



Biodefense in the Intelligence Age

An Action Plan for AI-Powered Biological Resilience

June 2026

The pace of scientific discovery is accelerating. In the Intelligence Age, scientists and public-interest institutions can use advanced AI to synthesize evidence, reason across scientific domains, and analyze complex data. These capabilities can help expert teams move faster from question to validated evidence, identify emerging threats¹ earlier, and develop new countermeasures that protect people.

Advanced AI capabilities for biology are improving rapidly and becoming increasingly available across the scientific ecosystem. In April 2026, OpenAI [introduced GPT-Rosalind](#), our frontier reasoning model built to support research across biology, drug discovery, and translational medicine. We believe advanced AI can help scientists better understand biology, develop new therapies, and improve human health across a wide range of diseases – and Rosalind is a reflection of our commitment to achieve transformational impact that improves health outcomes.






But the same capabilities that can help scientists better understand disease, develop new therapies, and improve human health also have implications for biological security. In May 2026, [we announced Rosalind Biodefense](#) to help trusted developers to build new biodefense and pandemic preparedness capabilities. In the face of a growing array of biological threats, and as advanced AI systems become more capable and more widely available, societies will need these capabilities to strengthen resilience and preparedness faster than they increase risk. OpenAI's strategy is to provide trusted public-interest institutions with these capabilities so they can use them effectively for defense, preparedness, and resilience before biological crises occur. We believe the best way to strengthen biological security is to equip responsible defenders with advanced capabilities while developing the safeguards, evidence, and governance needed for their safe deployment.

OpenAI is launching this biodefense action plan to help build a more resilient biological future: one in which societies can detect threats sooner, develop countermeasures more rapidly, and respond to crises with greater confidence and coordination. We will advance capability and safety together through trusted access, expert partnership, rigorous evaluation, and strong safeguards. This work complements OpenAI's broader life-sciences efforts, including accelerating discovery for cancer, infectious disease, and other areas of human health.

This is broader than any single model release or program. GPT-Rosalind and Rosalind Biodefense are the first implementation of this broader strategy and represent our commitment to work with governments, public-health institutions, research organizations, and responsible life-sciences partners in strengthening biological resilience. Through trusted partnerships, they will help generate evidence about how advanced AI can strengthen biological resilience while informing the safeguards, evaluations, and governance needed for responsible deployment. This action plan expands that work into a long-term agenda for AI-powered biological resilience.

¹ By biological threats, we mean naturally occurring outbreaks, accidental releases, and deliberate misuse of biological agents or enabling information that could harm people, agriculture, or ecosystems.

Our plan consists of 5 pillars:

				
1. Equipping defenders through trusted access	2. Accelerating medical countermeasures	3. Building earlier warning systems	4. Strengthening diagnostics, preparedness, and response	5. Measure impact, risk, and resilience

Our Commitments

OpenAI's biodefense strategy rests on five commitments: strengthening biological resilience, expanding trusted access for qualified defenders, advancing safeguards alongside capability, partnering with institutions that can validate and operationalize these tools, and building the evidence base needed for responsible deployment.

Strengthen biological resilience. OpenAI will support trusted experts and institutions working to speed medical countermeasures, enhance threat detection, improve diagnostics and response, and deploy useful biological AI capabilities for public benefit.

Expand Trusted Access. OpenAI will ensure that advanced biological AI capabilities are available to trusted defenders working in the public interest. We will build trusted-access pathways for institutions with the mission, expertise, governance, and safeguards needed to use these capabilities responsibly.

Develop safeguards and evaluations. OpenAI will advance capability and safety together. Biological capability evaluations, access controls, monitoring, abuse investigation, model-safety research, and partner feedback can reduce misuse risk while creating the conditions for responsible defensive deployment. This work is guided by our Preparedness Framework, which evaluates increasingly capable biological AI systems and informs how advanced capabilities are deployed responsibly.

Partner with institutions that can validate and deploy these capabilities. OpenAI does not view biodefense as a technology problem alone. Effective biological resilience depends on public-health institutions, scientific expertise, laboratory infrastructure, emergency-management systems,



manufacturing capacity, and international cooperation. We will work with public-health agencies, national laboratories, defense-science organizations, research funders, universities, responsible companies, and nonprofit biosecurity organizations to strengthen these capabilities with useful tools, evidence, and partnerships.

Build and share evidence. OpenAI will build and share evidence responsibly. We will develop evaluation methods that measure risk, safeguards, and public-benefit outcomes, and we will publish high-level findings and lessons where doing so strengthens the field without creating information hazards.

Why This Matters Now

The biological sciences are becoming increasingly important to economic prosperity, public health, and national security. The world faces a growing range of biological challenges: emerging infectious diseases, antimicrobial resistance, threats to food systems, supply-chain vulnerabilities, and the possibility of deliberate misuse of biotechnology. Many of these challenges are driven by factors independent of AI, including globalization, scientific advances, ecological change, and the growing interconnectedness of modern societies.

At the same time, advances in AI are becoming increasingly relevant to biology. We know from our own evaluations that advanced models can meaningfully assist biology workflows, and we expect these capabilities to continue diffusing across the ecosystem. OpenAI's response is to build the strongest defensive applications with trusted partners before less responsible actors can use similar capabilities for harm. Responding effectively requires the ability to detect threats early, understand them quickly, and develop countermeasures at unprecedented speed.

Biological crises can arise naturally, accidentally, or deliberately, and their consequences can reach across families, health systems, food supplies, economies, and national security. Advanced capabilities therefore must be developed responsibly and deployed with care, while remaining available to trusted defenders working in the public interest. Biology is physical: useful progress depends on laboratories, materials, manufacturing, clinical systems, public-health operations, and trust. That is why it is important to not only accelerate defensive uses of AI but also build a broader biodefense ecosystem, such as improving synthesis screening and supply chain resilience.

Preparedness and progress are not opposing goals. Effective safeguards, rigorous evaluations, and a clear understanding of risk make it possible to accelerate beneficial uses responsibly. By testing capabilities, validating mitigations, and working with trusted experts, we build the confidence needed to deploy advanced biological AI systems in support of public health, preparedness, and

resilience. . OpenAI's Preparedness Framework tracks biological and chemical capabilities alongside other frontier risks; this action plan applies that same approach to strengthening biological resilience before the next crisis.

Action 1: Equip Defenders Through Trusted Access

OpenAI will expand responsible access to advanced life-sciences capabilities for trusted defenders working on biological resilience. That includes government science and public-health teams, national laboratories, defense-science organizations, trusted academic groups, responsible companies, and nonprofit biosecurity organizations with the expertise and governance to use them well.

Trusted access is the deployment approach for this work. It gives trusted institutions a way to use stronger capabilities for clearly beneficial purposes while pairing them with access management, organizational governance, user controls, appropriate monitoring, and the ability to narrow or revoke access as circumstances change.

Rosalind Biodefense will be one early pathway under this plan, within a broader agenda for resilient biological systems. Through this effort, trusted institutions will be able to use GPT-Rosalind and future life-sciences models to support high-impact biodefense work, including countermeasure development, threat assessment, diagnostic development, early warning and detection, and response planning.

OpenAI has launched a Government Trusted Access for our frontier AI capabilities in life sciences for trusted government partners. The pathway will focus on bounded mission work, authorized users, strong data protections, and expert review of model-supported scientific outputs. Its purpose is to ensure that advanced biological AI capabilities are available to trusted defenders working in the public interest, with safeguards proportionate to the sensitivity of the work.

Action 2: Accelerate Medical Countermeasures

OpenAI, working with key partners, will help accelerate medical countermeasure research against naturally emerging, accidental, and deliberate biological threats, including through Rosalind Biodefense. Scientists, public-health experts, and biodefense teams increasingly have access to AI systems that can help synthesize evidence, identify promising research directions, analyze data, and accelerate routine scientific work. Used responsibly, these tools can help expert teams address bottlenecks that matter for public health and national security.

OpenAI will focus on projects where expert partners can validate outputs and turn model-supported insights into real scientific progress. The goal is to help capable teams move more quickly from scientific questions to validated evidence, from evidence to promising interventions, and from

promising interventions to countermeasures that can be evaluated and delivered responsibly. OpenAI will prioritize defensive and public-benefit applications consistent with our policies and safeguards.

Governments can partner in this work by identifying priority threat areas, connecting useful AI capabilities to public-health and biodefense research programs, funding validation environments, strengthening manufacturing and delivery pathways, and helping translate successful pilots into durable countermeasure capacity.

Action 3: Build Early Warning Systems

OpenAI, working with key partners, will support early-warning systems that help expert teams detect, interpret, and respond to biological signals sooner. Potential applications include metagenomic sequencing, anomaly detection across structured and unstructured data, triage of pathogen or sequence information, and synthesis of clinical, genomic, agricultural, environmental, and open-source indicators.²

Government, public-health, and agricultural-health partners can help define the most important operational workflows, provide appropriate validation environments, and determine where AI improves speed, signal quality, or coordination without creating new risks. This work should strengthen existing surveillance and response institutions, not displace them.

Action 4: Strengthen Diagnostics, Preparedness, and Response

OpenAI will support expert teams preparing for and responding to biological events. AI systems can help with diagnostic development, interpretation of complex evidence, scenario planning, operational playbooks, resource allocation, epidemiological modeling, screening, non-pharmaceutical interventions (NPIs), and coordination across scientific, medical, public-health, and emergency-response communities.

Preparedness depends on more than medical countermeasures alone. AI may also help accelerate the development and deployment of non-pharmaceutical interventions and resilience measures, including improved respiratory protection, clean-air technologies, environmental monitoring, infection-control practices, and other approaches that can reduce transmission and protect communities during biological events.

The value of these tools will be determined through practical use under real constraints. This work can include testing model-supported workflows in preparedness exercises, tabletop scenarios, and

² Approved uses do not include designing, planning, optimizing, conducting, or troubleshooting dangerous gain-of-function research, as defined in EO 14292 or any successor U.S. government framework.

bounded operational settings, with expert review and clear measures of whether readiness improves.

Governments can use these collaborations to improve readiness before a crisis. Exercises can identify where AI makes response planning faster or more rigorous, where human review remains essential, and where institutions need better data systems, training, procurement pathways, or coordination mechanisms to sustain resilience over time.

Action 5: Measure Impact, Risk, and Resilience

OpenAI will help build up our understanding for how AI affects biological resilience. This work will include capability evaluations, safeguards evaluations, operational pilots with domain experts, real-world impact measurement, shared benchmarks, and policy-relevant evidence about where AI adds value and where additional controls are needed.

Evaluation must cover both sides of the ledger: where models could increase misuse risk and where expert use of AI can improve defensive outcomes. A durable biodefense strategy requires credible measurements that distinguish promising demonstrations from validated progress and help policymakers and institutions make informed decisions about deployment.

OpenAI will share high-level findings and methods publicly where doing so improves public understanding, supports responsible governance, and does not create information hazards. Governments, national laboratories, academic centers, standards bodies, and other responsible partners can help validate results, define shared benchmarks, and build evidence that is useful across the biosecurity community. OpenAI will continue supporting independent evaluations and third-party access for trusted biosecurity organizations and government AI institutes. This work builds on existing collaborations with national laboratories, biosecurity organizations, and government AI safety institutes to better understand both the risks and defensive applications of advanced biological AI capabilities.

Partnering for Biological Resilience

Biological resilience depends on capable institutions and trusted cooperation. OpenAI will prioritize partnerships where the mission is clearly defensive or public-benefit oriented; partners bring relevant scientific, public-health, national-security, or operational capability; projects are bounded enough to evaluate; model-supported work is subject to expert review and real-world validation; and lessons can improve both impact and safety.

The strongest government partnerships will be focused collaborations with institutions that can validate science, protect sensitive data, connect findings to public-health or biodefense operations, and use AI-supported insights in responsible public-interest programs. OpenAI's role is to provide useful capabilities, evidence, and support within that broader ecosystem while ensuring that advanced biological AI capabilities remain available to trusted defenders working in the public interest.

Governments can strengthen biological resilience by identifying priority mission areas, creating secure environments for evaluation, funding validation and testing infrastructure, modernizing procurement pathways, supporting manufacturing and surge capacity, and connecting AI-enabled workflows to existing public-health and biodefense institutions. Governments can also help define the biological-resilience challenges where advanced AI can make the greatest operational difference and help translate successful pilots into durable capability.

Public-health agencies, national laboratories, research funders, defense-science organizations, and emergency-management institutions can help translate AI capability into real-world resilience by providing domain expertise, data governance, evaluation capacity, laboratory and manufacturing pathways, and routes to operational use. These institutions play a critical role in determining where AI improves speed, rigor, or coordination while maintaining the scientific standards and safeguards necessary for responsible deployment.

Industry, academia, and nonprofit biosecurity organizations can help build the broader ecosystem through validated workflows, training, independent evaluation, shared benchmarks, and practical tools that make beneficial science easier and risky misuse harder. OpenAI will continue working

with these partners to build the evidence base, governance frameworks, and operational experience needed to strengthen biological resilience over time.

What Success Looks Like

Success means more than better AI models or new programs. It means stronger institutions, faster scientific response, more effective countermeasures, earlier warning, and a world better prepared for biological challenges before they become biological crises. OpenAI's work should be measured against those outcomes.

Preparedness and progress must reinforce each other. The same evaluations, access controls, monitoring systems, trusted-access programs, and expert partnerships that reduce misuse risk also help responsible institutions use AI safely for public benefit.

This is the promise of biodefense in the Intelligence Age: an opportunity to strengthen humanity's ability to prevent, detect, and respond to biological threats. Realizing that opportunity will require careful governance, rigorous science, trusted partnerships, and sustained investment in resilience. OpenAI is committed to doing its part.

About OpenAI

OpenAI's mission is to ensure that artificial general intelligence benefits all of humanity. We're building AI to help people solve hard problems because by helping with the hard problems, AI can benefit the most people possible – through more scientific discoveries, better healthcare and education, and improved productivity. We're off to a strong start, creating freely available intelligence used by more than 900 million people around the world each week. We believe AI will scale human ingenuity and drive unprecedented economic growth and new freedoms that help people accomplish what we can't even imagine today.

Cover image created with ChatGPT