defensive activities—from accelerating vaccine design to improving threat detection. Yet demand for such access is fragmented across HMG, industry and academia, meaning no single actor has sufficient leverage to negotiate priority access or favourable terms with frontier AI developers.

There is an opportunity to aggregate national security demand and negotiate privileged access for UK defenders. A single government broker could leverage existing security clearance infrastructure, build 'safe harbours' that address liability concerns for frontier AI developers, and negotiate 'Sovereign Tier' access agreements that provide UK defensive actors with early and enhanced access to frontier capabilities. This would help shift the offence–defence balance in the UK's favour and ensure national security actors are not technologically disadvantaged.

To do so, the UK should establish a Trusted Research Uplift and Secure Testing (UK TRUST) programme to privilege defenders with early access to frontier AI. UK TRUST would:

1. Build a 'Green Lane' for secure model access across academia, industry and HMG—including pre-approved API gateways to 'rails-free' AI models for vetted defensive researchers needing to do dual-use work [see UK SCAN above].
2. Promote two-way information exchange between national security 'Trusted Testers' and AI developers to provide feedback to improve AI model capabilities.
3. Coordinate and expand a programme of randomised controlled trials[15] to measure defensive uplift from frontier AI access, generating evidence to inform future AI design and defence investment.
4. Develop secure enclaves with locally-hosted model weights for the most sensitive national security applications.[16]

The private sector alone cannot solve this coordination problem because:

- Frontier AI developers face liability risks in providing unrestricted access;
- Individual government departments lack the scale to negotiate priority terms; and
- No market mechanism exists to aggregate and prioritise national security demand.

A strategic programme would cost £5–10mn per year for secure API infrastructure, vetting administration and running uplift trials. This investment would accelerate beneficial UK research, de-risk UK-first deployment for future frontier AI systems and demonstrate to allies that the UK is a trusted partner for sensitive capability development.

[15] See Microsoft 365 Copilot Experiment: Cross-Government Findings Report (GDS, 2025) and AI coding assistant trial: UK public sector findings report (DSIT & GDS, 2025)
[16] See Securing AI Model Weights (Nevo et al., 2024) for more details

Definitions	
Rails-free AI model	An AI model without the safeguards that prevent discussion of sensitive topics, enabling legitimate dual-use defensive research (such as pandemic pathogen countermeasure design)
Defensive uplift	Measurable improvement in defensive actors' ability to detect, prevent or respond to threats as a result of access to advanced AI capabilities
Secure enclave	An isolated computing environment with enhanced security controls, enabling work with sensitive AI models or data

Economic and strategic case

Economic value: *UK TRUST accelerates beneficial research and positions the UK as a priority deployment market for frontier AI.*

1. Early access to frontier AI accelerates UK research and development across defence, health security and other national priorities—generating economic returns and improving the UK's ability to respond to emerging threats.
2. By building trusted infrastructure for sensitive AI deployment, UK TRUST de-risks UK-first release for frontier AI developers. This makes the UK a more attractive market for early access to future capabilities and further AI-driven economic growth.
3. Evidence from defensive uplift trials would inform more targeted investment in AI-enabled national security capabilities, improving the efficiency of public spending.

Strategic value: *UK TRUST demonstrates UK leadership to allies, supports UK AI sovereignty and diversifies the UK's AI strategy.*

1. A well-functioning Trusted Tester programme creates virtuous feedback loops: UK defensive actors gain early access, provide valuable feedback to developers, and demonstrate to allies that the UK is an expert testbed for frontier capabilities.
2. UK TRUST diversifies the UK's AI strategy beyond the foundation model capability race. Even if the UK cannot match the U.S. or China in developing frontier models, it can lead the ‘defensive capability-integration race’—rapidly deploying AI to enhance existing national security functions.

Comparative advantage: *The UK has existing Trusted Tester infrastructure and close relationships between government and frontier AI developers.*

1. Google DeepMind, headquartered in London, operates one of the first Trusted Tester programmes globally for its advanced biological science capabilities. The UK can build on this existing relationship and expertise.
2. The AI Security Institute pre-deployment testing programme demonstrates that the UK can operate trusted access programmes at scale and provides unique value to frontier AI developers for enhancing AI safety. UK TRUST would expand on this model and unlock new public–private partnerships across government to enhance UK national security with AI.

Implementation approach

Governance

UK TRUST requires coordination across multiple parts of government. The Sovereign AI Unit should lead negotiations with frontier AI developers for 'Sovereign Tier' access agreements. Cabinet Office support is critical to prioritise appropriate vetting for non-government security researchers (Government Security Group) and coordinate cross-departmental activities (National Security Secretariat). AISI and NCSC should provide technical infrastructure for secure access and enclave development, and establish monitoring and auditing functions for query logs to provide an additional safety layer. Importantly, it should be clear to all parties that rail-free access is for defensive purposes to avoid mischaracterisation by adversaries.

Clear ministerial backing is essential to enable rapid coordination. UK TRUST should establish a simple, fast-track vetting process for defensive researchers, both in and out of government, identifying the minimum assurance required for different national security domains. While aspects of chemical and biological defence are especially sensitive, empowering university researchers with frontier AI does not necessarily require the highest level of clearance.

Operational model

Note: Budgetary and human resource estimates are approximate

UK TRUST should launch with a core team of 6 FTE (programme management, vetting specialists and technical infrastructure leads) and an annual budget of £5–10mn, split across secure infrastructure, vetting administration and uplift trials.

Time	Personnel (FTE)	Potential priority activities
Month 12	6	– Establish Green Lane API gateway pilot with one frontier AI developer – Launch first defensive uplift RCT (e.g. Dstl threat assessment acceleration) – Begin scoping secure enclave requirements with NCSC – Identify minimum clearance requirements for model access outside HMG
Month 24	10	– Expand Green Lane to additional frontier AI developers, including closed-weight AI scientific tools – Finalise first defensive uplift trial results, including external published summary – Pilot operational secure enclave for most sensitive applications
Month 36	12	– Green Lane access available across priority HMG departments, defence contractors and cleared academic researchers – Sovereign Tier access agreement negotiated with at least three major frontier AI developers – Evidence base established to inform future defensive AI investment

Funding should be drawn from Cabinet Office and DSIT budgets, where possible, with additional capacity from NCSC for technical development, as appropriate. The programme should explore cost-sharing arrangements with frontier AI developers, who benefit from structured feedback on their models' performance in high-stakes applications.

Conclusion

Managed access and safeguards will not prevent determined actors from using AI to uplift attempts to develop chemical and biological weapons. Open-weight release of frontier-capable models makes this almost certain over the next two years.

The UK must pursue defensive acceleration—not as an alternative to responsible AI governance, but as a necessary complement to it. This is vital to achieve security in the short- and medium-term against a growing threat landscape.

This report outlines a framework for operationalising defensive acceleration, through:

1. Strategic prioritisation of the highest-consequence threats;
2. Differential technological development, identifying where state intervention can accelerate progress; and
3. Relentless focus on real-world deployment to defensive actors.

We have applied this framework to chemical and biological defence—identified by the Strategic Defence Review as the most urgent and essential research activity—and generated 22 promising interventions, developing concept notes for four highest-priority options:

- **UK SONAR** would provide AI-enabled environmental biosurveillance to detect novel pathogens before they cause widespread harm.
- **UK PROVE** would establish a national testbed to validate next-generation biosecurity technologies and link them to procurement.
- **UK SCAN** would develop AI-enabled screening systems to secure access to biological tools and materials whilst accelerating legitimate research.
- **UK TRUST** would aggregate national security demand and negotiate privileged frontier AI access for UK national security actors.

These interventions are mutually reinforcing: UK SCAN builds a capability that enables UK TRUST, which in turn accelerates the activities of UK SONAR and UK PROVE. Together, these programmes address critical gaps across the detect–prevent–respond spectrum and would cost approximately £40–65mn per year. This offers significant value for money given the security benefits of enhanced UK chemical and biological defence.

Defensive acceleration can genuinely strengthen both national security and economic competitiveness simultaneously. The UK possesses unique comparative advantages to lead in this space and build exportable capabilities for international partners, generating revenue and strengthening allied biosecurity capabilities. Shifting the offence–defence balance now is crucial, before the next crisis forces reactive measures that cost far more in lives and resources.

Realising this potential requires solving critical market failures. Firms lack incentives to invest in next-generation biosecurity technologies and funding for large-scale trials is scarce without guaranteed procurement. Frontier AI developers face liability risks in providing unrestricted access to defensive researchers. Government intervention—through challenge funding, standard-setting, demand aggregation and safe harbours—can unlock private investment and accelerate capabilities that the market alone will not deliver.

References

Acelas A., Palya H., Flyangolts K., Fady P–E., & Nelson C., Forthcoming. 'Evaluating AI-Assisted Customer Verification for DNA Synthesis Screening'. *The Centre for Long-Term Resilience*.

AI Security Institute, 2025. *Managing risks from increasingly capable open-weight AI systems.* Available at: https://www.aisi.gov.uk/blog/managing-risks-from-increasingly-capable-open-weight-ai-systems

Ameneiros, A. et al., 2025. 'IBBIS Global DNA Synthesis Map'. Available at: https://globalsynthesismap.bio/project

Anthropic, 2025. *Claude Opus 4.5 System Card*. Available at: https://assets.anthropic.com/m/64823ba7485345a7/Claude-Opus-4-5-System-Card.pdf

Association of the British Pharmaceutical Industry, 2024. *The road to recovery for UK industry clinical trials*. Available at: https://www.abpi.org.uk/publications/the-road-to-recovery-for-uk-industry-clinical-trials/

Aveggio, C., Patel, A.J., Nevo, S. & Webster, K., 2025. 'Exploring the Offense-Defense Balance of Biology: Identifying and Describing High-Level Asymmetries.' *RAND Corporation.* Available at: https://doi.org/10.7249/PEA4102-1

Blueprint Biosecurity, 2025. *Blueprint for Far-UVC*. Available at: https://blueprintbiosecurity.org/far-uvc-preprint-post/

Cabinet Office, 2025. *National Security Strategy 2025: Security for the British People in a Dangerous World*. Available at: https://www.gov.uk/government/publications/national-security-strategy-2025-security-for-the-british-people-in-a-dangerous-world/

Carter, S. R. & Butchello, G., 2026. 'A Framework for Managed Access to Biological AI Tools'. *Nuclear Threat Initiative.* Available at: https://www.nti.org/analysis/articles/a-framework-for-managed-access-to-biological-ai-tools/

Clifford, M., 2025. 'What is def/acc anyway?'. *def/acc*. Available at: https://defacc.substack.com/p/what-is-defacc-anyway

Department for Science, Innovation and Technology & Government Digital Service, 2025. *AI coding assistant trial: UK public sector findings report*. Available at: https://www.gov.uk/government/publications/ai-coding-assistant-trial/ai-coding-assistant-trial-uk-public-sector-findings-report

Department for Science, Innovation and Technology, 2025. *Safety and security risks of generative artificial intelligence to 2025, Annex B*. Available at: https://www.gov.uk/government/publications/frontier-ai-capabilities-and-risks-discussion-paper/safety-and-security-risks-of-generative-artificial-intelligence-to-2025-annex-b

Department for Science, Innovation and Technology, 2025. *Digital and Technologies Sector Plan.* Available at: https://www.gov.uk/government/publications/digital-and-technologies-sector-plan

Department for Science, Innovation and Technology, 2025. *Seoul Ministerial Statement for advancing AI safety, innovation and inclusivity: AI Seoul Summit 2024.* Available at: https://www.gov.uk/government/publications/seoul-ministerial-statement-for-advancing-ai-safety-innovation-and-inclusivity-ai-seoul-summit-2024/seoul-ministerial-statement-for-advancing-ai-safety-innovation-and-inclusivity-ai-seoul-summit-2024

Department for Work and Pensions, 2025. *The cost of working age ill-health and disability that prevents work*. Available at: https://www.gov.uk/government/statistics/the-cost-of-working-age-ill-health-and-disability-that-prevents-work/the-cost-of-working-age-ill-health-and-disability-that-prevents-work

Foreign, Commonwealth and Development Office, 2024. *UK sanctions Russian troops deploying chemical weapons on the battlefield.* Available at: https://www.gov.uk/government/news/uk-sanctions-russian-troops-deploying-chemical-weapons-on-the-battlefield

Frontier Model Forum, 2025. *Preliminary Taxonomy of AI-Bio Misuse Mitigations.* Available at: https://www.frontiermodelforum.org/issue-briefs/preliminary-taxonomy-of-ai-bio-misuse-mitigations/

Government Digital Service, 2025. *Microsoft 365 Copilot Experiment: Cross-Government Findings Report*. Available at: https://www.gov.uk/government/publications/microsoft-365-copilot-experiment-cross-government-findings-report

Intelligence Analysis, 2025. *Explaining Uncertainty in UK Intelligence Assessment*. Available at: https://www.gov.uk/government/publications/explaining-uncertainty-in-uk-intelligence-assessment/explaining-uncertainty-in-uk-intelligence-assessment

Ministerie van Defensie, 2025. *Russia further intensifies its use of chemical weapons in Ukraine.* Available at: https://english.defensie.nl/latest/news/2025/07/04/russia-further-intensifies-its-use-of-chemical-weapons-in-ukraine

Ministry of Defence, 2025. *The Strategic Defence Review 2025 - Making Britain Safer: secure at home, strong abroad.* Available at: https://www.gov.uk/government/publications/the-strategic-defence-review-2025-making-britain-safer-secure-at-home-strong-abroad

Nevo, S. et al., 2024. 'Securing AI Model Weights'. *RAND Corporation.* Available at: https://doi.org/10.7249/RRA2849-1

OpenAI, 2025. *GPT-5 System Card*. Available at: https://cdn.openai.com/gpt-5-system-card.pdf

OpenAI, 2025. *Preparing for future AI capabilities in biology.* Available at: https://openai.com/index/preparing-for-future-ai-capabilities-in-biology/

Parthemore, C. & Weber, A., 2026. 'Op-ed: Biodefense is Core Defense'. *Defense News*. Available at: https://councilonstrategicrisks.org/2026/01/08/op-ed-biodefense-is-core-defense/

Rose, S. & Nelson, C., 2023. 'Understanding AI-Facilitated Biological Weapon Development'. *The Centre for Long-Term Resilience.* Available at: https://doi.org/10.71172/nm7j-qzt1

UK Civil Aviation Authority, 2025. *UK Airport Data 2025*. Available at: https://www.caa.co.uk/data-and-analysis/uk-aviation-market/airports/uk-airport-data/uk-airport-data-2025/

UK Health Security Agency, 2021. *Wastewater testing coverage data for 19 May 2021: EMHP programme*. Available at: https://www.gov.uk/government/publications/wastewater-testing-coverage-data-for-19-may-2021-emhp-programme

U.S. Department of State, 2024. *Imposing New Measures on Russia for its Full-Scale War and Use of Chemical Weapons Against Ukraine*. Available at: https://2021-2025.state.gov/new-measures-in-response-to-russias-full-scale-war-and-use-of-chemical-weapons-against-ukraine/

U.S. Department of State, 2025. *2025 Arms Control Treaty Compliance Report*. Available at: https://www.state.gov/wp-content/uploads/2025/04/2025-Arms-Control-Treaty-Compliance-Report_Final-Accessible.pdf

U.S. National Human Genome Research Institute, 2022. *DNA Sequencing Costs: Data*. Available at: https://www.genome.gov/about-genomics/fact-sheets/DNA-Sequencing-Costs-Data

Webster, T., Moulange, R., Del Castello, B., Walker, J., Zakaria, S. & Nelson, C. 2025. 'Global Risk Index for AI-Enabled Biological Tools'. *The Centre for Long-Term Resilience & RAND Europe.* Available at: https://doi.org/10.71172/wjyw-6dyc

About The Centre for Long-Term Resilience (CLTR)

CLTR is an independent think tank with a mission to transform global resilience to extreme risks. We do this by working with governments and other institutions to improve relevant governance, processes, and decision-making. We are based in the heart of Whitehall at Public Hall. Read more about our work at www.longtermresilience.org

The Centre for Long-Term Resilience is a non-profit company registered in England and Wales under the name Alpenglow Group Limited (12308171). Our registered office address is 71-75 Shelton Street, Covent Garden, London, England, WC2H 9JQ.

Contact us

info@longtermresilience.org