

Deterrence by Resilience

Securing the United States
and Promoting Innovation
in the Age of Bioconvergence

August 2026

**Elizabeth
Sherwood-Randall**

*Ashton B. Carter Visiting Professor
Harvard Kennedy School*

J. Michael McQuade

*Director, Program on Emerging
Technology, Scientific
Advancement & Global Policy*



HARVARD Kennedy School

BELFER CENTER

for Science and International Affairs



HARVARD Kennedy School

BELFER CENTER

for Science and International Affairs

Deterrence by Resilience

Securing the United States
and Promoting Innovation
in the Age of Bioconvergence

Contributing Authors

Elizabeth Sherwood-Randall, IB3 Project Lead

Michael McQuade, Beth Cameron, Ashish Jha, Shankar Sundaram, Trent Buatte, Chris Li,
Dylan Rogers, Samuel Steuart, Elizabeth Guo, Shawn Haq, June Lee, Jacob Levine,
Henrik Nolte, Rebekah Reed

Initiative on Bioconvergence, Biosecurity, and Bioresilience (IB3)

Program on Emerging Technology, Scientific Advancement, and Global Policy

Contents

Executive Summary	7
Chapter 1. Introduction: The Deterrence by Resilience Framework	13
Chapter 2. Mapping the Bioconvergence Threat	17
A. The Promise and Peril of Bioconvergence	17
B. Bioconvergence Risk	21
C. Trends in Biological AI Systems	24
D. Evaluating Risks from Biological AI Systems.	25
E. Analogizing Bioconvergence Risks	34
Chapter 3. Prevention	37
A. Bio Risks of General Purpose AI Models: Current State and Critical Gaps	37
Findings and Recommendations	53
B. Biological Data Governance	56
Findings and Recommendations	59
C. Governance of Biological AI Tools.	60
Findings and Recommendations	63
D. Nucleic Acid Synthesis Screening and Laboratory Regulation.	64
Findings and Recommendations	68
Chapter 4. Detection	70
A. Diagnosing Failures in the Existing Biological Monitoring Systems: The COVID-19 Pandemic as a Stress Test	71
B. Building Bio-Monitoring Systems for Continuous Readiness.	75
C. Building the System: Infrastructure and Integration	77
D. Financing and Governing a Biological Monitoring System	81
E. Conclusion.	83
Findings and Recommendations	83
Chapter 5. Attribution	90
A. Attribution as Deterrence	90
B. Current Challenges for Attribution	91

C. Capability Inventory: Where the United States Stands Today	95
D. The US Policy Context	97
E. Conclusion.	98
Findings and Recommendations	98
Chapter 6. Readiness and Response	101
A. Readiness	101
Findings and Recommendations	111
B. Response	113
Findings and Recommendations	130
Chapter 7. Resilience.	133
A. The Interdependent Relationship Between Response and Resilience.	133
B. Building Bioresilience Through AI, Functional Data, and Response Infrastructure . .	135
C. Industrial Resilience: Strengthening the Private Sector and Supply Chain	140
D. Expanding Global Resilience	147
Findings and Recommendations	152
Conclusion: A Call to Action.	157
Appendix 1. Consolidated Findings and Recommendations	161
Appendix 2. Legislative Proposals on Biosecurity and AI.	182
Appendix 3. Perspective on Liability for Bio-risks of AI Models	197
Appendix 4. Further Discussion of the Defense Production Act.	200
Appendix 5. Gaps for NSBRR and NIBR Implementation.	205

Acknowledgments

This initiative has benefited from generative contributions by and collaborations with many people. In our research phase, we consulted extensively with leading scholars and experts at numerous universities and research institutions, current and former senior government officials, pioneers in AI and biotechnology companies, innovative private sector CEOs and their teams, and others in related fields. I am grateful to everyone who shared their knowledge and perspective with us.

We were advised by a deeply experienced core group, including Beth Cameron, Drew Endy, Jill Hruby, Ashish Jha, Eric Lander, John MacWilliams, Jason Matheny, Ernie Moniz, and Shankar Sundaram. Each gave unstintingly of their time and ideas throughout the process and contributed their wisdom and substantive expertise at critical junctures. Several also authored key elements of the report, including Dr. Jha on biological detection systems, Dr. Cameron on attribution and on global cooperation, and Dr. Sundaram on building a national data foundry. Ben Buchanan, Matt McKnight, David Relman, and Nikki Romanik provided invaluable insights on understanding, evaluating, and meeting novel challenges. Jack Goldsmith generously illuminated the potential legal architecture that could advance our objectives, inspired inquiry into a wide range of legal opportunities, supervised students in pursuit of relevant research, and continually challenged conventional thinking.

Our brilliant research team of multidisciplinary Harvard University graduate students included Chris Li, who served as research coordinator while completing his doctorate in synthetic biology, along with Harvard Kennedy School students Trent Buatte and Dylan Rogers, and Harvard Medical School student Sam Steuart. Each brought tremendous energy, curiosity, creativity, and competence to our work. At the start of 2026 we added a group of excellent Harvard Law School students, including Elizabeth Guo, Shawn Haq, June Lee, Jacob Levine, and Henrik Nolte. With Trent's creative and deft guidance, this "law pod" expanded our consideration of legislative and regulatory options; he also brought prodigious talent to organizing and polishing the full report. Notably, each of the graduate students involved in IB3 is listed as a contributing author because of their essential roles in researching and drafting elements of the document. Our undergraduate research assistant, Sofia Detjen, prepared useful background materials and helped keep track of the many research and writing assignments that we generated in our weekly team meetings.

Meghan O'Sullivan, the director of the Harvard Kennedy School's Belfer Center for Science and International Affairs, first suggested I lead this initiative and has been a great champion of our efforts. Michael McQuade, director of Belfer's Emerging Technologies, Scientific Advancement, and Global Policy Program, has been my expansive thought partner and essential collaborator. With his rigorous physicist's brain, steady temperament, and generosity of spirit, he has made our endeavor possible and added immeasurable value from start to finish. Rebekah Reed, the program's associate director, improved the final product with her fine intellect and editing skills. Shannon Felton Spence raised our game with her spirited command of strategic communications and production savvy.

This report bears the imprint of the late Ash Carter, my close colleague for three decades, who served as US Secretary of Defense from 2015 to 2017 and, in the final chapter of his consequential

career, as director of the Belfer Center. Ash's pioneering work on nuclear security and threat reduction in the early 1990s—at another pivotal moment in human history—set the gold standard for this kind of project. We have modeled our initiative on the thorough research, technically informed analysis, strategic vision, and cool pragmatism that were hallmarks of his approach to tackling the hardest national security problems.

Finally, a word on a nonhuman partner: AI. Our team has experimented with several LLMs, including Microsoft CoPilot, GPT 5.5, and Claude Sonnet 4.6 and Opus 4.7, in a variety of ways to aid in preliminary investigation, organization, and streamlining. Any information generated by AI tools has been thoroughly fact-checked and validated against verified sources. The bottom line is that human beings have written every word of this report, and we bear the responsibility for it.

Elizabeth Sherwood-Randall

Cambridge, Massachusetts

August 2026

Executive Summary

The world has entered a transformational era defined by the accelerating convergence of biology, artificial intelligence (AI), bioengineering, and digital infrastructure. This process, which this report describes as **bioconvergence**, is revolutionizing how biological capabilities are discovered and developed, and it promises revolutionary advances in health, science, and economic productivity. Simultaneously, however, it is reshaping the biological threat landscape in ways that challenge traditional national security frameworks and create potentially catastrophic risks to the homeland and to allies and partners around the world.

Bioconvergence is generating a more distributed, broadly accessible, and rapidly evolving threat environment. Capabilities once limited to state-sponsored programs are increasingly available to new actors through advances in AI uplift for biological knowledge, AI-enabled biological design, declining costs of DNA synthesis and sequencing, and laboratory automation. These developments not only lower barriers to innovation but also lower barriers to entry and misuse. Traditional approaches to deterrence, including deterrence by denial and deterrence by punishment, are increasingly insufficient.

This report advances a new strategic framework for balancing the opportunities and risks of bioconvergence in the United States: **deterrence by resilience**. Deterrence by resilience seeks to reduce adversaries' confidence that biological misuse will succeed, scale, or produce a strategic advantage. It combines steps across the full lifecycle of biological threats—prevention, detection, attribution, readiness, response, and long-term resilience—into a mutually reinforcing, virtuous cycle that raises the expected costs of misuse while reducing the damaging consequences of biological incidents. Within this framework, innovation and security are not competing objectives; rather, the infrastructure that enables American leadership in biotechnology and AI also enables national biosecurity. This is not to say that government involvement will have no costs—regulation often carries compliance costs—but that increased safety also generates benefits for both the public and the private AI-enabled biotechnology ecosystem. On an ongoing basis, it will be necessary to evaluate the implications of regulatory actions for innovation and, to the greatest degree possible, to incentivize safety measures.

Organization of This Report

This report proceeds in six substantive chapters.

Chapter 2 Mapping the Bioconvergence Threat

This chapter assesses the evolving threat landscape created by bioconvergence, examining how advances in biotechnology and AI interact with longstanding risks from natural outbreaks, accidents, and deliberate biological misuse. It presents a structured assessment of potential threat actors and scenarios, and concludes that biological risks increasingly resemble a hybrid of the nuclear and cyber domains while remaining distinct from both.

Chapter 3 Prevention

This chapter examines prevention strategies, including governance of advanced AI models, biological data, AI-enabled biological tools, DNA synthesis, and laboratory security. It also examines private-sector incentives mechanisms. It argues that effective prevention will require a combination of targeted regulation, market incentives, public-private collaboration, and international coordination.

Chapter 4 Detection

This chapter confronts the reality that prevention alone will not eliminate bio-convergence risks, and focuses on the need to develop national and international infrastructure to detect biological threats as early as possible. It examines the creation of a persistent national biological monitoring architecture capable of identifying both known and novel biological threats. It emphasizes the importance of wastewater surveillance, genomic monitoring (itself advanced using cutting-edge AI systems), distributed laboratory networks, and sustainable financing mechanisms to support continuous readiness.

Chapter 5 Attribution

This chapter makes the case that as AI-enabled biological design expands, the ability to rapidly determine whether an incident is natural, accidental, or deliberate, and to identify its origin, will be critically important to public health and national security. It outlines measures to improve biological forensics, bioengineering-signature detection, and the integration of public health, intelligence, and law-enforcement capabilities.

Chapter 6 Readiness and Response

This chapter addresses the dual needs for readiness—what is characterized in this report as a “warm start” to meet biological threats—and response—how public institutions act to protect human health when a threat presents itself. It examines the institutional, industrial, and scientific capacities required to prepare for and manage biological emergencies, proposing reforms to public-private surge capacity, crisis decision-making, medical countermeasure development, and clinical trials infrastructure.

Chapter 7 Resilience

This chapter assembles key elements of what America needs to build resilience to biological risks—meaning the sustained ability to reduce harmful effects and recover quickly from a significant biological incident, whether natural, accidental, or deliberate. It argues that enduring biosecurity will depend on sustained investment in biotechnology innovation, biological data infrastructure, and AI-enabled scientific capabilities. In particular, it tackles the intertwined challenges to US security and the private sector presented by China’s growing strength in select areas of biotechnology research and its increasing dominance in components of the pharmaceutical supply chain, including active pharmaceutical ingredients, key starting materials, and drug manufacturing. Meeting this challenge requires industrial policy to support US firms as well as cooperation with select global partners to onshore and “friend-shore” strategic supply chains.

Primary Recommendations

This report advances a comprehensive agenda for strengthening American biosecurity and competitiveness in the age of bioconvergence. While each of the report's recommendations is important on its own, this report encourages a holistic approach to managing AI-enabled biological risks, wherein the interconnected risks associated with new advances in AI models, biodesign tools, autonomous laboratories, and related convergent capabilities are addressed as an interactive ecosystem. The report's principal recommendations are organized around the core elements of the deterrence by resilience framework:

Prevention

- Establish AI-biosecurity guardrails for frontier models and laboratory systems. Create a federal regulatory framework for advanced AI models and autonomous laboratory capabilities that—independently or when combined—exceed defined compute and biological capability thresholds, with requirements tailored to biological design, synthesis, and dual-use research risks.
- Mandate oversight, testing, and incident reporting for high-risk AI systems. Require pre-deployment biosecurity evaluations, post-deployment misuse monitoring, incident reporting; empower a federal entity—such as the Center for AI Standards and Innovation (CAISI)—to set standards for risk assessment, red teaming, model evaluation, and safety testing of advanced AI systems.
- Align compliance incentives with regulatory obligations. Pair binding safety requirements with incentives such as tax credits, procurement preferences, liability safe harbors, and priority access to federal testbeds for organizations that implement independent audits, robust red teaming, and sequence-screening safeguards.
- Create a centralized governance framework for biological datasets. Establish a dedicated federal office to oversee biological data governance, expand sensitivity-based access controls for high-risk datasets, and require greater transparency from AI developers regarding biological training data to improve security while supporting innovation.
- Strengthen DNA synthesis screening and control access to biological tools. Mandate customer and sequence screening using function-based criteria, implement know-your-customer requirements for benchtop DNA synthesizers, and mandate technical safeguards that block sequences of concern.
- Modernize biosafety and biosecurity regulations for high-risk research. Apply risk-tiered biosafety requirements to all laboratories, including cloud labs; establish a unified statutory framework for high-risk biological research, and provide resources for implementing it; and require biosafety and biosecurity compliance as a condition of federal research funding.

Detection

- Establish a nationwide biological monitoring system. Set statutory requirements for biomonitoring coverage, detection timelines, and interoperable data standards, while scaling wastewater surveillance and pathogen-agnostic metagenomic sequencing.
- Strengthen federal biosecurity coordination and rapid capability development. Build a dedicated HHS biosecurity team with cross-government coordination responsibilities and provide flexible authorities (e.g., Other Transaction Authority) to test, evaluate, and deploy emerging monitoring technologies.
- Expand genomic monitoring through a distributed laboratory network. Authorize long-term federal partnerships with sequencing companies and laboratory networks to provide ubiquitous and continuous genomic surveillance coverage.
- Create a Biological Early Warning Trust Fund. Establish a dedicated, multiyear funding mechanism for biomonitoring and testing infrastructure, including automatic emergency funding triggers and matching grants for state and local governments to support surveillance, data sharing, and surge capacity.

Attribution

- Build a national biological attribution capability. Invest in the tools, infrastructure, and data-sharing systems needed to rapidly determine whether a biological event is natural, accidental, or deliberate, including the ability to detect novel engineering signatures and securely share across health, agriculture, intelligence, and law enforcement agencies.
- Operationalize advanced bio-attribution technologies. Authorize and fund the acceleration and transition of the Intelligence Advanced Research Projects Activity (IARPA) Finding Engineering-Linked Indicators (FELIX) program into a standing operational capability. This should include development of a national reference library of engineering-signatures covering laboratories and platforms known to work with dangerous pathogens; expansion of AI-based attribution tools capable of identifying biological agents designed entirely or substantially using AI systems; and integration of these tools with existing intelligence and law enforcement databases.

Readiness

- Designate the biotechnology ecosystem as a critical infrastructure sector. Establish the biotechnology ecosystem as a formal critical infrastructure sector; assign DHS, DOD, and HHS as co-Sector Risk Management Agencies; and create mechanisms for sustained public-private threat information sharing and resilience planning.
- Create a National Strategic Bio Readiness Reserve (NSBRR). Build a public-private surge capacity reserve modeled on the Civil Reserve Air Fleet to ensure rapid access to critical biotechnology capabilities during emergencies, including discovery and design, compute power, advanced development, and biomanufacturing.

Response

- Streamline and accelerate federal clinical trials. Establish a National Clinical Trials Emergency Authority; require major hospitals to maintain standing clinical trial infrastructure with pre-negotiated Indefinite Delivery, Indefinite Quantity (IDIQ)-style contracts; create an emergency NIH research reserve fund; and streamline contracting and trial activation during declared crises.
- Invest in capacity to rapidly field personal protective equipment (PPE), tests, therapeutics, and prototype vaccines and adaptable platform technologies, targeting prioritized viral families with significant epidemic or national security risk potential. Sustained funding will be needed for associated programs led by the Department of Defense (DOD), Biomedical Advanced Research and Development Authority (BARDA), and Coalition for Epidemic Preparedness Innovations (CEPI).

Resilience

- Create a public-private National Institute for Bioresilience (NIBR) to integrate AI, biological data, and national security missions. Establish and fund a public-private institution with dedicated compute and data infrastructure to develop frontier biological AI capabilities, support continuous threat characterization, build in biosecurity and biosafety measures, and accelerate countermeasure development.
- Establish Bio Emergency Services Teams (BEST), conduct regular preparedness exercises, and connect field response operations with national AI-enabled biological analysis and decision-support capabilities.
- Reinvest in US biotechnology innovation and workforce development. Restore and expand long-term federal investments in research, talent, and scientific infrastructure while strengthening incentives for domestic biotechnology R&D, pharmaceutical discovery, production, and commercialization.
- Launch a Senior Leaders Biosecurity Network to lead international cooperation on biotechnology security and resilience. Expand capacity-building partnerships and develop common standards for AI-enabled biotechnology, DNA synthesis screening, pathogen surveillance, and supply-chain resilience.
- Manage strategic competition for critical supply chains. Pursue a structured risk reduction dialogue with China that addresses growing dangers of proliferating AI-enabled biotechnologies, akin to dialogues on countering nuclear terrorism. Simultaneously, prioritize policies that improve the competitiveness, speed, and scalability of US-based biotechnology development—including accelerating clinical trials, supporting advanced manufacturing, and speeding regulatory review.
- Reduce pharmaceutical supply-chain vulnerabilities and expand domestic manufacturing. Assess critical dependencies on foreign biotechnology inputs, establish national security review mechanisms for sensitive biotechnology transactions, and incentivize domestic manufacturing capacity. Build support for a Boost American Bio Manufacturing Act modeled on the CHIPS and Science Act to reduce growing

dependencies on China via a comprehensive industrial policy supporting the US biotech and pharmaceutical industries.

For ease of access, the full set of the report's detailed recommendations is compiled in Appendix 1.

Introduction: The Deterrence by Resilience Framework

We are living in a transformational era defined by simultaneously accelerating advances in biology, artificial intelligence (AI), bioengineering, and digital infrastructure. This process, which this report refers to as *bioconvergence*, is revolutionizing how biological capabilities are developed, accessed, and applied. It promises massive advances in healthcare, science, and economic productivity. It is also changing the biological threat landscape in ways that defy existing security frameworks and create potentially existential risks to human health and US national security.

Bioconvergence has produced a democratized, offense-dominant threat environment. Capabilities that were once confined to highly resourced, state-run programs are increasingly distributed across borders and sectors. The digitization of biology, AI-enabled design, falling costs of synthesis and manufacturing, and the global diffusion of biotechnology tools and expertise are lowering barriers to entry while accelerating the pace of innovation. These forces offer the advantage to the offense: Biological systems self-replicate and scale rapidly, while defense against these threats is diffuse, resource-intensive, and reactive. Traditional approaches to deterrence—which are designed for domains with clear attribution, centralized control, and high barriers to entry—are increasingly irrelevant to emerging biological threats.

The emerging proliferation challenge of biological threats is qualitatively different from threats posed by nuclear weapons. It is persistent rather than episodic, distributed rather than centralized, and most importantly, the capabilities are deeply entwined with the same innovation ecosystems that underpin economic growth, national security, and public health. As a result, strategies that rely primarily on restriction, exclusion, or negotiated forbearance in exchange for protection (such as extended deterrence) are insufficient. A new framework that enables the establishment of deterrence by resilience is needed to meet this dynamic moment and build a more secure future.

Bioconvergence

Bioconvergence refers to the accelerating integration of advances in biology, artificial intelligence, bioengineering, and digital infrastructure that is transforming how biological capabilities are discovered, designed, scaled, and deployed.

About This Report

This report is the product of the Initiative on Bioconvergence, Biosecurity, and Bioresilience (IB3), housed within the Program on Emerging Technology, Scientific Advancement, and Global Policy at the Harvard Kennedy School of Government's Belfer Center for Science and International Affairs. The report benefits from extensive interviews, conversations, and convenings with individuals at

the forefront of the bioconvergence challenge, including current and former public officials; CEOs and technical advisors at leading AI, biotechnology, and pharmaceutical companies; pioneers in synthetic biology, machine learning, biosecurity, and public health; university-based scholars and think tank experts; and philanthropists funding relevant research.

The report's purpose is to advance a science-based, policy-relevant framework for managing biological risks in the era of AI-enabled biotechnologies. This framework is designed to enable the United States to sustain leadership in innovation while strengthening deterrence against misuse of novel biological capabilities. Because AI and biological innovations, including novel biological threats, are not limited to the United States, it will also be necessary to work with allies and partners, engage international institutions, and forge new dialogues, in particular with China, to manage emerging risks.

The central strategic challenge is this dual imperative: The United States must continue to lead in AI and biotechnology, both domains that are increasingly central to national power and societal well-being. At the same time, it is imperative to reduce the prospects that these capabilities will be used for pernicious purposes. As has been the case across time, innovation creates advantages, but it also exposes new vulnerabilities. Constraining innovation too rigidly risks forfeiting the societal and economic benefits promised by supercharged biotechnologies and could also dangerously restrain the development of the very capabilities needed to detect, deter, and mitigate emerging threats. For the private sector, securing their technologies from generating or exacerbating potentially catastrophic risks will differentiate their value and build public trust. Managing the tension between innovation and threat generation at this frontier will require deftness and sustained collaboration between government and the private sector. If done effectively, promoting innovation will advance the American edge in countering and withstanding AI-enabled biological risks.

The core analytic claim of this report is that, in the biological domain, **innovation infrastructure is security infrastructure**. A robust ecosystem of AI, biotechnology, data, manufacturing, and public-private coordination is the foundation not only of competitiveness, but of deterrence and resilience. Without maintaining a comprehensive technological edge, the United States will lack the leverage necessary to reduce adversarial confidence in the benefits of pernicious employment, shape global norms, raise the costs of misuse, and respond effectively when prevention fails.

Multiple reports, studies, and articles, many of which are referenced in this report, describe the growing risks associated with AI-enabled biotechnologies and propose a variety of ways to address them. This report seeks to expand upon this body of work, providing a technically informed assessment of risks, an analysis of the full scope and implications of the bioconvergence challenge, and a comprehensive and synergistic set of recommendations for action by the executive and legislative branches of the US government, the private sector, the international community, and philanthropy.

The information and analysis in this report are current as of July 2026. However, advancements in this domain are accelerating at a rate that far exceeds the speed of policymaking. New technologies and capabilities of AI and biotechnology emerge every day. For example, during the research and development of this report, METR updated its estimates of frontier AI agents' autonomous task-completion horizons. In its revised evaluation, the estimated post-2023 doubling interval for performance on a task suite focused primarily on software engineering, machine learning, and cybersecurity was approximately 4.3 months, compared with approximately seven months in its

earlier, longer-term analysis.¹ And in July 2026, University of Minnesota researchers reported a chemically defined synthetic cell capable of genome replication, growth, resource acquisition, and genetically encoded division.² Although urgent action is needed to match the scope and pace of such change, the recommendations in this report take into account this dynamic environment and have been developed to be adaptable over time.

The Strategic Framework: Bio-deterrence by Resilience

A new model of biodeterrence for a new challenge

This report makes the case that the United States needs to invest in a new model of bio-deterrence by resilience. Bio-deterrence by resilience is vitally important to the safety and security of the American people—and of the planet. This new paradigm focuses on both reducing the potential effects of biological misuse and improving recovery outcomes in order to systematically diminish adversarial confidence that pernicious actions will succeed, scale, or deliver strategic benefits. It calls for working to establish global guardrails for AI-enabled biotechnology and to be skeptical about their effectiveness. Therefore, it argues for persistent efforts to counter the prospect of generating malicious effects by limiting the impacts of attacks, raising the expected consequences of aggression, and investing in sustained national strengths at the intersection of AI, biotechnology, and bioengineering. As systems become more resilient—rendering them harder to exploit and quicker to rebound—the deliberate use of biological attacks becomes more costly and less attractive.

If implemented at scale, the recommendations in this report will systematically reduce adversary confidence that misuse of novel biological tools will succeed, scale, or deliver strategic benefits. Together, they combine upstream prevention, rapid detection, effective response, and sustained recovery to raise the expected costs and risks of biological misuse.

Bio-deterrence by resilience is a strategy for reducing biological risks in an environment characterized by low barriers to entry, uncertain attribution, and rapid technological change.

Traditional deterrence approaches emphasize deterrence by denial—preventing an adversary from obtaining capabilities or achieving its objectives in using such capabilities—and deterrence by punishment—threatening to impose pariah status or to use massive retaliation to dissuade anyone from using such capabilities. In the biological domain, this is insufficient. Attribution of biological attacks may be delayed or contested; attacks may be covert or ambiguous; and perpetrators may be nonstate actors or risk-tolerant states. Under these conditions, deterrence weakens when adversaries believe they can act without being identified, stopped, or denied meaningful gains.

¹ Model Evaluation & Threat Research [METR], “Time Horizon 1.1,” Jan. 29, 2026, <https://metr.org/blog/2026-1-29-time-horizon-1-1/>.

² Nathaniel J. Gaut et al., “A Chemically Defined Synthetic Cell Capable of Growth and Replication,” *bioRxiv*, July 2, 2026, <https://doi.org/10.64898/2026.07.01.735724>.

Deterrence by resilience shifts the focus from retaliation and denial to reducing the potential effects of biological misuse and improving recovery outcomes in order to systematically diminish adversarial confidence that pernicious actions will succeed, scale, or deliver strategic benefits. As systems become more resilient—rendering them harder to exploit and quicker to rebound—the deliberate use of biological attacks becomes more costly and less attractive.

In practice, deterrence by resilience is not a single capability but a self-reinforcing virtuous cycle linking prevention, response, and recovery. As systems become more resilient—in essence, harder to exploit and quicker to rebound—the deliberate use of biological attacks becomes less attractive, even in the absence of clear attribution or punitive threat to retaliate. Over time, this dynamic can shape global norms, new regimes, investment decisions, and adversary behavior, strengthening deterrence through persistent advantage rather than relying exclusively on prevention, which is unrealistic, or on the threat of nuclear retaliation, which may or may not be credible.

Mapping the Bioconvergence Threat

A. The Promise and Peril of Bioconvergence

The era of bioconvergence is defined by the accelerating integration of biotechnologies, artificial intelligence, and bioengineering. This process is transforming how biological capabilities are discovered, designed, and applied. It promises major advances in health, science, and economic productivity. It is also reshaping the biological threat landscape, challenging our understanding of both positive and pernicious potential.

Bioconvergence

The accelerating integration of advances in biology, artificial intelligence, bioengineering, and digital infrastructure that is transforming how biological capabilities are discovered, designed, scaled, and deployed. Its drivers are mutually reinforcing: DNA sequencing and synthesis have collapsed the cost of reading and writing genetic information; gene editing tools like CRISPR have dramatically broadened access to precise genome modification; AI systems trained on biological data can now design novel proteins and predict molecular interactions at scales previously unimaginable; and automation is closing the loop between digital design and physical synthesis. The same convergence that enables historic advances in medicine and agriculture is lowering barriers to the creation and misuse of dangerous biological agents.

The promise is already visible. mRNA vaccine platforms enabled the fastest vaccine development in history during the COVID-19 pandemic.³ Gene therapies are now approved for previously untreatable diseases. AlphaFold has enabled predictions of the structures of virtually all known proteins. In 2024, Demis Hassabis and John Jumper received half of the Nobel Prize in Chemistry for protein-structure prediction; David Baker received the other half for computational protein design.⁴ Agentic AI systems capable of autonomously generating hypotheses, coordinating

3 Sandy Cohen, "The Fastest Vaccine in History," UCLA Health, 2020, <https://www.uclahealth.org/news/article/the-fastest-vaccine-in-history>; Elie Dolgin, "How COVID Unlocked the Power of RNA Vaccines," *Nature*, 2021, <https://www.nature.com/articles/d41586-021-00019-w>.

4 Google DeepMind, "AlphaFold," accessed April 2026, <https://deepmind.google/science/alphafold/>.

experiments, and analyzing data are beginning to compress research timelines that historically ran for decades.⁵

But the risks are also accelerating alongside the benefits. The Department of Homeland Security has warned that the intersection of AI and biotechnology could “create novel risks to the homeland” by enabling malicious actors to better “conceptualize and conduct” biological weapons attacks.⁶ A group of thirty-five biosecurity professionals has called for “urgent attention” to risks arising from bioconvergence.⁷ Anthropic CEO Dario Amodei has written that biology is “the area I’m most worried about” in the context of AI-enabled catastrophic risk.⁸

These warnings reflect a novel risk environment that requires a new analytical framework. This chapter explores that framework through three broad steps: establishing the baseline of biological risk that already exists independent of AI and biotechnology; mapping how technological advances layer new and qualitatively different threats on top of that baseline; and presenting three analytically grounded threat scenarios that illustrate how these dynamics could manifest in practice.

Baseline Biological Risk

Biological risk did not begin with advanced technology. Before the first large language model was trained or the first CRISPR edit was made, the world already faced persistent, serious threats from naturally occurring diseases, laboratory accidents, and deliberate misuse of biological agents. Evaluating the additional risks that AI and biotechnology introduce requires first understanding the magnitude of what already exists, and resisting the error of treating bioconvergence as a threat that begins from a zero-baseline risk.

Baseline Biological Risk

The persistent, background level of biological threat that exists independent of AI and advanced biotechnology. It encompasses naturally occurring disease arising from ordinary ecological dynamics, laboratory accidents and containment failures, and deliberate biological misuse that predates the current era of AI-enabled design. The baseline is not static: Structural changes in human activity, global connectivity, and laboratory proliferation are increasing it. Understanding the baseline is essential for calibrating the additional risks that bioconvergence introduces, and for allocating biosecurity resources proportionate to the full spectrum of threats.

-
- 5 Sam Rodriques and Michael Hinks, “Kosmos: An AI Scientist for Autonomous Discovery,” *Edison*, Nov. 5, 2025, <https://edisonscientific.com/articles/announcing-kosmos>.
 - 6 US Department of Homeland Security, “Fact Sheet: DHS Advances Efforts to Reduce the Risks at the Intersection of Artificial Intelligence and Chemical, Biological, Radiological, and Nuclear (CBRN) Threats,” 2024, https://www.dhs.gov/sites/default/files/2024-04/24_0429_cwmd-dhs-fact-sheet-ai-cbrn.pdf.
 - 7 Nuclear Threat Initiative, *The Convergence of Artificial Intelligence and the Life Sciences: Safeguarding Technology, Rethinking Governance, and Preventing Catastrophe* (Nuclear Threat Initiative, 2023), https://www.nti.org/wp-content/uploads/2023/10/NTIBIO_AI_Executive-Summary_FINAL.pdf.
 - 8 Dario Amodei, “The Adolescence of Technology: Confronting and Overcoming the Risks of Powerful AI,” January 2026, <https://www.darioamodei.com/essay/the-adolescence-of-technology>.

The baseline operates across three distinct pathways: naturally occurring disease, laboratory accidents, and deliberate misuse. All three are changing in ways that make the current moment more dangerous than prior eras, independent of the specific threat posed by bioconvergence.

Natural Threats

The COVID-19 pandemic provides the starkest recent demonstration of what biological threats can do at scale. The pandemic caused at least 7 million reported deaths worldwide,⁹ while the IMF estimated cumulative global output losses of approximately \$13.8 trillion through 2024.¹⁰ It also exposed the fragility of international public health infrastructure and the inadequacy of public health systems built for case documentation and reactive response rather than early warning. COVID-19 was not an anomaly; it was a preview of what zoonotic disease events can cause under the structural conditions that characterize the modern world.

Zoonotic Disease

A disease caused by a pathogen that originates in an animal reservoir and successfully crosses the species barrier to infect humans. This transmission event is called spillover. The majority of emerging infectious diseases affecting humans are zoonotic in origin: COVID-19, SARS, MERS, Ebola, HIV, and most influenza strains all emerged through animal-to-human spillover. The frequency and consequences of spillover events are shaped by the degree and nature of human contact with animal reservoir populations, making land use, agricultural practices, and habitat encroachment significant drivers of zoonotic disease risk.

The structural conditions driving zoonotic spillover are worsening.¹¹ An estimated 1.7 million undiscovered viruses circulate in mammalian and avian reservoir populations; the ecological conditions that determine when and how frequently these cross into human hosts are being altered by human activity at accelerating rates.¹² Agricultural intensification and the expansion of human settlements into previously wild areas increase the frequency of contact between humans and wildlife reservoirs, compressing the ecological buffer that historically slowed spillover.¹³ Deforestation fragments habitats, concentrating reservoir animals and intensifying transmission dynamics within those populations before spillover occurs. Climate change is expanding the geographic range of vector-borne diseases, including enabling the spread of mosquito species carrying dengue viruses, malaria parasites, and Zika virus into previously temperate zones, and altering the seasonal patterns that govern vector reproduction.¹⁴ Urbanization increases population

9 World Health Organization, "WHO COVID-19 Dashboard," accessed August 2, 2026, <https://data.who.int/dashboards/covid19/deaths>.

10 International Monetary Fund, "New IMF Staff Paper Strategy to Manage the Long-Term Risks of COVID-19," Apr. 5, 2022, <https://doi.org/10.5089/9798400205996.001>.

11 Anthony S. Fauci and David M. Morens, "Emerging Pandemic Diseases: How We Got to COVID-19," *Cell* 182 (2020): 1077-92, <https://doi.org/10.1016/j.cell.2020.08.021>.

12 Dennis Carroll et al., "Building a Global Atlas of Zoonotic Viruses," *Bulletin of the World Health Organization* 96 (2018): 292-94, <https://pmc.ncbi.nlm.nih.gov/articles/PMC5872013/>.

13 Bryony A. Jones et al., "Zoonosis Emergence Linked to Agricultural Intensification and Environmental Change," *Proceedings of the National Academy of Sciences* 110 (2013): 8399-8404; Kate E. Jones et al., "Global Trends in Emerging Infectious Diseases," *Nature* 451 (2008): 990-93.

14 Intergovernmental Panel on Climate Change, *Climate Change 2023: Synthesis Report* (Intergovernmental Panel on Climate Change, 2023); Fauci and Morens, "Emerging Pandemic Diseases," 1077-92.

density and therefore transmission rate once a pathogen enters human populations; globalization ensures that local spillover events can seed outbreaks on multiple continents within days.

Natural threats require no adversarial actor, no advanced technology, and no deliberate intent. They constitute the baseline of biological risk, and that baseline is rising.

Biological Accidents

Accidents also contribute to baseline risk. Laboratory accidents, unintentional releases, and the unintended consequences of well-meaning research have repeatedly had public health consequences. In 1979, the accidental release of anthrax spores from a Soviet military research facility at Sverdlovsk killed at least sixty-four people in the surrounding city.¹⁵ In 2004, viral exposures were documented for two researchers at the Institute of Viral Disease Control at the Chinese Center for Disease Control in Beijing, triggering a SARS outbreak that required emergency public health intervention.¹⁶ In 2014, laboratory safety incidents at the US Centers for Disease Control and Prevention exposed dozens of employees to viable anthrax and resulted in the unintended transfer of an H5N1-contaminated sample to a USDA laboratory, prompting a congressional investigation and substantial internal reforms.¹⁷ The debate persists about whether SARS-CoV-2 emerged through zoonotic spillover or a research-related laboratory incident.¹⁸

The risk of accidents is growing. The number of BSL-3 and BSL-4 laboratories has expanded significantly over the past two decades, with new high-containment facilities established in more countries and institutional contexts, often with variable safety cultures and oversight regimes.¹⁹

Biological Misuse

The deliberate use of biological agents as weapons predates modern biotechnology by centuries. In 1763, British officers conspired to distribute smallpox-contaminated blankets to Indigenous peoples during Pontiac's Rebellion.²⁰ During the First World War, Germany ran a covert biological sabotage program, infecting Allied horses and mules with anthrax and glanders to disrupt supply lines.²¹ Imperial Japan's Unit 731 conducted human experimentation during the Second World War through deliberately infecting prisoners with plague, cholera, and anthrax.²²

More recent cases demonstrated ongoing interest in weaponizing biology. The Soviet Union's Biopreparat program, which weaponized smallpox on an industrial scale, only came to light in the

15 Matthew Meselson et al., "The Sverdlovsk Anthrax Outbreak of 1979," *Science* 266 (1994): 1202–8, <https://www.science.org/doi/10.1126/science.7973702>.

16 Dennis Normile, "Mounting Lab Accidents Raise SARS Fears," *Science* 304 (2004): 659–61, <https://www.science.org/doi/10.1126/science.304.5671.659>.

17 US Government Accountability Office [GAO], *High-Containment Laboratories: Comprehensive and Up-to-Date Policies and Stronger Oversight Mechanisms Needed to Improve Safety*, GAO-16-305 (March 2016), <https://www.gao.gov/assets/gao-16-305.pdf>.

18 World Health Organization, Scientific Advisory Group for the Origins of Novel Pathogens (SAGO), "Report on the Origins of COVID-19," June 27, 2025, <https://www.who.int/news/item/27-06-2025-who-scientific-advisory-group-issues-report-on-origins-of-covid-19>.

19 Jocelyn Kaiser, "Growing Number of High-Security Pathogen Labs Around the World Raises Concerns," *Science*, Mar. 17, 2023, <https://www.science.org/content/article/growing-number-high-security-pathogen-labs-around-world-raises-concerns>.

20 Colin G. Calloway, *The Scratch of a Pen: 1763 and the Transformation of North America* (Oxford University Press, 2006).

21 Mark Wheelis, "First Shots Fired in Biological Warfare," *Nature* 395, no. 6699 (1998): 213, <https://doi.org/10.1038/26089>.

22 Hal Gold, *Unit 731 Testimony: Japan's Wartime Human Experimentation Program* (Tuttle Publishing, 1996).

1990s.²³ The Aum Shinrikyo cult in Japan, which carried out the 1995 Tokyo subway sarin attack, made multiple attempts to develop and deploy biological agents, including anthrax and botulinum toxin.²⁴ The 2001 Amerithrax attacks killed five people, infected seventeen others, and caused billions of dollars in remediation costs.²⁵

These cases reveal a consistent pattern: Motivated actors have repeatedly sought biological weapons capability, and some have succeeded. Defenders have historically benefited from a negative correlation between capability and motivation. States with the capacity to weaponize biology have largely chosen not to because of the opprobrium that use would engender, blowback risk, or both. Nevertheless, although the United States forswore biological weapons development in 1969,²⁶ other countries have continued to pursue their own programs.²⁷ AI and biotechnology are now further shifting this equilibrium by radically democratizing access to potentially lethal capabilities, including those who are less bound by the same inhibitions or motivations that have historically constrained the most capable nations.

B. Bioconvergence Risk

Bioconvergence risk builds on the baseline described above. It emerges from rapid, mutually reinforcing advances in biotechnology, artificial intelligence, and bioengineering. These quickly evolving fields are converging at a pace that is literally hard to fathom: In one recent evaluation, a leading US AI model outperformed every human expert tested on a benchmark of practical virology knowledge.²⁸

Risks from Biotechnology

Three key developments have driven the modern biotechnology revolution.

First, large-scale genomic sequencing has made it possible to read the genetic code of organisms at high speed and low cost. The first human genome cost at least \$500 million to sequence.²⁹ Whole genome sequencing can now be done for as little as \$200.³⁰

23 Cesar Bartolome, “Biopreparat: The Soviet Biological Weapons Programme,” University of Kent, 2022, <https://research.kent.ac.uk/kcwa/biopreparat-the-soviet-biological-weapons-programme/>.

24 Philipp C. Bleek, “Revisiting Aum Shinrikyo: New Insights into the Most Extensive Non-State Biological Weapons Program to Date,” Nuclear Threat Initiative, 2011, <https://www.nti.org/analysis/articles/revisiting-aum-shinrikyo-new-insights-most-extensive-non-state-biological-weapons-program-date-1/>.

25 Federal Bureau of Investigation, “Amerithrax or Anthrax Investigation,” accessed April 2026, <https://www.fbi.gov/history/famous-cases/amerithrax-or-anthrax-investigation>.

26 Richard Nixon, “Statement on Chemical and Biological Defense Policies and Programs,” Nov. 25, 1969, The American Presidency Project, <https://www.presidency.ucsb.edu/documents/statement-chemical-and-biological-defense-policies-and-programs>.

27 Arms Control Association, “Biological Weapons Convention: Overview and Key Provisions,” 2023, <https://www.armscontrol.org/factsheets/bwc>; Office of the Director of National Intelligence, *Annual Threat Assessment of the US Intelligence Community* (February 2025), <https://www.odni.gov/files/ODNI/documents/assessments/ATA-2025-Unclassified-Report.pdf>.

28 For example, in April 2026, SecureBio reported that a pre-release version of GPT-5.5 achieved the highest score of any model it had tested on its Virology Capabilities Test, which assesses practical virology knowledge, and outperformed all human experts who took the evaluation. SecureBio, “SecureBio’s Pre-Release Assessment,” Substack, <https://securebio.substack.com/p/securebios-pre-release-assessment>, Apr.23, 2026.

29 National Human Genome Research Institute, “The Cost of Sequencing a Human Genome,” accessed April 2026, <https://www.genome.gov/about-genomics/fact-sheets/Sequencing-Human-Genome-cost>.

30 Illumina, “Illumina’s Revolutionary NovaSeq X Exceeds 200th Order Milestone in First Quarter 2023,” 2023, <https://emea.illumina.com/company/news-center/press-releases/2023/4b48580b-42d4-419b-8512-5adccb069836.html>.

Second, DNA synthesis has become cheaper, faster, and more accurate. Synthesis is the ability to write, rather than merely read, genetic sequences: to specify a desired DNA sequence and have it produced to order. The cost per base pair has fallen from approximately \$4 in the early 2000s to \$0.10 today.³¹

Third, the CRISPR-Cas9 gene editing system, developed in the early 2010s, provided researchers with a simple, precise, and relatively low-cost tool for modifying the genomes of living organisms. CRISPR has dramatically expanded the range of experiments accessible to researchers without specialized infrastructure, placing precise genome modification within reach of a far broader scientific community. Beyond CRISPR, advances in bioengineering have accelerated synthetic biology; for example, researchers from the University of Minnesota announced in June 2026 that they had for the first time created an entirely synthetic cell, which was capable of feeding, growing, copying its DNA, and dividing.³²

Each advance unlocked new opportunities and compounded the baseline biological risks described above. Together, these capabilities enabled what researchers now call generative biology: the ability to design and construct biological systems that do not exist in nature.

Generative Biology

The convergence of AI, automation, and synthetic biology that makes it possible to design and construct biological systems that do not exist in nature. Rather than working with what evolution produced, researchers can now specify a desired biological function and use AI-powered tools and other techniques to generate the sequences, proteins, or organisms engineered to achieve it. The same capability can accelerate drug discovery while also lowering barriers to engineering agents with specified harmful properties.

Access to Biotechnology

The implications of generative biology depend on who can access it. As noted above, the modern biotechnology revolution has been defined by democratization as well as innovation at the frontier, and the tools necessary for generative biology are reaching a far wider range of actors than the institutional research settings that dominated the landscape of the past. Oversight frameworks built for centralized laboratory science have not kept pace with this expanding distribution.

Benchtop DNA synthesizers provide an example of the governance gap. Traditional DNA synthesis requires specialized equipment housed in commercial facilities, which are subject, at least in principle,³³ to screening of orders for dangerous sequences. A new generation of compact, desktop-scale devices now enables on-site synthesis of DNA without reliance on commercial providers and without any regulatory oversight of the sequences produced.³⁴ The result is a capability that

31 Prashant Singh et al., “The Falling Cost of DNA Synthesis,” *Nature Biotechnology* (2025).

32 Whipple, “Synthetic Cell Offers First Glimpse of Life That Is ‘Built Not Born.’”

33 As discussed further in Chapter 3, most prominent companies comply with voluntary screening standards, such as those established through the International Gene Synthesis Consortium.

34 Nuclear Threat Initiative, “Benchtop DNA Synthesis Devices: Capabilities, Biosecurity Implications, and Governance,” May 10, 2023, <https://www.nti.org/analysis/articles/benchtop-dna-synthesis-devices-capabilities-biosecurity-implications-and-governance/>; Lena

was previously subject to institutional and commercial chokepoints now operating outside them entirely.

Contract research organizations (CROs) and emerging cloud labs represent a parallel shift in access. CROs conduct research on behalf of clients, with the client providing instructions and the CRO providing expertise, equipment, and labor. Cloud labs promise to extend this model further: Clients design experiments digitally, and automated laboratory systems execute them with minimal human involvement. Together, these services lower expertise and equipment barriers to frontier biotechnology, allow actors to outsource the most skill-intensive steps in a research workflow, and introduce oversight challenges for which no current regulatory framework is adequate.

Risks from Artificial Intelligence

Democratized access to biotechnology is part of the bioconvergence risk equation. Advancing AI capabilities are another, and the two interact in ways that are qualitatively different from either alone. Continued improvements in AI, driven by scaling compute, algorithmic innovation, and access to high-quality biological data, are further lowering barriers to entry and raising the ceiling of harm that motivated actors can cause. There are at least five categories of AI systems with meaningful dual-use biological applications:

- **Large Language Models (LLMs):** A subset of foundation models trained on internet data that can output natural language or other multimodal content and be adapted for a variety of tasks. LLMs can respond to natural language queries across biology, chemistry, and medicine; synthesize complex technical literature; and explain and troubleshoot experimental protocols.
- **Biological Foundation Models (BFMs):** A subset of foundation models trained on biological sequence data that can be adapted for a variety of research tasks. Evo 2 is an example of a BFM, trained on 9 trillion nucleotides and capable of prediction and design tasks across DNA, RNA, and proteins.³⁵ ESM3 is another example, trained on data from 2.78 billion proteins.³⁶
- **Biological Design Tools (BDTs):** Models trained on biological sequence data typically optimized for a specific biological task or adaptable to a narrow range of applications. AlphaFold 3 is an example of a BDT, trained on experimentally derived protein structures and sequences.³⁷ RoseTTA-Fold-All-Atom is a freely available equivalent that can model entire biological assemblies.³⁸
- **AI Science Agents:** AI systems capable of reasoning, planning, and taking actions to accelerate the scientific method. Co-Scientist is an example of an AI science agent,

Kroepke, "Regulatory Gaps in Benchtop Nucleic Acid Synthesis Create Biosecurity Vulnerabilities," Arms Control Association, 2025, <https://www.armscontrol.org/blog/2025-11-24/regulatory-gaps-benchtop-nucleic-acid-synthesis-create-biosecurity>.

35 Garyk Brixi et al., "Genome Modeling and Design Across All Domains of Life with Evo 2," bioRxiv, Feb. 21, 2025, <https://doi.org/10.1101/2025.02.18.638918>.<https://doi.org/10.1101/2025.02.18.638918>

36 EvolutionaryScale, "ESM3: Simulating 500 Million Years of Evolution with a Language Model," June 25, 2024, <https://www.evolutionaryscale.ai/blog/esm3-release>.

37 Josh Abramson et al., "Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3," *Nature*, May 8, 2024, <https://www.nature.com/articles/s41586-024-07487-w>.

38 Rohith Krishna et al., "Generalized Biomolecular Modeling and Design with RoseTTAFold All-Atom," *Science*, Mar. 7, 2024, <https://www.science.org/doi/10.1126/science.adl2528>.

a multiagent system that autonomously proposes, ranks, and iterates on hypotheses for a given scientific problem.³⁹ Kosmos from Edison Scientific is another example, conducting literature reviews and running scientific workflows autonomously.⁴⁰

- Autonomous Laboratories. Autonomous laboratories are “self-driving laboratories” that integrate AI with automated capabilities, such as robotics, to conduct scientific experiments without human intervention.⁴¹

C. Trends in Biological AI Systems

Four trends compound the biosecurity risks that biological AI systems pose.

First, **capabilities are advancing rapidly**. For example, while this report was being developed, Model Evaluation & Threat Research, a nonprofit that evaluates advanced AI systems and studies the risks that they may pose, found that leading AI agents were improving faster than previously estimated. In 2025, they estimated that the duration of tasks AI systems could autonomously complete without human oversight would double approximately every seven months.⁴² By 2026 that estimate had shrunk to 4.3 months.⁴³ The progression from GPT-4.5, classified as “medium” biological capability by OpenAI, to GPT-5, classified as “high,” occurred within six months. Benchmarks evaluating general scientific reasoning are increasingly saturated, meaning the evaluations designed to measure frontier capability are being outpaced by capability itself.

Second, **significant dual-use capabilities are diffusing** across a widening group of actors. The academic norms of transparency and reproducibility create substantial pressure to publish model weights and code alongside research, resulting in many BFM and BDTs being released as open source. Highly capable systems, including Evo 2 and RoseTTAFold-All-Atom, have been released under permissive licenses enabling download, modification, and redistribution without oversight.⁴⁴ Once released, their distribution is effectively irreversible. An attempt by Google DeepMind to release AlphaFold 3 under restricted access prompted a coordinated open letter from scientific leaders criticizing the decision; within six months, Google DeepMind released the full source code and model weights.⁴⁵ This diffusion is not limited to US-origin tools or US-based actors. China has made substantial investments in AI-enabled biology. BioMap’s xTrimoPGLM, a 100 billion-parameter protein language model, has demonstrated strong performance across structure, function, and interaction prediction tasks, with reports of competitive accuracy to leading US

39 Google DeepMind Co-Scientist Team, “Co-Scientist: A Multi-Agent AI Partner to Accelerate Research,” Google DeepMind, May 19, 2026, <https://deepmind.google/blog/co-scientist-a-multi-agent-ai-partner-to-accelerate-research/>.

40 Edison Scientific, website, <https://edisonscientific.com/>.

41 National Institute of Standards and Technology [NIST], “Autonomous Laboratories,” accessed July 5, 2026, <https://www.nist.gov/autonomous-laboratories/>.

42 Thomas Kwa et al., “Measuring AI Ability to Complete Long Tasks,” METR, Mar. 19, 2025, <https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/>.

43 METR, *Time Horizon 1.1*, January 2026.

44 Arc Institute, Evo 2, GitHub repository, accessed July 5, 2026, <https://github.com/ArcInstitute/evo2>; Baker Laboratory, RoseTTAFold-All-Atom, GitHub repository, accessed July 5, 2026, <https://github.com/baker-laboratory/RoseTTAFold-All-Atom>.

45 Stephanie Wankowicz et al., “AlphaFold 3 Transparency and Reproducibility,” Zenodo, 2024, <https://zenodo.org/records/11391920>; Ewen Callaway, “AI Protein-Prediction Tool AlphaFold3 Is Now More Open,” *Nature*, Nov. 14, 2024, <https://www.nature.com/articles/d41586-024-03708-4>.

systems on antibody-target binding.⁴⁶ ByteDance has launched Protenix, a competing molecular structure prediction model.⁴⁷ The Chinese National Genebank aggregates genomic data providing infrastructure for government-directed biological AI development that has no US counterpart.⁴⁸

Third, **safeguards are not keeping pace** with the capabilities and diffusion. Safeguards are used in an estimated three percent of BFMs and BDTs.⁴⁹ Model-level safeguards, including data curation, refusal training, and targeted unlearning, face fundamental limitations: Models can often infer harmful information from adjacent training data even when specific sequences are excluded; refusal training can be reversed through adversarial fine-tuning; and open-source systems can be modified to remove safeguards entirely. System-level safeguards, including classifiers and user verification, are similarly circumventable through sophisticated jailbreaks and false identities.

Fourth, the convergence of autonomous laboratories with LLMs, BDTs, and AI science agents will enable the physical **creation of novel pathogens without a human in the loop**. This could markedly accelerate deliberate and accidental biological risks.

D. Evaluating Risks from Biological AI Systems

The capacity to effectively evaluate the risks of AI-enabled biological capabilities is being developed, but it is a lagging rather than a leading indicator of what is happening. For example, efforts to evaluate how much “uplift” current AI systems already provide to biological threat actors have produced mixed results, although since 2024 it has become increasingly clear that the risks are real and growing.

Uplift

The meaningful increase in a user’s ability to accomplish a task that results from access to an AI system: the difference between what someone could achieve with the system’s assistance compared to without it, relying only on existing resources such as internet searches, textbooks, or prior knowledge and expertise. In biosecurity contexts, uplift is a central measure in AI capability assessments. The relevant question is not whether a model can engage with sensitive biological information, but whether it materially advances a user’s ability to acquire knowledge or complete tasks that would otherwise be beyond their reach. Uplift is defined relative to an individual’s capability, i.e., equipping a highly trained researcher with additional cutting-edge techniques and turning a less skilled individual into a fluent practitioner are both examples of uplift.

46 Bo Chen et al., “xTrimoPGLM: Unified 100-Billion-Parameter Pretrained Transformer for Deciphering the Language of Protein,” *Nature Methods* 22 (2025), <https://doi.org/10.1038/s41592-025-02636-z>.

47 Protenix Team et al., “Protenix-v1: Toward High-Accuracy Open-Source Biomolecular Structure Prediction,” *bioRxiv*, Feb. 6, 2026, <https://doi.org/10.64898/2026.02.05.703733>.

48 Weiwen Wang et al., “The China National GeneBank Sequence Archive (CNSA) 2024 Update,” *Horticulture Research* 12, no. 5 (2025), <https://doi.org/10.1093/hr/uhaf036>.

49 Pablo Villalobos and David Atanasov, “Announcing Our Expanded Biology AI Coverage,” *Epoch AI*, Jan. 29, 2025, <https://epoch.ai/blog/announcing-expanded-biology-ai-coverage>.

For example, an initial benchmark for AI uplift on biological tasks was set by a 2024 RAND red-team study in which researchers produced operational plans for specific biological attack scenarios with and without LLM assistance. RAND found “no statistically significant difference in the viability of plans generated with or without the assistance of the current generation of large language models.”⁵⁰ This finding was significant but is now dated: LLMs are far more capable than they were in late 2023, when the study was conducted.

LLM’s capacity to support uplift appears to be advancing rapidly. A 2025 wet lab study by Lawrence Livermore National Laboratory found results “consistent with skills-based uplift” but had a sample size of only ten.⁵¹ A 2026 controlled evaluation comparing novices with LLM access against those who only had access to the internet (but not LLMs) across in-silico tasks (computational questions rather than wet lab experimentation) found that LLM-assisted participants were 4.16 times more accurate than the internet-only control group. In the study, accuracy was measured by scoring performance on a number of benchmark tasks related to virology and molecular biology—in formats including multiple-choice questions, short-answer questions, and coding prompts.⁵²

Taken together, these studies outline the current risk landscape with important specificity. AI systems can lower information barriers substantially, particularly for computational and in-silico tasks. Tacit knowledge, process knowledge, and hands-on laboratory skills remain significant barriers for wet-lab work. The gap between these findings matters: It defines the boundary between what sophisticated nonstate actors can achieve today and what they may be able to achieve as laboratory automation continues to advance.

Spotlight: Bio-Uplift in Claude Mythos 5 and Claude Fable

As the frontier advances, this gap will be further reduced. In its June 2026 system card—a technical manual explaining how an AI system was built, how it was tested, how it works, and how it should be safely used—for Claude Mythos 5 and Claude Fable 5, Anthropic classified the models as having “CB-1” capabilities, or the ability to “significantly help individuals or groups with basic technical backgrounds (e.g., undergraduate STEM degrees) create/obtain and deploy chemical or biological weapons with serious potential for catastrophic damages.”⁵³ While the company judged the model to fall short of the higher “CB-2” threshold covering novel weapons production, it described that determination as “a much less clear and obvious judgment than with previous models,” concluding that the model “can likely accelerate well-resourced expert teams at novel bioweapon development, and materially increase their chances of success.”⁵⁴ This marks a notable shift from Anthropic’s earlier Mythos Preview system card, which, just two months prior in April 2026, characterized the ability of the model to approach the CB-2 classification as a clearer negative.

50 Protenix Team et al., “Protenix-v1: Toward High-Accuracy Open-Source Biomolecular Structure Prediction,” bioRxiv, Feb. 6, 2026, <https://doi.org/10.64898/2026.02.05.703733>.

51 Luca Righetti, “Research Note: Five Lessons from Having Helped Run an AI-Biology RCT,” METR, Feb. 19, 2026, <https://metr.org/blog/2026-02-19-five-lessons-from-ai-biology-rct/>.

52 Chen Bo Calvin Zhang et al., “LLM Novice Uplift on Dual-Use, In Silico Biology Tasks,” arXiv, Feb. 26, 2026, <https://arxiv.org/abs/2602.23329>.

53 Anthropic, *Claude Fable 5 / Mythos 5 System Card* (June 9, 2026), 18, <https://www.anthropic.com/claude-fable-5-mythos-5-system-card>.

54 Ibid. at 18.

The real-world implications are already evident in testing. The Claude Fable 5 and Mythos 5 System Card describes a beneficial-use tabletop exercise in which two-person teams of generalist biologists using the models produced protocol designs that expert graders assessed as comparable to or better than those produced by specialist teams.⁵⁵ In one evaluated task, the generalist team completed roughly sixteen hours of model-assisted work that graders estimated could otherwise have required forty to ninety-five working days without AI assistance.⁵⁶ Anthropic assessed the residual catastrophic risk from novel weapons production as “low, but higher than for any previous model, and with significant uncertainty.”⁵⁷

This development in frontier capability reinforces this chapter’s central concerns. The failure modes the evaluators documented in models—hallucinated citations, weak constraint carryover across long sessions, and poor recovery from errors⁵⁸—are closely related to the critical remaining constraint of the tacit knowledge barrier. However, Anthropic itself notes several of these limitations “can likely be alleviated through better harnessing and prompting,” suggesting the safeguards may not constitute durable protection.

Given the consequences of biological misuse at scale, serious investment is required across all layers of the deterrence by resilience framework. The three threat scenarios below ground that imperative, making visible which capability and governance gaps enable harms in each threat pathway.

Threat Scenarios

Scenario analyses help illustrate the threats emerging from bioconvergence’s rapidly evolving risk landscape and identify the interventions most likely to reduce those threats. Scenarios are not prediction; they are a method for identifying which capabilities, actor types, and governance gaps determine outcomes across a range of plausible futures.

To develop a credible and analytically grounded set of scenarios, IB3 ran a structured analysis process in January 2026. Senior scientists, national security officials, and biosecurity specialists from across government, academia, and industry submitted reasonable worst-case scenarios at the intersection of AI and biotechnology. More than forty distinct submissions were clustered and de-duplicated into a shortlist spanning a range of actors, capabilities, and outcomes. Shortlisted scenarios were then assessed for probability by 2030 and impact if realized.

This framework reveals a threat space more varied in both mechanisms and consequences than the mass-casualty pandemic paradigm that tends to dominate public and policy discussion of biological risk. Analysis of each scenario highlights important threat vectors and remaining barriers, both of which inform the prevention, detection, response, and recovery architecture described in Chapters 4 through 7. It is important to remember that these scenarios

55 Ibid. at 23.

56 Ibid.

57 Ibid. at 18.

58 Ibid. at 23.

are what experts can imagine today; continued rapid progress across AI, biotechnology and engineering may bring further new concerning futures into view.

Relevant Actors

The expert submissions surfaced four broad categories of threat actor, distinguished by capability, motivation, and the nature of the uplift they would receive from advances in biotechnology and AI.

The first is **sophisticated state and state-adjacent actors**: nation-state programs with purpose-built research facilities, dedicated scientific personnel, and the institutional continuity to pursue ambitious biological design objectives over time. For instance, Russia and China maintain infrastructure and expertise broadly relevant to biological weapons development; North Korea and Iran are assessed as sustaining offensive programs. For these actors, the risk does not lie in access to information. State programs have the resources to develop expert teams and access nonpublic scientific materials. The relevant question for sophisticated state actors is whether biological design tools and biological foundation models can raise the ceiling of what is technically achievable, producing agents more precisely engineered, less subject to the transmissibility-lethality evolutionary tradeoff, and better suited to strategic objectives. If AI-enabled design makes biological weapons more controllable, more targetable, less attributable and less prone to blowback, their strategic attractiveness increases substantially, and the logic that has historically limited state use becomes weaker.

The second category is **organized nonstate groups with meaningful resources**: large terrorist organizations capable of recruiting scientific expertise, or well-funded criminal syndicates motivated by economic gain or hired attack. Several expert submissions described scenarios in which criminal actors used protein language models fine-tuned on toxicological datasets to engineer novel toxins that bypass existing screening controls. This tier is the fastest-changing part of the threat landscape. Previously constrained by information and expertise barriers, organized nonstate actors are most directly enhanced by LLM-driven information democratization. Experts at the Centre for Long-Term Resilience (CLTR) therefore identify “somewhat capable groups”—those with some scientific background but not frontier expertise—as likely to see the greatest near-term uplift from advancing AI trends, driven specifically by information access and LLM-enabled experimental troubleshooting.⁵⁹

The third category is **skilled individuals and small cells**: lone wolves, biohackers, or disgruntled insiders exploiting legitimate research access. The 2001 Amerithrax attacks demonstrate how a single insider with laboratory access can inflict large-scale harm without advanced design capabilities. Aum Shinrikyo shows a complementary lesson: A well-motivated group with genuine scientific expertise failed not for lack of effort but because of specific knowledge gaps that current AI systems could plausibly have bridged. For this category, AI-assisted troubleshooting and information synthesis is the main risk.

A fourth category, distinct from the others, is **the accidental actor**: the well-intentioned researcher whose work produces dangerous unintended consequences, or the inadequately secured laboratory from which a modified pathogen escapes. As high-containment facilities proliferate and AI-accelerated research workflows compress the timelines over which dangerous biology can

59 Centre for Long-Term Resilience, “The Near-Term Impact of AI on Biological Misuse,” July 26, 2024, <https://www.longtermresilience.org/reports/the-near-term-impact-of-ai-on-biological-misuse/>.

be conducted, the potential for consequential accidents expands. Establishing and maintaining best-practice containment measures is critical, in the American homeland and around the world.

Relevant Capabilities

The threat actors described above differ in motivation, resources, and institutional backing. What determines what each can achieve is the set of capabilities available to them. Two intersecting domains define that question: advanced biotechnology and AI.

Biotechnology advances: The cost of DNA synthesis has fallen from approximately \$4 per base pair in the early 2000s to \$0.10 today; whole genome sequencing, which cost hundreds of millions of dollars two decades ago, can now be done for \$200. CRISPR-Cas9 has placed precise gene editing within reach of researchers without specialized equipment. Benchtop DNA synthesizers enable on-site synthesis without requiring commercial screening protocols. Contract research organizations and cloud labs allow actors to outsource experimental work to facilities with expertise, equipment, and labor, reducing the need for in-house infrastructure and human capacity. The published scientific literature now includes detailed protocols for reverse genetics—reconstructing a live pathogen from its genome sequence alone—that were previously confined to specialized facilities. Crucially, several experts described scenarios that required no AI capability at all: serial passage—the iterative exposure of animals to a pathogen to select for enhanced transmissibility—relies on standard laboratory equipment and published protocols that have existed for decades. Ordered-fragment assembly—acquiring separately ordered DNA fragments from multiple commercial providers and combining them to reconstruct a dangerous pathogen—is similarly a conventional technique whose limiting factor has been knowledge access rather than specialized equipment. AI accelerates and extends access to these pathways, though it does not yet create the biological threat from scratch. A helpful tool in this respect is RAND’s Global Risk Index for AI-Enabled Biological Tools.⁶⁰

AI Advances: Overlaid on this foundation, the four AI system categories interact with biotechnology platforms in differentiated ways across actor tiers. LLMs primarily benefit actors at the lower capability tiers by lowering information and tacit knowledge barriers that previously required specialist expertise: synthesizing across sources, troubleshooting experimental failures in real time, and providing step-by-step guidance tailored to an actor’s specific circumstances and available resources. Biological foundation models and biological design tools primarily extend the capabilities of actors with existing wet-lab infrastructure, enabling design objectives previously out of reach, including the engineering of novel sequences that do not resemble known pathogens and therefore bypass sequence-based screening controls. AI science agents compound all three AI system classes in integrated workflows, potentially enabling actors with limited specialist knowledge to execute multistep biological research workflows that previously required coordinated expert teams. The most technically proximate near-term scenarios rely on conventional biotechnology accelerated by LLM information uplift; the highest-consequence scenarios depend on the convergence of state-level AI design capability with wet-lab infrastructure.

60 RAND, “AI Risk Index,” accessed July 5, 2026, <https://www.rand.org/randeurope/research/projects/2024/ai-risk-index.html>.

Range of Possible Outcomes

The experts assessed the shortlisted scenarios across five impact dimensions: human cost (fatalities and casualties), economic cost (direct monetary harm), environmental cost (geographic extent and reversibility of damage), governance and political cost (disruption to public services and confidence in public institutions and leaders), and geopolitical cost (degree of international stress, conflict, or escalation risk). Scenarios were also assessed for probability of occurring by 2030 in the absence of additional intervention.

Across the long list of scenarios evaluated, threat vectors and outcomes varied substantially. On human cost, scenarios ranged from targeted assassination agents causing tens of casualties to engineered pandemic pathogens with civilizational impacts. On economic cost, scenarios ranged from pharmaceutical and food supply chain disruption worth hundreds of millions of dollars to the multitrillion-dollar losses associated with large-scale pandemic events. On governance and political cost, several scenarios described events that caused mass panic and erosion of institutional trust independent of physical casualties. A designer toxin released in a crowded venue might cause limited direct fatalities while generating sustained public panic around the apparent invisibility of the threat, a dimension that conventional casualty-focused threat assessments risk underweighting. On geopolitical cost, scenarios included creating attributional ambiguity, triggering Biological Weapons Convention withdrawal, provoking biological weapons arms race dynamics, or inducing nuclear escalation risk, outcomes that depend less on the scale of biological harm than on the strategic context in which the event occurs and the speed with which attribution can be established.

The three composite scenarios presented below were selected on three criteria: breadth of threat actor profiles represented, diversity of harm mechanisms illustrated, and strength of grounding in near-term technical feasibility. The first reflects the near-term nonstate, LLM-enabled, information-barrier-limited pathway. The second reflects the state-level, BDT-enabled strategic degradation scenario at the high-consequence end of the design capability spectrum. The third reflects a distinct mechanism, agricultural and food security disruption, illustrating that the biological threat space extends well beyond the direct casualty paradigm. Together, the three scenarios do not exhaust the space of plausible threats to 2030. Instead, they help illuminate its contours and the actor and capability dynamics critical to understanding it.

Scenario 1: Release of a Highly Pathogenic Virus by a Network of Politically Motivated Lone Wolves

A decentralized network of politically motivated “lone wolves” uses jailbroken, open-source LLMs to identify means of delivering existing pathogens into large human populations. Benign-appearing fragments of a pathogen are successfully ordered from multiple commercial DNA synthesis providers. A single expert individual is ultimately able to assemble and release a highly pathogenic virus into an urban area.

This scenario is the closest to current technical feasibility. It does not require capabilities that do not yet exist; it requires chaining them across three sequential steps, each of which presents

meaningful but declining barriers.

The first step is **information acquisition**. This is the barrier most clearly within current reach. Open-source LLMs with meaningful biological knowledge are freely downloadable, and jailbreaking them, while increasingly difficult, remains achievable by motivated actors with modest technical sophistication. The more significant development is not that these models can be coaxed into providing dangerous information, but that they can synthesize fragmented public scientific knowledge in ways that previously required expert acquisition and interpretation. A lone actor can now receive a coherent, tailored explanation of a complex biological procedure without the graduate-level training at a credentialed institution that once served as an implicit gatekeeping function.

The second step is **material acquisition**. This means obtaining the physical biological or genetic materials needed for the work, including commercially synthesized DNA fragments. Experts have warned that a malicious actor could divide a genome across orders from different suppliers, making each individual order less likely to raise screening concerns. This vulnerability is particularly important because synthetic DNA fragments can then be assembled into an infectious virus using techniques known as reverse genetics, which have become more accessible as the underlying methods and materials have become more widely available. In 2020, researchers showed that infectious SARS-CoV-2 could be reconstructed within one week of receiving most of the synthetic DNA fragments they needed via standard commercial channels.⁶¹

The critical remaining barrier is **tacit knowledge**. This is the embodied, hands-on laboratory skill that allows a researcher to troubleshoot failures, handle dangerous biological material safely, and successfully execute protocols rather than merely understand them theoretically. AI-provided information does not currently substitute for this kind of practical knowledge and judgment. There is also a fundamental biological constraint: The traits that a hostile actor would find most desirable for enhanced transmissibility and increased lethality are frequently in evolutionary tension. A pathogen optimized for one characteristic tends to perform less effectively on the other.

Scenario 2: A Hostile Nation-State Releases a Stealth Respiratory Pathogen into an Adversary's Population

A sophisticated nation-state program leverages biological design tools to engineer a “stealth” respiratory pathogen. The agent is designed to have negligible acute symptoms but causes severe debility or death after a long incubation period of more than six months. The attacker targets the military readiness and civilian workforce of key adversaries while masking the biological nature of the event to avoid direct attribution.

The nation-state scenario represents a qualitatively different category of threat from the lone wolf scenario, distinguished not only by scale but by design ambition. Where nonstate actors are most constrained by tacit knowledge and material access, state programs are constrained primarily by the complexity of the design objective itself.

61 Thi Nhu Thao et al., “Rapid Reconstruction of SARS-CoV-2 Using a Synthetic Genomics Platform,” *Nature* 582 (2020): 561-65, <https://doi.org/10.1038/s41586-020-2294-9>.

The scenario's technical requirements are demanding but within reach of a well-resourced program. The **core design challenge** is engineering an agent that combines three attributes that are individually achievable but collectively difficult: **high transmissibility**, to seed the attack across a target population; **delayed onset of severe symptoms**, to preserve operational surprise and prevent early containment; and **non-attribution**, to prevent the event from being recognized as a deliberate rather than a naturally occurring event.

The scenario's strategic logic is also its **most significant deterrence challenge**. An agent designed for a six-month or longer incubation period is specifically engineered to prevent recognition before widespread harm has occurred. By the time debility and mortality manifest in an adversary's military and civilian population, the release event is months in the past; the pathogen has spread through ordinary human movement across borders; and forensic determination that the event was deliberate, rather than a natural outbreak, requires exactly the genomic surveillance and attribution science infrastructure that most nations do not currently have. The scenario is most directly deterred not by preventing the release, which a determined, well-resourced state program could likely achieve, but by developing early detection and attribution capability that raises the probability of rapid identification, and response capacity that limits the strategic payoff of the attack.

Practical deterrents constrain the scenario's near-term probability. The risk of **blowback** is a significant deterrent for any transmissible pathogen. At the present time, genetic targeting of specific ethnic populations is extremely challenging: Neighboring populations that bear the greatest historical animus toward each other are frequently among the most genetically similar. And the development and testing of a novel engineered pathogen at a state-program scale would require a **physical and biosafety infrastructure** that is difficult to conceal entirely from intelligence collection.

Scenario 3: An Ideologically Motivated Terrorist Group Targets Large-Scale Agriculture to Impact Food Supply

A terrorist group releases a genetically modified plant pathogen, such as a novel corn rust, specifically targeting the staple crops of an adversary nation. The disease spreads through global trade channels and is compounded by climate stressors, decimating food supply chains.

The agricultural scenario is the most underappreciated of the three, in part because it does not conform to the pandemic casualty paradigm that can dominate biosecurity analysis. Its consequences are measured primarily in economic impact, food security, and geopolitical leverage rather than direct fatalities, but these effects can be severe, cascading, and difficult to reverse.

Existing natural plant pathogens already demonstrate the scale of harm this pathway can produce. Wheat stem rust, a fungal pathogen, destroys an estimated 5.5 million metric tons of wheat

annually.⁶² Rice blast is estimated to destroy enough rice each year to feed 60 million people. These losses occur without deliberate adversarial intervention.

A modified plant pathogen engineered for enhanced virulence against specific crop varieties, or for resistance to existing fungicide and treatment protocols, could produce losses substantially exceeding the natural baseline. The economic consequences of even a credible threat of agricultural contamination are significant independent of actual crop damage: Export restrictions triggered by pathogen detection can cause immediate food price spikes in import-dependent nations, and the uncertainty created by a confirmed biological attack on food supply would generate cascading effects on commodity markets, food security, and diplomatic relationships that extend far beyond the affected crop.

What makes the agricultural scenario **particularly difficult to govern** is that effective harm **does not require advanced biological engineering**. An adversary does not necessarily need to engineer a novel pathogen: Acquiring and cultivating existing, highly virulent plant pathogen strains, then introducing them into a target nation's agricultural system through trade shipments or covert release, requires substantially less technical sophistication than engineering a human pathogen and may achieve comparable economic effect. Climate change is already expanding the geographic range and seasonal window of several significant agricultural pathogens, potentially reducing the adversarial design burden by improving environmental conditions for spread. However, there is one factor that may limit this specific bioconvergence threat in the near term: AI tools are trained primarily on human biology and bacterial genomics rather than plant biology and, as of 2026, have not yet been effectively applied to plants.⁶³

Aggregate Assessment: Drivers of Bioconvergence Risk

Taken together, analysis of the three threat scenarios identifies structural features of the current bioconvergence risk landscape and points to the most critical gaps in the biosecurity architecture.

- First, biological events have historically been avoided because the pool of actors with both the capability and the intent to weaponize biology has been very small. Those who understood the dangers most intimately have generally been the least likely to deploy them for harm, a pattern that often reflects the restraining effect of scientific knowledge itself, as well as blowback and reputational risks. Advances in AI and biotechnology are broadening access to a much broader and less expert range of actors not bound by the same norms and inhibitions.
- Second, tacit knowledge is a significant remaining barrier, but a shrinking one. Explicit knowledge barriers have been substantially lowered by LLMs. Laboratory automation and AI-based troubleshooting now threaten to lower the tacit knowledge barrier further, compressing the gap between theoretical understanding and operational capability. The timeline over which this barrier provides meaningful protection is shrinking.

62 Food and Agriculture Organization of the United Nations, "Wheat Stem Rust - Ug99," accessed April 2026, <https://www.fao.org/agriculture/crops/rust/stem/rust-report/stem-ug99racettksk/en/>; Deborah Pagliaccia et al., "Genetic Structure of the Rice Blast Pathogen (*Magnaporthe oryzae*) over a Decade in North Central California Rice Fields," *Microbial Ecology* 75 (2017): 310–17.

63 Lin Yang et al., "Artificial Intelligence-Driven Plant Bio-Genomics Research: A New Era," *Tropical Plants* 4 (February 2025), <https://doi.org/10.48130/tp-0025-0008>.

- Third, nation-state-level programs cannot be deterred by commercial channel interventions alone. Measures such as DNA synthesis screening may be effective against less sophisticated nonstate actors but nation-state programs have the means to circumvent such chokepoints, and must be deterred through investments that decrease the expected consequences of an attack, primarily through detection, attribution, response, and recovery capacity.
- Fourth, the absence of a multilayer, ubiquitous, and continuous detection and response infrastructure magnifies the risk of all three scenarios. A sophisticated biological attack against a nation with robust detection and response capabilities would have far lower expected impact than the same attack against a nation that cannot identify the threat until it has already spread through the population.

Understanding why the governance preparation and response structure that this risk landscape requires is so different from prior frameworks is clarified by examining what prior technologies, particularly nuclear and cyber, can and cannot teach us about managing potentially catastrophic risks.

E. Analogizing Bioconvergence Risks

Reasoning by analogy is a natural starting point for governance design. Bioconvergence resembles nuclear and cybersecurity threats in some respects and diverges in important ways. Mapping both clarifies which governance tools translate and which must be built from scratch, and guards against the false confidence of treating bioconvergence as merely a variant of problems that have previously been managed.

The Nuclear Analogy

The nuclear weapons analogy is the most commonly invoked in discussions of catastrophic biological risks. Like nuclear weapons, certain biological threats carry existential potential. Like nuclear weapons, states are a relevant threat actor. Four differences demonstrate where the analogy breaks down.

Physical chokepoints. Nuclear weapons development requires highly specialized materials—specifically enriched uranium or plutonium—whose production is capital-intensive, technically demanding, and leaves physical and radiological signatures. These create impediments to acquisition as well as interdiction points at which proliferation can be detected and interrupted, as reflected in decades of nonproliferation initiatives led by the US and other governments globally. There is no clear biological equivalent of a uranium enrichment facility or a plutonium reactor. Biological agents are often found in nature or can be synthesized from commercially available reagents using equipment with widespread legitimate uses. They are ubiquitous and varied.

Separation of civil and military applications. The 1953 “Atoms for Peace” proposal, followed by the 1968 Non-Proliferation Treaty (NPT), distinguished between nuclear power and nuclear weapons programs. IAEA safeguards are intended to ensure that nuclear material is not diverted from civil to military applications. But in biology, the same tools, reagents and techniques that

underpin vaccine development or cancer research can be used to cause harm. Separating peaceful and malicious use is challenging.

International architecture and norms. The five nuclear powers of the Cold War have remained vigilant about the security of their nuclear arsenals and the spread of fissile materials. After the Cold War, the Nunn-Lugar Cooperative Threat Reduction Program effectively reduced the threat of “loose nukes.” Such nuclear nonproliferation efforts have been sustained by successive US administrations. Additionally, the NPT has helped to keep the number of nuclear powers to nine. Biology does not have the same history of cooperative threat reduction between states, and the Biological Weapons Convention is vastly weaker than the NPT, with no verification mechanism, inspectorate, or enforcement authority.

Commercial competition. There has been little pressure to profit from nuclear weapons. Nuclear materials and delivery systems have remained the preserve of states, controlled within classified government programs. The opposite is true of bioconvergence. Advances in AI, biotechnology, and engineering are being led by the private sector. This divergence and its implications are discussed further in Chapter 3.

The Cyber Analogy

The cyber domain offers a second instructive comparison. Like bioconvergence, cyber threats are diffuse, attribution-resistant, fast-moving, and shaped by an offense-defense dynamic in which new capabilities continuously require new defenses. Three features of the biological threat context nonetheless challenge this analogy.

The actor population. Cyberattacks can be monetized. Ransomware, financial theft, and intellectual property exfiltration fund a large and self-sustaining criminal ecosystem that continuously drives innovation in offensive techniques. No comparable business model exists for biological weapons. The deliberate use of biological agents remains unattractive except for actors motivated by ideology, geopolitics, or mass destruction for its own sake. This fundamentally changes the threat actor population: The cyber landscape includes a vast criminal middle tier with no biological equivalent.

Incentives for defense. Cybersecurity costs are partially internalized within private firms because those firms are themselves the victims of cyberattacks. This creates commercial incentives for private cybersecurity investment that have partially corrected the market failure and driven a large private security industry. Biosecurity costs are almost entirely externalized. The harm from a biological attack falls on the public and governments, not on individual biotechnology or AI firms whose practices may have contributed to the risk.

Self-replication and scale of harm. The consequences of large-scale cyberattacks are measured primarily in economic disruption or degraded military capability, though it is possible to envision incidents of serious societal disruption. Biological attacks pose similar threats, but also add new risks due to the self-replicating nature of biological agents. Cyberattacks can cause cascading failures, but biological attacks can spread through entire populations after one release event. The direct consequences of a single successful biological attack could therefore rise to millions of deaths.

Drawing on both analogies

Both analogies point toward the need for greater governance domestically and globally. Like nuclear threats, bioconvergence demands sustained attention to state-level programs and a fit-for-purpose international legal architecture. Like cyber threats, biosecurity needs layered technical defenses, public-private coordination, and an acceptance that offense and defense will continuously co-evolve. Unlike either analogy, the bioconvergence challenge presents a dual-use landscape in which the same tools that could enable biological threats are essential to the prevention, readiness and resilience efforts that civilization depends upon.

The stark reality is that governing bioconvergence is and will be harder than governing either nuclear or cyber threats. The absence of physical chokepoints, the totality of the dual-use problem, commercial pressures to build market value in an intensely competitive landscape, and the inadequacy of existing international architecture make the policy challenge genuinely novel.

The decisions made in the near term about how AI and biotechnology are governed, how biological data are managed, how detection infrastructure is built, and how recovery capacity is sustained will determine whether bioconvergence's consequences are catastrophic or strategically managed. The chapters that follow describe what must change for humanity to navigate the era of bioconvergence.

Prevention

Prevention in the context of the deterrence by resilience framework means creating and enhancing barriers to stop bad actors from misusing the technology of bioconvergence. Unlike traditional security approaches that rely primarily on punishment after the fact, prevention shapes the technological, regulatory, and institutional environment to make harmful applications of AI-enabled biotechnologies harder to achieve while preserving beneficial innovation.

Prevention operates across three intersecting domains:

- governance of AI tools with biological applications;
- regulation of biological data that trains those tools; and
- the physical and procedural controls on DNA synthesis, laboratory access, and cloud-based research services.

Each domain has distinct gaps, and the challenges in each are growing. The cost of designing and acquiring biological capabilities is falling rapidly. AI tools that once required institutional resources and specialized expertise are now accessible to a widening range of actors. This democratization of capability creates structural risks that need to be addressed by government and private sector action.

A. Bio Risks of General Purpose AI Models: Current State and Critical Gaps

Despite the scope of the threats to national interests, the United States' approach to preventing bioconvergence misuse is fragmented across multiple regulatory domains, enforcement mechanisms, and jurisdictions. Federal oversight relies heavily on voluntary industry standards, federal funding conditions, and narrow list-based approaches that do not recognize novel threats. Initiatives at the US state level are endeavoring to address some gaps, particularly for AI safety, but also will likely create a patchwork of inconsistent requirements and be faced with federal challenges. International frameworks provide additional pressure but lack enforcement mechanisms for the most dangerous applications.

Managing the Tension Between Innovation and Regulation

Pitting innovation against regulation is particularly problematic in the biosecurity context. For AI systems that can accelerate biological design and scientific discovery, the relevant policy choice should not be between innovation or regulation, but rather about which governance model most effectively advances benefits while reducing catastrophic risks. In principle, and based on best

established biosecurity practices, public policy can set rules and guidelines for managing shared risks that make innovation more durable by increasing trust, reducing systemic vulnerabilities, and creating predictable paths to deployment. Certainly, regulation imposes costs on private firms and can, if excessive, impede competitiveness. However, companies generating novel possibilities that also give rise to new and potentially existential risks should bear some responsibility to manage those risks. Technically informed and appropriately tailored public policy can and should balance security mandates and incentive structures in order to preserve the leading edge of US industry innovation and simultaneously protect public interests.

Market forces often correct for externalities without government intervention, but markets can also underprovide public goods such as public health and security.⁶⁴ This is evident in the behavior of AI and biotechnology companies. AI companies have implemented varying, voluntary policies to test their models for biological threats, including Anthropic, OpenAI, and DeepMind, but the path is not linear and there is conflicting information about how effective or comprehensive their guardrails against misuse are or will be. Further, in some cases they have actually rolled back safety and transparency measures,⁶⁵ and a recent independent expert report graded frontier companies' safety and security measures as C+ at best, with several failing completely.⁶⁶ The bottom line is that most biotechnology firms also do not prioritize biosecurity in the development of products because their principal focus is on delivering results and generating profits.

In public health, government regulation and incentives have shaped the advancement of US companies that meet certain standards. For example, in pharmaceuticals, medical devices, and biologics, regulations have created frameworks within which firms could invest, test, commercialize, and scale products with greater public legitimacy.⁶⁷ This is particularly salient for AI, about which recent surveys indicate a majority of American consumers are concerned, and over 60 percent distrust companies to develop responsibly.⁶⁸ Consumer trust is key to companies' financial futures, and building it starts with demonstrating consistent, reliable, and verifiable safety measures.

The false tradeoff between innovation and regulation is particularly notable in biotechnology, where government funding and conditions can simultaneously make research faster, improve scientific quality, strengthen oversight, and reduce biosecurity risks.⁶⁹ Research ecosystems with security-by-design can improve scientific progress, company viability, and biosecurity.⁷⁰ For example, a recent trailblazing analysis examined three bioconvergence case studies: secure and

64 Suerie Moon, John-Arne Rottingen, and Julio Frenk, "Global Public Goods for Health: Weaknesses and Opportunities in the Global Health System," *Health Economics, Policy and Law* 12 (2017): 195–205.

65 For example, in 2025, OpenAI released GPT-4.1 without publishing safety information in a system or model card. Maxwell Zeff, "OpenAI Ships GPT-4.1 Without a Safety Report," *TechCrunch*, Apr. 15, 2025, <https://techcrunch.com/2025/04/15/openai-ships-gpt-4-1-without-a-safety-report/>.

66 Future of Life Institute, "AI Safety Index: Summer 2026," July 2026, <https://futureoflife.org/wp-content/uploads/2026/07/AI-Safety-Index-Summer-2026-Digital.pdf>.

67 For example, federal oversight of biologics dates back to the Biologics Control Act of 1902, and the FDA's science-based regulation of biological products supported safer and more effective products over time. US Food and Drug Administration, "The History of Biologics Regulation," Apr. 24, 2019, <https://www.fda.gov/about-fda/histories-fda-regulated-products/history-biologics-regulation>.

68 Pew Research Center, "Americans and AI 2026: Chatbots, Smart Devices and Views on Impact," June 17, 2026, <https://www.pewresearch.org/internet/2026/06/17/americans-and-ai-2026-chatbots-smart-devices-and-views-on-impact/>; Pew Research Center, "Key Findings About How Americans View Artificial Intelligence," Mar. 12, 2026, <https://www.pewresearch.org/short-reads/2026/03/12/key-findings-about-how-americans-view-artificial-intelligence/>.

69 Casey Avesaggio, Sarah R. Carter, Harshini Mukundan, Elizabeth E. Cameron, and Steph Guerra, "Aligning Innovation and Security in AI-Enabled Biotechnology: A Framework for Designing and Funding Mutually Reinforcing Approaches," *Frontiers in Microbiology* 17 (2026): 1–6.

70 Ibid.

tiered data-sharing infrastructures for datasets; data provenance and metadata tracking; and continual evaluation and benchmarking of biological tools. It found that tailored approaches can ensure that biosecurity measures mutually reinforce innovation within the technology itself.⁷¹

The relevant challenge now is how to build governance for AI-enabled biotechnology that is proportionate to capability, responsive to dual-use benefits and risks, and sensitive enough to promote the technology breakthroughs necessary to maintain US competitiveness as well as the sharp scientific edge that will be needed to respond effectively to weaponized biological threats.

Federal Approach: Legislative and Executive Action

Legislative Landscape

The US Congress has yet to pass legislation imposing any safety requirements on AI models or on biological risks associated with the models. Instead, the legislative approach remains disaggregated and sector-specific. Although lawmakers have introduced hundreds of AI-related bills since 2019, only a small fraction have become law, and most enacted measures focus on appropriations, research and development, or institutional coordination rather than enforceable safety standards.⁷²

No federal law creates a centralized regulatory authority for AI oversight, or confers clear rulemaking authority to impose such standards. Proposed legislation addressing AI safety spans a wide spectrum from more stringent approaches, such as requiring pre-deployment safety testing and regulatory approval,⁷³ to proposals that would preempt state-level AI regulations.⁷⁴ Other proposals emphasize accountability mechanisms, including third-party audits, liability frameworks, and regulatory “sandboxes” to test AI systems under controlled conditions.⁷⁵ However, these proposals remain largely under development.

Within this broader legislative landscape, relatively few measures directly address the biological risks associated with advanced AI systems. However, Members of Congress have attempted to address bioconvergence risks with several proposed pieces of legislation:

⁷¹ Ibid.

⁷² For example, the National Artificial Intelligence Initiative Act established a coordinated federal framework for AI research and policy, directing agencies such as the National Institute of Standards and Technology to develop voluntary safety standards updated on a periodic basis and tasking agencies including the Department of Energy and the National Science Foundation with advancing AI research, development, and workforce development. William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021, Pub. L. No. 116-283, div. E, 134 Stat. 3388 (2021) (codified at 15 U.S.C. § 9401). Similarly, the AI in Government Act and the Advancing American AI Act focus on internal federal governance: They require agencies to document AI use cases, develop procurement safeguards, and identify best practices to mitigate bias and discriminatory impacts in AI systems. Consolidated Appropriations Act, 2021, Pub. L. No. 116-260, div. U, title I, 134 Stat. 1182 (2020) (codified at 40 U.S.C. § 11301 note); James M. Inhofe National Defense Authorization Act for Fiscal Year 2023, Pub. L. No. 117-263, title LXXII, 136 Stat. 2395 (2022). The CHIPS and Science Act further reinforces this approach by funding AI research and supporting the development of technical standards for trustworthy AI. CHIPS and Science Act of 2022, Pub. L. No. 117-167, 136 Stat. 1366.

⁷³ Algorithmic Accountability Act, H.R. 5511, 117th Cong.

⁷⁴ American Artificial Intelligence Leadership and Uniformity Act, H.R. 5388, 119th Cong.

⁷⁵ SANDBOX Act, S. 2750, 119th Cong. (2025); TEST AI Act of 2025, S. 1633, 119th Cong. (2025); Marsha Blackburn, The Republic Unifying Meritocratic Performance Advancing Machine Intelligence by Eliminating Regulatory Interstate Chaos Across American Industry Act (TRUMP AMERICA AI Act): Section-by-Section Analysis (Washington, DC: US Senate, n.d.), <https://www.blackburn.senate.gov/services/files/C43D3B19-391B-4EB6-84C1-0FC37EEBBA4D>.

- The **Artificial Intelligence and Biosecurity Risk Assessment Act**, introduced in 2023, would have required the Department of Health and Human Services (HHS) Administration for Strategic Preparedness and Response to regularly assess the risk of AI being “used intentionally or unintentionally to develop novel pathogens, viruses, bioweapons, or chemical weapons.”⁷⁶ The legislation has not been reintroduced in the 119th Congress.
- The **Strategy for Public Health Preparedness and Response to Artificial Intelligence Threats Act** would require HHS to produce a national strategy for preparedness, response, and biodefense against the misuse of AI, especially where AI could facilitate biological or chemical weapons or other national health-security threats.⁷⁷
- The **Generative AI Terrorism Risk Assessment Act** would require DHS to assess terrorist threats arising from generative AI, including use related to chemical, biological, radiological, or nuclear weapons.⁷⁸
- Finally, the recently introduced **Biosecurity Modernization and Innovation Act**, while not directly addressing AI models, seeks to prevent AI-enabled biosecurity threats by requiring gene synthesis providers to screen orders and customers, create a biotechnology governance sandbox at the National Institute for Standards and Technology (NIST), and require an assessment of gaps in federal biosecurity authority.⁷⁹

(See **Appendix 2** for a more comprehensive table of proposed legislation concerning biosecurity and AI.)

In September 2025, the statutorily mandated National Security Commission on Emerging Biotechnology (NSCEB) published its report on the future of American biotechnology, which included detailed recommendations to Congress around six pillars: Prioritize biotechnology at the national level; mobilize the private sector to get US products to scale; maximize the benefits of biotechnology for defense; out-innovate our strategic competitors; build the biotechnology workforce of the future; and mobilize the collective strengths of our allies and partners. The NSCEB identified key areas where Congress in particular could act decisively to protect US national security and American industry. To date, comprehensive legislation has not been passed.

In sum, US legislative efforts to address AI and biosecurity have not achieved the necessary scope and scale that this moment demands. Existing laws prioritize coordination, research, and voluntary standards, while proposed measures—particularly those focused on biological risks—center on assessment and preparedness rather than enforceable constraints. As discussed throughout this report, legislative action to address the risks of bioconvergence needs to be cross-sectoral—spanning AI models, biotechnology, and public health infrastructure—as well as comprehensive—**across the entire life cycle of AI-enabled bioconvergence challenges.**

76 Artificial Intelligence and Biosecurity Risk Assessment Act, S. 2399, 118th Cong. (2023); Artificial Intelligence and Biosecurity Risk Assessment Act, H.R. 4704, 118th Cong. (2023).

77 Strategy for Public Health Preparedness and Response to Artificial Intelligence Threats Act, S. 501, 119th Cong. (2025).

78 Generative AI Terrorism Risk Assessment Act, H.R. 1736, 119th Cong. (2025).

79 Biosecurity Modernization and Innovation Act of 2026, S. 3741, 119th Cong. (2026).

Executive Actions

Executive branch actions across multiple administrations have increasingly recognized that advances in artificial intelligence, particularly highly capable foundation models, may create or amplify biological and biosecurity risks.⁸⁰ The US executive branch has deployed a range of incentives to encourage AI companies to adopt biosecurity safeguards. These include procurement tools, national security authorities, funding conditions, and support for voluntary standards.

AI Safety

In 2019, the president directed the development of voluntary technical standards for reliable, robust, trustworthy, secure, portable, and interoperable AI systems.⁸¹ The following year, the administration extended this framework to govern federal agencies' adoption of safe AI applications.⁸² In 2023, Executive Order 14110, *Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence*,⁸³ directed 100 actions from fifty federal entities organized in eight policy pillars. It required companies developing or intending to develop dual-use AI models to report to the government on model training, testing, and data ownership; and entities that acquire, develop, or possess potential large computing infrastructure to report to the government on the location and amount of computing power. It also provided the first federal government directions on the convergence of AI and biotechnology, providing detailed tasks to DOD, OSTP, and NIST on biosecurity, synthetic nucleic acids, sequencing screening, and databases.⁸⁴

EO 14110 was followed in October 2024 by National Security Memorandum (NSM) 25, which mapped the EO's requirements onto national security agencies.⁸⁵ NSM-25 also provided detailed taskings regarding biological risks and AI, including testing by NIST focused on the potential for AI models to promote the development of biological and chemical weapons; the creation by NIST of voluntary standards for publishing computational biological and chemical models; and Committee on Foreign Investment in the US review of foreign investments in AI.

80 The White House, *National Biodefense Strategy* (2018); National Security Commission on Artificial Intelligence, *Final Report* (March 2021); Exec. Order No. 14110, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," 88 Fed. Reg. 75,191 (Oct. 30, 2023).

81 Exec. Order No. 13859, "Maintaining American Leadership in Artificial Intelligence," Feb. 11, 2019.

82 Exec. Order No. 13960, "Promoting the Use of Trustworthy Artificial Intelligence in the Federal Government," Dec. 3, 2020; Office of Management and Budget, *Guidance for Regulation of Artificial Intelligence Applications*, M-21-06 (Nov. 17, 2020).

83 Exec. Order No. 14110, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence."

84 These requirements included: DHS: Evaluate and assess potential risks related to critical infrastructure adoption and use of AI and consider ways to mitigate vulnerabilities; develop a framework to conduct structured evaluation and stress testing of the nucleic acid procurement screening systems. DOD: Study and report on the use of AI in biology and the potential to increase biosecurity and national security risks and make recommendations to mitigate such risks. OSTP: Develop a framework for providers of synthetic nucleic acids to encourage such providers to implement procurement screening mechanisms, including standards and incentives. OSTP and NIST: Develop procurement screening specifications and guides, and best practices for sequence-of-concern database management and conformity-assessment, for use by nucleic acid sequence providers. All agencies that fund life science research: Establish requirements that life science researchers receiving federal funding must procure synthetic nucleic acids from manufacturers adhering to the framework established in the EO.

85 National Security Memorandum 25, *Advancing the United States' Leadership in Artificial Intelligence; Harnessing Artificial Intelligence to Fulfill National Security Objectives; and Fostering the Safety, Security, and Trustworthiness of Artificial Intelligence* (Oct. 24, 2024), <https://bidenwhitehouse.archives.gov/briefing-room/presidential-actions/2024/10/24/memorandum-on-advancing-the-united-states-leadership-in-artificial-intelligence-harnessing-artificial-intelligence-to-fulfill-national-security-objectives-and-fostering-the-safety-security/>.

In January 2025, the president rescinded EO 14110 and ordered the creation of a new AI Action Plan, which was published in July 2025.⁸⁶ NSM-25 was left in place. The AI Action Plan was a nonbinding framework for AI priorities including the following elements relevant to bioconvergence:

- Acknowledgment that AI systems may pose “novel national security risks in the near future in areas such as cyberattacks and the development of chemical, biological, radiological, nuclear, or explosives (CBRNE) weapons”;
- Recommendation that the government evaluate frontier AI systems for national security risks, including CBRNE and cyber risks;
- Identification of investment in biosecurity as a component of AI strategy;
- Recommendation that “all institutions receiving [f]ederal funding for scientific research to use nucleic acid synthesis tools and synthesis providers that have robust nucleic acid sequence screening and customer verification procedures.”
- Creation of enforcement mechanisms for this requirement rather than relying on voluntary attestation.

The priorities have yet to be translated into new policies. In March 2026, the White House released its *National Policy Framework for AI*, which proposes federal legislation that would create a single national AI standard, largely preempt stricter state AI laws, accelerate AI infrastructure development, protect AI companies from some third-party liability, and promote a light-touch regulatory approach. Safeguards for biological risks were absent from this proposed AI legislative framework. In June 2026, following growing concerns about the cyber capabilities of new AI models like Anthropic’s Mythos,⁸⁷ the president issued an executive order on US federal government review of advanced AI models.⁸⁸ The new EO initiates a cybersecurity review process under which developers of covered frontier AI models may provide the US government access to their models up to thirty days before public release; this review remains voluntary and limited to cybersecurity risks, leaving out other potentially catastrophic risks including AI-enabled biological risks.⁸⁹

Biosecurity Initiatives

Several US initiatives and strategies have tackled bioconvergence with a whole-of-government approach, spanning the Bush, Obama, Trump, and Biden administrations. Most recently, in 2022, following the COVID-19 pandemic, the National Biodefense Strategy and Implementation Plan envisioned a detailed whole-of-government strategy with agency leads and milestones to improve surveillance, detection, forecasting of biological threats; prevent bio-accidents; strengthen biodefense infrastructure; and enable rapid response.⁹⁰ The 2022 executive order on bioeconomy

86 Exec. Order No. 14148, “Initial Rescissions of Harmful Executive Orders and Actions,” Jan. 20, 2025; Exec. Order No. 14179, “Removing Barriers to American Leadership in Artificial Intelligence,” Jan. 20, 2025.

87 Gordon M. Goldstein, “Six Reasons Claude Mythos Is an Inflection Point for AI and Global Security,” Council on Foreign Relations, Apr. 15, 2026, <https://www.cfr.org/articles/six-reasons-claude-mythos-is-an-inflection-point-for-ai-and-global-security>.

88 Exec. Order No. 14409, “Promoting Advanced Artificial Intelligence Innovation and Security,” June 2, 2026.

89 Council on Foreign Relations, “Assessing Trump’s Executive Order on AI Oversight,” June 2, 2026, <https://www.cfr.org/articles/assessing-trumps-executive-order-on-ai-oversight>.

90 National Security Memorandum 15, National Biodefense Strategy (Oct. 18, 2022).

also recommended a new effort at HHS to oversee biosecurity and biosafety as part of the innovation cycle.⁹¹

Two years later, in 2024, the White House Office of Science and Technology Policy (OSTP) issued a government-wide Framework for Nucleic Acid Synthesis Screening that directed federal agencies to condition covered life-sciences funding on the procurement of synthetic nucleic acids from providers or manufacturers adhering to federal screening standards. Recipients were required to publicly attest they comply with the Framework; screen all orders of synthetic nucleic acids; verify all customers ordering “Sequences of Concern”; report potentially illegitimate orders; and have adequate cybersecurity. Because of limits on executive-only action, these restrictions have only been imposed on entities receiving federal funding, leaving other private actors unregulated.⁹²

Before 2024, US oversight of high-consequence biological research operated under multiple policy instruments. The 2012 and 2014 Dual Use Research of Concern (DURC) policies governed research involving select agents with dual-use implications. The 2017 Department of Health and Human Services (HHS) Potential Pandemic Pathogen Care and Oversight Framework (P3CO) required departmental-level review of federally funded research reasonably anticipated to create, transfer, or use pathogens with enhanced pandemic potential (PEPPs). In May 2024, the White House issued the US Government Policy for Oversight of DURC and PEPP, which attempted to harmonize oversight across agencies and clarify ambiguities surrounding gain-of-function⁹³ research. The policy created two new categories of risk and required risk assessments to determine whether research handles sequences of concern or generates pathogens with enhanced pandemic potential.⁹⁴

In May 2025, EO 14292, Improving the Safety and Security of Biological Research, superseded both the Framework for Nucleic Acid Synthesis Screening and DURC policy, requiring both to be revised.⁹⁵ The new EO directed the Office of Science and Technology Policy (OSTP) to strengthen top-down oversight, audits, and enforcement of gain-off-function research.

The EO directed OSTP to develop a strategy addressing oversight gaps for certain non-federally funded research and ordered the termination or suspension of federal funding for “dangerous gain-of-function research.”⁹⁶ This included research conducted in foreign jurisdictions with inadequate oversight, specifically China. OSTP has not yet published the revised policy on biological research or the strategy for oversight of non-federally funded research, leaving significant gaps in biosecurity until both are finalized and promulgated.

91 Exec. Order 14081, Advancing Biotechnology and Biomanufacturing Innovation for a Sustainable, Safe, and Secure American Bioeconomy, 87 Fed. Reg. 56849 (Sept. 15, 2022).

92 Extending these safeguards to all laboratories would likely require federal legislation.

93 Gain-of-function refers to a genetic mutation in an organism that confers a new or enhanced ability, such as transmissibility or virulence. This can occur naturally or be intentionally induced through experimentation. Todd Kuiken, “Oversight of Gain-of-Function Research with Pathogens: Issues for Congress, Congressional Research Service,” Report R47114, July 1, 2025.

94 White House Office of Science and Technology Policy, *United States Government Policy for Oversight of Dual Use Research of Concern and Pathogens with Enhanced Pandemic Potential* (May 6, 2024), <https://aspr.hhs.gov/S3/Documents/USG-Policy-for-Oversight-of-DURC-and-PEPP-May2024-508.pdf>.

95 Exec. Order No. 14292, “Improving the Safety and Security of Biological Research,” May 2025, <https://www.whitehouse.gov/presidential-actions/2025/05/improving-the-safety-and-security-of-biological-research/>.

96 EO 14292 defines “dangerous gain-of-function research” as research that enhances transmissibility or pathogenicity of infectious agents and could reasonably be expected to cause significant societal consequences. The EO enumerated outcome-based criteria for identifying covered work.

Export Control Measures

Export controls have also remained mixed. In January 2025, the Department of Commerce issued the AI Diffusion Rule, which was designed to limit the spread of advanced AI chips and models to certain nations, including through setting export controls on advanced closed-weight AI model weights, and created a three-tiered system for licensing semiconductor exports.⁹⁷ In May 2025, the secretary of commerce announced the department would rescind the proposed rules imposing export controls on certain AI models.⁹⁸ However, while the Trump administration did not let export controls on advanced chips and AI model weights enter into force, it has kept in place Biden-era rules restricting the export of certain DNA synthesis tools and software (discussed further below).⁹⁹

Opportunities for Expanding Federal Incentives

In the absence of comprehensive federal legislation, the government has two immediate levers that do not require new statutory authority: **procurement and financial incentives**. Federal contracts, grants, cloud credits, and research partnerships represent substantial commercial value to frontier AI developers. Conditioning eligibility for those resources on meeting baseline AI-bio safety requirements could create incentives for compliance with voluntary norms. This creates market-wide incentives without waiting for legislation and reaches firms regardless of whether they are subject to state AI laws or federal sector-specific regulation.

Under the current incentive structure, safety investment can be a competitive disadvantage: Firms that invest in biosecurity safeguards bear costs that competitors who do not invest in them avoid. Positive incentives can reverse this dynamic and be a powerful complement to penalizing noncompliance. Tax credits for qualifying safety expenditures, accelerated procurement preference for firms that exceed baseline requirements, liability safe harbor for demonstrated good-faith compliance, and priority access to federal testbeds and compute infrastructure could make safety investment commercially attractive rather than commercially costly. Combined with penalties for noncompliance, these measures could form a federal architecture that shapes the market without mandating a specific governance model. Section 6 addresses how to develop this architecture in partnership with the private sector.

Conclusion

In all, the federal approach to regulating the biological risks of AI models is at best a patchwork. Neither the legislative nor executive branches have comprehensively addressed the risks. However, the executive branch has consistently identified bioconvergence as a foundational national security risk and has sought to leverage executive powers over federal funding, procurement, and export controls to impose requirements for DNA sequence screening and potentially dangerous life science research. These policies nonetheless fall short of responding to the full lifecycle for

97 US Department of Commerce, Framework for Artificial Intelligence Diffusion, 90 Fed. Reg. 4544 (Jan. 15, 2025).

98 US Department of Commerce, “Department of Commerce Announces Rescission of Biden-Era Artificial Intelligence Diffusion Rule, Strengthens Chip-Related Export Controls,” May 13, 2025, <https://www.bis.gov/press-release/department-commerce-announces-rescission-biden-era-artificial-intelligence-diffusion-rule-strengthens>.

99 US Department of Commerce, Controls on Certain Laboratory Equipment and Related Technology to Address Dual Use Concerns About Biotechnology, 90 Fed. Reg. 4612 (Jan. 16, 2025).

threats posed by bioconvergence, particularly by not addressing underlying AI model development, safety, testing, and deployment.

State-Level Innovation in AI Safety

At the state level, emerging legislative activity, particularly in California and New York, has become a significant driver of AI safety practices, including those relevant to biological risk. In the absence of comprehensive federal regulation, states have begun to fill the governance gap with targeted requirements focused on transparency, accountability, and safety testing for high-capability AI systems. These efforts do not yet constitute a unified regulatory regime, and in fact could lead to fragmentation, but consistent with states providing the laboratories for policy,¹⁰⁰ they are creating a credible pathway toward regulation that shapes AI companies' behavior.

California

California has taken the lead in AI regulation, enacting laws that require developers of generative AI systems to implement transparency measures such as digital watermarking of AI-generated content¹⁰¹ and disclosure of training data sources.¹⁰² While these measures are not specific to biosecurity, they establish expectations around traceability and accountability that are directly relevant to managing the risk of bioconvergence. In September 2025, California enacted the Transparency in Frontier Artificial Intelligence Act,¹⁰³ marking the first US law to broadly regulate safety practices for large AI systems. The legislation also marked a major compromise: The California governor previously vetoed a 2024 AI safety bill that would have imposed stricter requirements that opponents felt would stifle innovation.¹⁰⁴

The act centers on risk governance for catastrophic risk where AI contributes to large-scale harm including the creation or release of chemical or biological weapons. Rather than banning models presenting catastrophic risks, the statute requires developers¹⁰⁵ to implement and publish a "frontier AI framework" detailing how they identify, assess, and mitigate these risks. This includes establishing internal thresholds for dangerous capabilities, adopting mitigation strategies, and implementing cybersecurity protections.

The law also imposes transparency and reporting requirements. Developers must publish public transparency reports describing model capabilities, intended uses, and restrictions, and must report "critical safety incidents" to state authorities within twenty-four hours for imminent threats

100 New State Ice Co. v. Liebmann, 285 US 262, 311 (1932) (Brandeis, J., dissenting) ("It is one of the happy incidents of the federal system that a single courageous State may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country.").

101 California Senate Bill 942, California AI Transparency Act (2024).

102 California Assembly Bill 2013, Generative Artificial Intelligence Training Data Transparency Acts (2024).

103 California Senate Bill 53, Transparency in Frontier Artificial Intelligence Act (2025). S.B. 53 applies to "large frontier developers," defined by both a computational threshold and firm size (over \$500 million in annual revenue). This design targets the most capable models while avoiding burdens on smaller firms. Importantly, the law does not explicitly limit its jurisdiction to California-based entities, raising the possibility of extraterritorial application where firms have sufficient nexus to the state.

104 From the private sector, Meta and OpenAI publicly opposed the legislation, while Google, Microsoft, and Anthropic proposed substantial edits that were not taken.

105 The law applies to "large frontier developers," defined by both a computational threshold of models trained above 10^{26} integer or floating-point operations (FLOPS) and firm size (over \$500 million in annual revenue).

and fifteen days for other incidents. The California attorney general can enforce the law through civil penalties of up to \$1 million per violation.

From a biosecurity perspective, California's approach is notable for embedding dual-use risk management into general AI governance rather than treating it as a separate regulatory domain. By requiring firms to internally identify catastrophic risk, including biological misuse, the state essentially converts biosecurity from an externality into a compliance obligation.

New York

In December 2025, New York enacted the Responsible AI Safety and Education (RAISE) Act, which builds on California's model while introducing important differences in scope and liability. Like California's Transparency in Frontier Artificial Intelligence Act, the RAISE Act targets advanced AI systems and emphasizes developer responsibility for managing risks associated with their models.¹⁰⁶ However, it adopts a slightly different framing, focusing on "critical harm" rather than "catastrophic risk."¹⁰⁷ Critical harm includes the creation or use of chemical, biological, radiological, or nuclear weapons. Under the law, before deploying a new frontier model developers must publicly disclose a safety protocol detailing: protections and procedures that reduce the risk of "critical harm"; cybersecurity protections; and testing procedures to evaluate if the model poses an unreasonable risk of critical harm or could be used to create another frontier model in a manner that would increase the risk of critical harm. Developers must also document and retain detailed test results for their frontier models and implement safeguards to prevent unreasonable risk of critical harm.

The Preemption Debate

The proliferation of state-level AI laws has triggered an emerging debate over federal preemption. Industry actors and some policymakers argue that a patchwork of state regulations could impose significant compliance costs, create legal uncertainty, and hinder innovation. Major technology companies, including OpenAI, Google, and Meta, have expressed concerns that divergent state requirements may lead to inconsistent obligations across jurisdictions.¹⁰⁸

In response, several federal proposals have sought to limit or preempt state authority in AI regulation. In 2025, the House of Representatives included in the One Big Beautiful Bill Act a provision imposing a ten-year moratorium on new state AI laws; however, the Senate subsequently voted 99-1 to strip the moratorium from the budget bill. In December 2025, the president issued Executive Order 14365, Ensuring a National Policy Framework for Artificial Intelligence, which seeks to create a "minimally burdensome national policy framework for AI."¹⁰⁹ The White House also formed a Department of Justice task force to challenge state AI laws, mandating a ninety-day

106 The act applies to "large developers," defined as "a person that has trained at least one frontier model and has spent over \$100 million in compute costs in aggregate in training frontier models." It defines "frontier model" as an AI model trained using 10^{26} computational operations or an AI model using knowledge distillation of another frontier model and costing over \$5 million.

107 California's "catastrophic risk" framework emphasizes AI's material contribution to a harm (outcome focused), while New York's "critical harm" standard emphasizes unreasonable risk (risk focused).

108 OpenAI, "Comments to White House Office of Science and Technology Policy on AI Action Plan" (2025) (urging federal preemption of state AI laws to avoid "bogging down innovation"); CIO Dive, "Meta Launches Lobbying Effort to Target State AI Laws," Sept. 24, 2025, <https://www.ciodive.com/news/meta-lobbying-state-ai-laws/761047/> (reporting Meta and Google concerns about patchwork state regulation).

109 Exec. Order No. 14365, "Ensuring a National Policy Framework for Artificial Intelligence," Dec. 11, 2025.

review of state AI laws to identify measures the administration considers onerous, and potentially conditioning federal funding and grants on states refraining from enforcing flagged AI laws.

In March 2026, the White House released its proposed National Policy Framework for AI legislation, which requests Congress “preempt state AI laws that impose undue burdens.” Arguing that states should “not be permitted to regulate AI development” and should “not be permitted to penalize AI developers for a third party’s unlawful conduct involving their models.”¹¹⁰

The preemption debate is especially salient for AI systems with biosecurity implications. The cross-border nature of biological risk and the national security dimensions of AI suggest the need for a unified federal approach. However, in the absence of federal legislation, state laws are currently the primary source of binding safety obligations on AI developers. Congressional preemption of state laws governing the development of AI models could be beneficial if patchwork state legislation is replaced by a unified, federal standard for assessing bio-risks in AI models. Or, in the absence of an alternative federal approach, preemption could increase bioconvergence risks by eliminating the few mandatory AI safety measures that exist in the United States.

International Governance of Bioconvergence Risk

International regulatory frameworks increasingly shape the conduct of US-based AI companies operating globally through binding regulation, biotechnology governance, and multilateral norms.

The EU AI Act

The EU AI Act is the world’s most comprehensive AI regulatory framework to date, reflecting a European governance model centered on precaution, fundamental rights, and risk-based oversight. The act organizes AI systems into four risk categories: “Unacceptable risk” systems, such as social scoring and manipulative AI targeting vulnerable groups, are prohibited outright. “High-risk” systems—including applications in healthcare, employment, education, law enforcement, and critical infrastructure—are permitted only if providers satisfy extensive requirements related to risk management, data quality, documentation, auditability, cybersecurity, and accuracy. “Limited risk” systems are primarily subject to transparency obligations, while “minimal risk” systems remain largely unregulated.

A major development during negotiations was the inclusion of rules governing general-purpose AI (GPAI) models, including foundation models and large language models. GPAI providers must produce technical documentation, provide downstream-use instructions, and comply with EU copyright and transparency obligations. Models classified as posing “systemic risk”—typically those with very high computational power or widespread deployment—face heightened requirements, including adversarial testing, risk evaluations, and incident reporting. Importantly, the act excludes from its scope AI systems used exclusively for military, defense, or national security purposes as well as certain scientific research and development activities.

Although the regulation formally entered into force in August 2024, implementation has lagged. Technical standards necessary for compliance—particularly for high-risk systems—have been

110 The White House, *A National Policy Framework for Artificial Intelligence* (March 2026), <https://www.whitehouse.gov/wp-content/uploads/2026/03/03.20.26-National-Policy-Framework-for-Artificial-Intelligence-Legislative-Recommendations.pdf>.

delayed until at least late 2026 or early 2027, contributing to industry pressure for delayed enforcement.

The EU AI Act does not separately regulate biological AI models. Instead, regulation depends largely on use case. Biological AI integrated into medical devices or regulated products may qualify as “high-risk,” while stand-alone research models may not. Some biological AI systems may also qualify as GPAI models if regulators interpret the definition broadly enough to include models capable of multiple biological tasks within a single sector. In such cases, providers could face obligations related to technical documentation, downstream-use guidance, and training-data transparency. Models designated as posing systemic risk—which could encompass many popular LLMs like Claude, ChatGPT, and Gemini—would face the most stringent requirements.

Overall, the EU AI Act exemplifies the “Brussels Effect,” whereby EU regulations influence global business practices through the EU’s market power and extraterritorial reach.

The EU’s Draft Biotech Act

In December 2025, the European Commission proposed an EU Biotech Act aimed at harmonizing biotechnology regulation, accelerating commercialization, and maintaining protections for health, ethics, environmental safety, and biosecurity. The proposal covers medicines, nucleic acid synthesis devices, biological substances, and AI models used in biological applications.

The draft act establishes a misuse-prevention framework for “biotech products of concern,” including synthetic DNA sequences capable of recreating dangerous pathogens. Suppliers would be required to verify legitimate customer need, conduct know-your-customer screening, and report suspicious transactions. Manufacturers of benchtop nucleic acid synthesis devices—including AI-assisted systems—would also be required to implement sequence-screening safeguards (see the discussion below for an explanation of similar proposals within the United States).

Importantly, the proposal creates a biosecurity-focused advisory group tasked with monitoring biological AI models throughout their lifecycle. Where a model is suspected of posing “biological systemic risk,” the advisory group could issue alerts to EU institutions and, where relevant, escalate concerns through the EU AI Act governance structure, formally linking AI regulation and biosecurity oversight.

Multilateral Norms and Voluntary Frameworks

Nonbinding international frameworks also influence AI governance expectations. The OECD AI Principles, endorsed by the United States and dozens of other countries, emphasize that AI systems should remain robust, secure, and safe throughout their lifecycle and should not create unreasonable safety or security risks. Similarly, the G7 Hiroshima Process Guiding Principles and Voluntary Code of Conduct encourage developers to adopt risk-mitigation measures and pay particular attention to chemical, biological, and nuclear risks. In the realm of biosecurity, the International Biosecurity and Biosafety Initiative for Science (IBBIS) has produced a common mechanism for nucleic acid synthesis screening as an important global, public/private voluntary effort to enhance adherence to DNA synthesis screening guidelines (discussed further below).¹¹¹

¹¹¹ International Biosecurity and Biosafety Initiative for Science, “Common Mechanism,” accessed July 5, 2026, <https://ibbis.bio/our-work/common-mechanism/>.

The International Bio Funders Compact also represents an international, voluntary agreement among life sciences funders to assess biosafety and biosecurity risks and address them as part of R&D funding decisions.¹¹²

Each of these remains voluntary and lacks enforcement mechanisms.

Private Sector Self-Regulation and Incentive Structures

Public-private partnership is essential to effective deterrence, especially prevention. Today, US frontier AI companies have committed to biosecurity safeguards, but these are voluntary and based on self-assessments. As a result, they are inconsistently applied and subject to revision without external oversight. This approach is insufficient to prevent the risks posed by bioconvergence discussed in Chapter 2. The government must partner with industry to create enforceable accountability mechanisms that have independent verification. Further, there are legitimate concerns that mandating requirements may stifle private sector innovation and disadvantage the US from a national security perspective. Policy makers need to take those concerns into consideration in formulating potential regulations and identify incentives that make safety investments beneficial for the relevant sectors.

Voluntary Regimes

US frontier AI companies have all publicly committed themselves to safety under voluntary regimes. Their policies have transformed from ad hoc trust-and-safety practices toward more formalized lifecycle safety regimes, although significant challenges remain. In 2023, fifteen leading technology firms endorsed the White House's voluntary commitments on the safe, secure, and transparent development and use of generative AI model technology.¹¹³ Those overarching principles address prerelease safety testing (including red teaming), information-sharing on risk and vulnerabilities within the industry, investment in cybersecurity, content provenance, and publicly reporting on model capabilities and limitations.

Leading AI companies have also adopted safety frameworks that establish internal thresholds for model capability and risk. While each company's internal safety policy differs in branding and terminology, they share several components of voluntary safeguards:¹¹⁴ These frameworks are designed to determine whether and under what conditions a model should be trained, scaled, or deployed. In 2023, Anthropic released its first Responsible Scaling Policy (RSP), which defined capability levels and prescribed additional safety and security controls as models approach more powerful thresholds.¹¹⁵ OpenAI followed suit later that year with its Preparedness Framework, which described risk categories (e.g., chemical, biological, cyber, autonomous replication) and linked

112 Nuclear Threat Initiative, "Bio Funders Compact," accessed July 5, 2026, <https://www.nti.org/about/programs-projects/project/bio-funders-compact/>.

113 The White House, "Voluntary AI Commitments," September 2023, <https://bidenwhitehouse.archives.gov/wp-content/uploads/2023/09/Voluntary-AI-Commitments-September-2023.pdf>. Endorsing companies include Amazon, Anthropic, Google, Inflection AI, Meta, Microsoft, OpenAI, Adobe, Cohere, IBM, Nvidia, Palantir, Salesforce, Scale AI, and Stability AI.

114 For a more comprehensive comparison of safety policies, see METR, "Common Elements of Frontier AI Safety Policies," December 2025, <https://metr.org/common-elements.pdf>; Frontier Model Forum, "Issue Brief: Components of Frontier AI Safety Frameworks," Nov. 8, 2024, <https://www.frontiermodelforum.org/updates/issue-brief-components-of-frontier-ai-safety-frameworks/>.

115 Anthropic, "Anthropic's Responsible Scaling Policy, Version 1.0," Sept. 19, 2023, <https://www-cdn.anthropic.com/1adf000c8f675958c2ee23805d91aaade1cd4613/responsible-scaling-policy.pdf>.

them to internal safeguards and escalation procedures.¹¹⁶ Similarly, Google DeepMind's Frontier Safety Framework is structured around capability evaluations and mitigation requirements.¹¹⁷

Despite the presence of voluntary safety frameworks among the major frontier AI companies, significant gaps remain. An Epoch AI review of approximately 370 biological foundation models found that fewer than 3 percent of the models' accompanying publications documented safeguards such as evaluations, access controls, or filtering or exclusion of training data.¹¹⁸ Companies evaluate their own compliance without independent validation. There is no external audit requirement, no government verification of threshold determinations, and no consequence for rolling back commitments. Existing safety standards, as limited as they are, are being weakened: For example, in early 2026, Anthropic's RSP Version 3.0 replaced pauses or stops with reporting and roadmap-based risk mitigation.¹¹⁹

Other voluntary initiatives have focused specifically on AI-enabled biology. For example, in May 2026 OpenAI announced GPT-Rosalind, a frontier AI model tailored for life sciences, drug discovery, and translational medicine.¹²⁰ To balance scientific acceleration with national security, OpenAI has committed to restrict public access to the model, and instead to sponsor its use for vetted biodefense organizations, the US government, and allied governments. This is paired with Rosalind Biodefense, a new initiative to enable the development of high-impact defensive applications of AI in the life sciences. The test for Rosalind's impact will be if it genuinely enhances public-private scientific advancement and helps trusted labs generate novel countermeasures more quickly, and if its restrictions prevent sophisticated AI models and agents from enabling pernicious bio activities. Importantly, this voluntary commitment by one company would be significantly strengthened if it were to be extended across AI labs. For example, the existing Frontier Model Forum could be used to coordinate a multicompany trusted access program for the most sophisticated biological models.

More generally, the private sector's implementation of safety measures varies from company to company in design and stringency, although they increasingly share a common architecture: (a) internal governance; (b) structured capability evaluations; (c) alignment techniques; (d) controlled deployment; and (e) post-release monitoring. Models are supposed to be evaluated pre-deployment against high-risk domains including biological and chemical weapons to reduce harmful outputs. To achieve this goal, AI developers use alignment techniques to shape model behavior to make it less likely to produce outputs that are dangerous or violate policy.

116 OpenAI, "Preparedness Framework (Beta)," Dec. 18, 2023, <https://cdn.openai.com/openai-preparedness-framework-beta.pdf>; OpenAI, "Our Updated Preparedness Framework," Apr. 15, 2025, <https://openai.com/index/updating-our-preparedness-framework/>.

117 Google DeepMind, "Frontier Safety Framework, Version 3.0," Sept. 22, 2025, https://storage.googleapis.com/deepmind-media/DeepMind.com/Blog/strengthening-our-frontier-safety-framework/frontier-safety-framework_3.pdf.

118 Villalobos and Atanasov, "Expanded Biology AI Coverage."

119 Anthropic, "Responsible Scaling Policy, Version 3.0," Feb. 23, 2026, <http://www-cdn.anthropic.com/e670587677525f28df69b59e5fb4c22cc5461a17.pdf>; Billy Perrigo, "Exclusive: Anthropic Drops Flagship Safety Pledge," *Time*, Feb. 24, 2026, <https://time.com/7380854/exclusive-anthropic-drops-flagship-safety-pledge/>.

120 OpenAI, "Strengthening Societal Resilience with Rosalind Biodefense," May 29, 2026, <https://openai.com/index/strengthening-societal-resilience-with-rosalind-biodefense/>.

Model Alignment Techniques

In the biosecurity context, the goal of model alignment is to reduce the likelihood that a model will provide detailed technical assistance with pathogen enhancement, synthesis routes, or other high-risk biological content, even when prompted by a user to do so.

The core alignment techniques used are:

- Data governance and filtering: Developers curate pretraining and fine-tuning datasets to exclude categories of illegal, privacy-violating, or extreme content.
- Post-training alignment: Instruction tuning, supervised fine-tuning (SFT), and reinforcement learning from human feedback (RLHF) are widely used to reduce harmful outputs.
- Red-teaming, or adversarial testing: Companies conduct internal and external red-teaming in biosecurity, cybersecurity, and other safety areas. Several companies, including OpenAI and Microsoft, have published public guidance on their red-teaming practices in addition to their frontier safety frameworks.

Deployment of sensitive information in models is controlled through tiered access and content filtering. And finally, models are monitored after their release.

Structured Capability Evaluations

Assessing when a model possesses or is approaching dangerous capabilities is essential to managing the model's risk. Most private firms have focused evaluations on high-risk domains: biological and chemical weapons, cyber offense, autonomous replication, and model deception or manipulation. METR identified “[b]iological weapons assistance” as one of the three most common capabilities addressed by AI companies.¹²¹ Company policies link high-risk domains to capability thresholds that trigger escalating safeguards. Models are examined using empirical measurements against benchmarks, which informs escalating requirements like additional model training, deployment restrictions, model access limitations, elevated internal reviews, and pauses in model development.

Beyond the model itself, AI companies often implement safety measures of the AI product with which users interface. This includes layering content filters and classifiers on top of models to prevent delivering certain responses, tiered access for vetted users, and model release to “sandboxes.”

Frontier AI companies also monitor their models after release. Particularly for LLMs, firms have enabled real-time abuse monitoring, policy updates to reduce jailbreak techniques, and reporting channels for flagged queries and content. In many ways, these efforts resemble traditional cybersecurity rather than software quality assurance by requiring constant mitigation based on adversary behavior.

¹²¹ METR, Common Elements of Frontier AI Safety Policies, 8.

Incentives for Self-Regulation

AI model safety architecture is neither standardized nor legally harmonized. It remains, to a significant extent, self-regulated. Even where technology companies largely self-regulate, public policy should focus on incentivizing companies to prioritize safety, particularly related to bioconvergence risks.

There are structural incentives to motivate self-regulation, and government policies can further influence these in several ways.

- 1. Voluntary commitments can push the adoption of market-wide best practices.**

In 2023 and 2024, leading AI developers signed the White House's voluntary safety commitments and formed their own 501(c)(3), the Frontier Model Forum, to identify industry best practices for significant risks to public safety and security. The US supported this by using NIST to develop voluntary safety standards and risk-management frameworks for AI systems.¹²²

- 2. Market incentives can drive safety.** Responsible AI governance is increasingly framed as a competitive advantage. Companies that demonstrate ethical, transparent, and risk-aware approaches may attract more customers and investment.¹²³ Trust sustains long-term adoption and profitability. In sectors where AI products directly intersect with public safety, such as healthcare, safety practices can become a deciding factor in procurement and deployment decisions, acting as a market incentive for companies.

- 3. Where market forces alone are insufficient to address public risks, government actions can help shape firms' behavior.** Such actions can include: tax incentives, procurement requirements for safety and beneficial treatment for those companies that adopt best practices; leveraging government financing to require safety provisions; liability protection and safe harbor for firms that take adequate safety measures; support for insurance markets to write down risks for companies that add safety measures; and export controls and foreign investment screening for certain products. Neither government action nor self-regulation alone will ensure security and safety for the convergence of AI and biology.

Lessons for Policymakers

The convergence of AI and biotechnology has fundamentally altered the prevention challenge by democratizing dangerous capabilities, compressing timelines from design to deployment, and creating novel threats that bypass existing controls. The current prevention architecture is insufficient. The recommendations below propose a prevention architecture built on four principles:

- **Capability-triggered obligations.** Safety and security requirements must scale with AI system capabilities and biological risk. Frontier models crossing defined thresholds for accessing biological design, synthesis support, autonomous

¹²² National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), Jan. 26, 2023, <https://www.nist.gov/itl/ai-risk-management-framework>.

¹²³ Kevin A. Bryan and Florenta Teodoridis, "Balancing Market Innovation Incentives and Regulation in AI: Challenges and Opportunities," Brookings Institution, Sept. 24, 2024, <https://www.brookings.edu/articles/balancing-market-innovation-incentives-and-regulation-in-ai-challenges-and-opportunities/>.

laboratory capabilities, and/or dual-use research potential must trigger mandatory safety, testing, and deployment obligations.

- **Government-industry partnership.** Neither regulation nor self-governance alone will succeed. Federal and state policy should create strong incentives for safety investment, including procurement preference, tax credits, liability safe harbors, incentives for producing safer and more secure tools, and research partnerships. The government should concurrently establish mandatory safety and security baselines and independent verification in those areas where self-regulation fails or introduces unacceptable risk.
- **Function over form.** Prevention measures must identify dangerous capabilities rather than rely on static lists. Function-based DNA synthesis screening, capability evaluations for AI models, and risk-based biological data classification can address threats that list-based systems miss and are more robust as capabilities evolve.
- **Preservation of beneficial innovation.** The same AI and biological capabilities that enable misuse also accelerate drug discovery, vaccine development, pathogen detection, and advances in agriculture, health, and climate. Prevention policy must channel innovation toward beneficial applications rather than constrain it categorically.

The window for prevention is narrowing. AI biological capabilities are advancing faster than governance structures. Each month without definitive action allows capabilities to diffuse more widely, making retroactive control harder and catastrophic misuse more likely. Leadership requires moving beyond voluntary frameworks toward enforceable standards. At the same time, the government cannot keep pace with rapidly changing technology to prevent harmful uses without close collaboration with industry.

Section A Findings and Recommendations

Finding 3.1.1: The United States lacks comprehensive federal oversight for AI systems with biological design capabilities. State-level regulations create compliance fragmentation. Current federal authorities do not adequately address AI-enabled biological threats.

Recommendation 3.1.1.1: Adopt a federal, capability-triggered AI-bio regime. Impose binding obligations on frontier AI developers whose models cross defined capability and compute thresholds relevant to biological design, synthesis, or dual-use research. Require pre-deployment biosecurity evaluations and post-deployment monitoring. Require model development companies to test qualifying models before release for biological and chemical misuse potential, document results, retain records, and retest after major updates.

Recommendation 3.1.1.2: Mandate public safety frameworks. Require AI companies to publish a frontier safety protocol describing dangerous-capability thresholds, mitigation measures, intended-use restrictions, and cybersecurity controls. Require companies to report critical safety incidents on short timelines.

Recommendation 3.1.1.3: Tie federal money to compliance with safety requirements. Make eligibility for federal contracts, grants, cloud credits, and research partnerships contingent on meeting baseline AI-bio safety requirements.

Finding 3.1.2: The US government lacks a central entity empowered to regulate AI safety. NIST is empowered to issue AI safety standards but only through voluntary standards, and does not have authority to promulgate binding regulations or enforce them. Within NIST, the Center for AI Standards and Innovation (CAISI) has the capacity to conduct national security assessments of advanced AI models but only on a voluntary basis.

Recommendation 3.1.2.1: Congress should authorize a central federal government entity empowered to issue regulations on safety testing for advanced AI models. Such regulatory authority must include the power to:

- Define covered models;
- Define risk thresholds;
- Evaluate frontier models;
- Conduct red-teaming of AI models;
- Review safety cases and incidents;
- Review risks to third parties; and
- Review the internal governance processes of AI labs.

Recommendation 3.1.2.2: Congress should authorize the central federal government entity to promulgate binding regulations specific to biological threats and AI models. This should include the regulation of AI models falling below defined compute thresholds for general purpose AI models but exceeding capabilities specific to biology or having been trained on biological data. In developing the requirements, the entity should consult with biology and public health experts at HHS, DOE national laboratories, and DOD. These bio evaluations should assess whether a model can:

- Generate actionable protocols for pathogen synthesis or modification;
- Assist nonexperts in overcoming tacit biological knowledge barriers;
- Enable optimization of delivery mechanisms, virulence, or immune evasion;
- Circumvent existing biosafety and biosecurity controls.

Evaluations should be independently reproducible and subject to audit by the designated federal entity. This entity would also work closely with the National Institute of Bioresilience, discussed in Chapter 7.

Recommendation 3.1.2.3: Congress should authorize CAISI to undertake the actions outlined in Recommendations 3.1.2.1 and 3.1.2.2.

Recommendation 3.1.2.4: Congress should create a mandatory AI-bio incident reporting regime. Covered AI developers, cloud providers, and relevant research institutions should be required to report material biological safety incidents, including:

- Demonstrated dangerous model capabilities;
- Unauthorized access to sensitive models;
- Evidence of malicious biological misuse;
- Failures of safeguards or containment mechanisms;
- Confidential risk information sharing.

The designated federal entity should maintain a secure incident database and produce anonymized lessons-learned reports to Congress.

Recommendation 3.1.2.5: Congress should create enforcement power along with compliance incentives. The designated federal entity should be authorized to levy civil fines for failure to comply and to make referrals to other federal regulators, such as the FTC.

Finding 3.1.3: Any AI-specific entity in the US government must be able to compete with tech companies to attract top talent and coordinate quickly with relevant government experts.

Recommendation 3.1.3.1: Congress should authorize the designated federal entity for AI safety to have excepted-service hiring authority for faster hiring and individualized compensation. This could be modeled after DARPA hiring authorities, DOE national laboratory appointments, NSF rotator programs, and US Digital Service hiring authorities.

Recommendation 3.1.3.2: Congress should establish an AI Safety Fusion Cell within CAISI. This would enable rapid information-sharing and colocation of national security and AI experts from Commerce, DOE, DOD, HHS, and the intelligence community.

Finding 3.1.4: AI-enabled biotechnology safety can be encouraged through federal incentives. Currently, contractors and grantees have no biosecurity obligations tied to receipt of federal resources. Conditioning federal funding on biosecurity compliance converts an externality into a cost firms must manage, creating market-wide incentives without waiting for legislation. Under current incentive structures, safety investment is a competitive disadvantage rather than a commercial advantage. Firms that invest in biosecurity safeguards bear costs avoided by competitors who do not, creating a race to the bottom. Positive incentives can reverse this dynamic without mandates.

Recommendation 3.1.4.1: Make eligibility for federal contracts, grants, cloud credits, and research partnerships contingent on meeting baseline AI-bio safety requirements.

Recommendation 3.1.4.2: Create positive incentives for safety investment. Examples could include: tax credits, accelerated procurement preference, liability safe harbors, and priority access to federal testbeds. These incentives would be available for firms that demonstrate they exceed baseline requirements through independent audits, robust red-teaming, and integration with synthesis screening safeguards.

Finding 3.1.5: Federal preemption of state AI laws could eliminate existing safety requirements without replacing them. State laws currently provide the only mandatory AI safety obligations in the United States. Preemption without replacement would widen rather than close governance gaps.

Recommendation 3.1.5.1: Congress should not preempt state AI safety laws unless it simultaneously installs a federal standard at least as protective on testing, transparency, and incident reporting related to biological risks as laws passed in California and New York.

Finding 3.1.6: International regulatory frameworks are increasingly shaping global AI-bio governance, but the United States lacks a coordinated strategy to respond to extraterritorial regulation and emerging international biosecurity standards. The European Union has established the most comprehensive international framework for regulating AI systems through the EU AI Act and the proposed EU Biotech Act. Because these regimes apply extraterritorially to companies operating within European markets, US-based AI firms developing biological AI capabilities may effectively become subject to foreign regulatory standards even in the absence of comparable domestic oversight.

Recommendation 3.1.6.1: Harmonize with international standards, where feasible. US rules should be interoperable with the EU-style approach to risk management, documentation, testing, and incident reporting, so firms are not forced to build separate compliance systems for different jurisdictions. The executive branch should create an interagency coordination body responsible for tracking international AI and biotechnology regulation, assessing implications for US competitiveness and national security, and coordinating engagement with allies on AI-bio governance issues. The mechanism should include participation from agencies responsible for commerce, national security, public health, biotechnology, and AI policy.

Recommendation 3.1.6.2: Promote international harmonization of AI-biosecurity safeguards through multilateral engagement. The United States should work through forums such as the OECD, G7, and other allied institutions to promote common standards for evaluating biological misuse risks in advanced AI systems. These efforts should prioritize shared approaches to model evaluations, incident reporting, biosecurity safeguards, and oversight of high-risk biological AI applications while reducing regulatory fragmentation across jurisdictions.

B. Biological Data Governance

Biological training data are the foundational input to AI-enabled biotechnology systems. The capabilities of biological AI models scale directly with the volume, quality, and diversity of the data on which they are trained. As a result, governing highly capable AI-enabled biotechnology systems requires governing those training datasets. Further, the relationship between specific data types and dual-use risk is not fully characterized, a challenge that argues for precautionary governance rather than deferred action.¹²⁴ The United States currently has no federal framework for governing biological AI models or evaluating dual-use risks associated with biological training data. This is a significant governance gap.

124 Ibid.

Biological Data

Biological data are information derived from the structure, function, or process of biological systems that is measured, collected, or aggregated for analysis. Biological data can be organized into three categories relevant for training AI models: data on DNA and protein **sequences**, data on protein **structures**, and data on the **function** of sequences and structures. Data in each of these categories are dynamic and constantly evolving.

Biological data used to train AI systems encompass three categories: data on DNA and protein sequences; data on protein structures (experimentally derived or AI-predicted); and data on the functional consequences of particular DNA sequences and protein structures.¹²⁵ Unlike AI systems trained in fields with established rules, like mathematics or logic, biology is an evolving field. As a result, data volume and quality are highly determinative of model capability.¹²⁶

For example, AlphaFold needed to be trained on the structures of roughly 150,000 proteins to be able to predict the structures of hundreds of millions of proteins.¹²⁷ Experts have identified both the lack of large volumes of biological datasets and the lack of high-quality datasets (i.e., data that are clean and accurately annotated) as a major bottleneck in projected improvements to AI biological tools in the near future.¹²⁸ And given this gap, AI firms are increasingly turning to synthetic data¹²⁹ —AI-generated biological data that replicate real-world datasets—to train their models.

The biological data landscape is not only complex, it is highly fragmented. More than 1,700 public and proprietary databases existed as of 2019, with heterogeneous data types, management structures, and access protocols.¹³⁰ Most pathogen sequence and function data are held by a small number of laboratories under varying access controls. This fragmentation both constrains beneficial research and creates inconsistent barriers to misuse. The Protein Data Bank (PDB), one of the few standardized, curated, publicly available repositories and the dataset on which AlphaFold was trained, is among the exceptions. A summary of the major biological datasets is provided in Table 3.1 below.

125 Allison Berke et al., *Data and AI-Enabled Biological Design* (June 2025), <https://www.rand.org/pubs/perspectives/PEA3886-1.html>.

126 National Academies of Sciences, Engineering, and Medicine, “Importance of Data in AI-Enabled Biological Models,” in *The Age of AI in the Life Sciences: Benefits and Biosecurity Consideration* (National Academies Press, 2025), <https://www.ncbi.nlm.nih.gov/books/NBK614603/>.

127 “Unlocking Life’s Building Blocks,” The Atlantic, <https://www.theatlantic.com/sponsored/google/unlocking-lifes-building-blocks-demis-hassabis/3867/>.

128 Nikki Teran and Jaime Yassif, “AlxBio Horizon Scan Winter 2025-2026,” Nuclear Threat Initiative, Mar. 3, 2026, <https://www.nti.org/analysis/articles/AlxBio-horizon-scan-winter-2025-2026/>.

129 Yun Gyeong Lee et al., “Synthetic Data Production for Biomedical Research,” *Osong Public Health and Research Perspectives* 16 (2025): 94–103.

130 Arlindo Oliveira, “Biotechnology, Big Data and Artificial Intelligence,” *Biotechnology Journal* 14 (2019), <https://doi.org/10.1002/biot.201800613>.

Table 3.1 Major Biological Datasets.

Data Type	Example Datasets	Methods of Access
Sequence data	GenBank (NCBI), European Nucleotide Archive, UniProt	Mostly public, open access
Structural data	Protein Data Bank, AlphaFold database	Fully public, open access
Functional data	KEGG pathway database	Fully public, open access
Human genomic/clinical data	NIH Database of Genotypes and Phenotypes (dbGaP), UK Biobank	Mostly controlled access (tiered, risk-contingent)
Electronic health records	Hospital databases, private data aggregators	Controlled (data use agreements, HIPAA-compliant sharing)
Biomedical literature	PubMed, academic journals	Mixed (subscription/license-based)

The disparate biological datasets are all subject to three trends that increase the need for an active federal response:

Commodification: Biological datasets are increasingly valued, traded, and licensed as proprietary AI training inputs, heightening the risk of high-value datasets reaching adversarial state actors without oversight.¹³¹

Fragmented tiering: The dual-use risks of biological datasets are driving a shift from open-access norms toward risk-contingent access frameworks. Without harmonization, this has the potential to both over-restrict low-risk data and under-restrict high-risk datasets.¹³²

Strategic asset recognition: Biological datasets are becoming increasingly understood as strategic assets, both in the United States and abroad. China’s National Genebank, which aggregates biological data with a government-controlled synthesis and analytics platform, is the most advanced example of treating biological datasets as national security infrastructure.¹³³

The increasing dual-use risks of biological datasets mean that more restrictive and more consistent access controls need to replace the open access norms prevalent today. AI models have made potential use cases for certain categories of biological data so dangerous that the US government has issued export controls on high-end cytometry and mass spectrometry equipment precisely to reduce adversary countries’ access to the datasets they can generate.¹³⁴ The goal is to maintain a competitive edge while raising the obstacles to misuse, while preserving the benefits of open science. Because data access is often irreversible once released, new and especially powerful datasets need to be adequately controlled before they diffuse widely. Recognizing the importance of these issues, the NSCEB has recommended that Congress establish a central office to manage

131 Nancy Connell et al., “Driving Responsible Innovation of AI, Life Sciences and Next Generation Biotech,” IEEE Standards Association, Jan. 3, 2021, <https://standards.ieee.org/beyond-standards/driving-responsible-innovation-of-ai-life-sciences-and-next-generation-biotech/>.

132 Bloomfield et al., “Securing Dual-Use Pathogen Data.”

133 Bo Wang et al., “The China National GeneBank—Owned by All, Completed by All and Shared by All,” *Yi Chuan* 41 (2019), <https://doi.org/10.16288/j.ycz.19-148>.

134 US Department of Commerce, Bureau of Industry and Security, “Controls on Certain Laboratory Equipment and Related Technology to Address Dual Use Concerns About Biotechnology,” Federal Register 90, no. 10 (Jan. 16, 2025): 4612–17, <https://www.govinfo.gov/content/pkg/FR-2025-01-16/pdf/2025-00723.pdf>.

and secure biological data.¹³⁵ And in June 2026, members of Congress in both chambers introduced legislation that would designate biological data a strategic national resource.¹³⁶

Tiered access frameworks based on risk, such as that used by NIH in its dbGaP, are a good model. Data are categorized by sensitivity. NIH requires restricted-use agreements and Data Access Committee approval for higher-risk datasets and limits users' ability to download the highest-risk data, which can be used only in cloud environments. While further research is needed to determine whether certain types of biological data are more likely to enable high-risk applications,¹³⁷ the close relation between the capabilities of AI biological models and their underlying training data makes policies around biological data central to leveraging the benefits while mitigating the risks from AI models in this field. Without robust efforts to harmonize access restrictions, ad hoc implementation of data access will further fragment the biological data landscape and increase the risk of misuse.

Sound biological data governance is also critical for innovation. The same fragmentation that creates inconsistent barriers to misuse also limits beneficial research applications that depend on biological data. A federal framework that creates clarity, risk-tiering, and trusted access pathways can support the continued growth in biotechnology research. This is the virtuous cycle that prevention policy must enable. Effectively governing what is dangerous makes it possible to promote what is beneficial.

Section B Findings and Recommendations

Finding 3.2.1: There is no government-wide framework to evaluate dual-use risks associated with biological training data. The capabilities of biological AI models depend directly on training data. Current data governance lacks risk assessment for dual-use applications. Commodification and restricted access trends require coordinated federal response.

Recommendation 3.2.1.1: NIH should expand its tiered access model for the Database of Genotypes and Phenotypes to cover a broader class of biological datasets involved in AI training.

Recommendation 3.2.1.2: Congress should require developers of frontier AI models to disclose categories of biological datasets used in training.

Recommendation 3.2.1.3: Congress should establish a central office within the US government responsible for creating harmonized standards for biological datasets, and storing and managing access to biological data. This should include coverage of synthetic biological datasets used in AI model training.

Recommendation 3.2.1.4: Establish standardized national risk-tiering and access-control protocols for biological datasets. The federal government should develop standardized criteria

135 National Security Commission on Emerging Biotechnology [NSCEB], *Charting the Future of Biotechnology* (April 2025), <https://www.biotech.senate.gov/wp-content/uploads/2025/10/NSCEB-%E2%80%93-Full-Report-%E2%80%93-Sep-30-.25.pdf>.

136 NSCEB, "Congress Takes Steps to Make Biological Data a Strategic National Resource," press release, June 11, 2026, <https://www.biotech.senate.gov/press-releases/congress-takes-steps-to-make-biological-data-a-strategic-national-resource/>.

137 For instance, studies suggest data on human-infecting pathogens significantly contributes to creating AI models that can design pathogens with enhanced pandemic potential. Doni Bloomfield et al., "Securing Dual-Use Pathogen Data of Concern," Proceedings of the NeurIPS Workshop on Biosecurity Safeguards for Generative AI (2026).

for categorizing biological datasets according to dual-use risk potential. NIH, NIST, and relevant national security agencies should jointly establish interoperable access-control standards across federally supported biological databases. Federal agencies should require standardized auditing, credentialing, and monitoring procedures for access to high-risk biological datasets.

Recommendation 3.2.1.5: Create trusted-access pathways for institutions that preserve beneficial biotechnology research. Federal agencies should establish expedited access mechanisms for accredited academic, medical, and public-interest researchers working with higher-risk biological datasets. NIH should develop standardized trusted-researcher credentialing systems for controlled-access biological data environments.

Finding 3.2.2: Biological datasets have become strategic national security assets, and access to high-risk biological datasets by adversary nations is a risk to US national security.

Recommendation 3.2.2: Congress should authorize federal review mechanisms for foreign access to high-risk biological datasets with national security implications. The Department of Commerce should evaluate whether certain categories of biological training data warrant export-control restrictions.

C. Governance of Biological AI Tools

As AI has evolved, the range of biological AI tools has expanded. Today, there are three major types:

- **Biological design tools (BDTs)** directly generate or evaluate DNA, RNA, and protein sequences and structures, and predict the function of genetic variants. Examples of BDTs include Evo2, a genomic foundation model whose largest version was trained on approximately 9.3 trillion single-nucleotide tokens drawn from genomic sequence data.¹³⁸ Evo2 can generate novel DNA and RNA sequences and predict the functional consequences of genetic variants.¹³⁹
- **Natural-language interfaces** for accessing biological design capabilities. These models, like MP4 and Pinal, accept plain English prompting for protein generation functions.
- **Robotic systems** to automate wet-lab work. These tools, such as OpenTrons, close the loop between digital design and physical synthesis.

As agentic AI systems are incorporated across all three categories, the capacity for autonomous biological work will increase.

138 Garyk Brixi et al., “Genome Modelling and Design Across All Domains of Life with Evo 2,” *Nature* 652 (2026): 1349–61, <https://doi.org/10.1038/s41586-026-10176-5>.

139 Other examples include AlphaProteo, AlphaFold 3, OpenFold 3, and Boltz-2. Andrii Markov, “Protein Engineering with AI: OpenFold3 vs. Boltz 2 vs. AlphaFold 3,” Blackthorn.ai, Sept. 29, 2025, <https://blackthorn.ai/blog/protein-engineering-with-ai/>; Teran and Yassif, “AlxBio Horizon Scan.”

Access Regimes

Historically, access to AI-enabled biological tools, like biological data, has been managed through voluntary corporate policies operating under academic open-source norms. Companies manage access to their AI biological models along a spectrum from fully open to fully closed. Increasingly, they are using tiered access gateways, but pressure by the research community has led to a weakening of some protections.¹⁴⁰ For example, Google DeepMind initially released its AlphaFold 3 protein structure prediction model via an application programming interface (API)—a web-based server for noncommercial researchers that restricts access to the underlying model, preventing users from running or adapting the model independently, rather than just querying it. In response to significant criticism by researchers, Google DeepMind released AlphaFold 3's inference code and made its model parameters available upon request for eligible noncommercial research, enabling qualified researchers to run the model independently under specified terms.¹⁴¹ Now, DeepMind provides general users with access to a version of AlphaFold 3 via an API but grants greater access to those users who demonstrate an academic or other noncommercial research affiliation.¹⁴² Similarly, EvolutionaryScale offers its ESM3 biological foundation model through a commercial API but provides a smaller, less capable open version (ESM3-open) for academic research.¹⁴³

For open-source models, attempts to build safeguards have proven inadequate. For instance, the creators of the Evo1 foundation model deliberately excluded some virus information from its training data to prevent the model from generating dangerous viral sequences.¹⁴⁴ But because the model's weights were published openly, researchers were able to bypass the safeguard by fine-tuning the model with withheld viral data, significantly increasing the model's ability to predict characteristics such as virulence and immune escape.¹⁴⁵ Despite academic community pressure for open-source research, some scientists have called for a better system of managed access to certain BDTs.

Information access is also a distinct difference from often-repeated analogies between AI and nuclear capabilities. In the nuclear context, for example, certain nuclear information is classified by operation of law from the moment it comes into existence.¹⁴⁶ But in the biosecurity space, information about and access to bio-LLMs, BDTs, autonomous labs, and convergence capabilities are not covered by legal restrictions and the data and information they generate are not “born” classified.

140 Sarah Carter et al., “Developing Guardrails for AI Biodesign Tools,” Nuclear Threat Initiative, November 2024, https://www.nti.org/wp-content/uploads/2024/11/NTIBio_Paper_Developing-Guardrails-for-AI-Biodesign-Tools_FINAL.pdf.

141 Ibid.

142 Callaway, “AI Protein-Prediction Tool AlphaFold3.”

143 EvolutionaryScale, “ESM3.”

144 Georgia Adamson and Gregory C. Allen, “Opportunities to Strengthen US Biosecurity from AI-Enabled Bioterrorism,” Center for Strategic and International Studies, August 2025, https://csis-website-prod.s3.amazonaws.com/s3fs-public/2025-08/250806_Adamson_AI-Enabled_Bioterrorism.pdf?VersionId=vS8T0bGMLr_IU7RTcXs8im_ndJuGz0cM.

145 Bloomfield et al., “Securing Dual-Use Pathogen Data.”

146 42 U.S.C. § 2014(y) (automatically classifying information as “Restricted Data” all data concerning (1) design, manufacture, or utilization of atomic weapons; (2) the production of special nuclear material; or (3) the use of special nuclear material in the production of energy).

Recent Research on BDT Access Restrictions

NTI Report. The issue of access to BDTs was addressed at length in a January 2026 report by the Nuclear Threat Initiative (NTI). The report lays out a detailed framework for designing a tiered access regime that balances the risk of particular BDTs against the desire for equitable access for academic researchers.¹⁴⁷ The NTI framework specifies four escalating risk levels based on the type of biological AI tool. Access requirements to the tools depends on whether they are accessed through an API or through a full model with its underlying code, training data, and model weights. At every level the framework identifies what type of user, institutional, and project verification is necessary.

The NIH's dbGaP model is another excellent example of tiered access controls. Data are categorized by sensitivity. NIH requires restricted-use agreements and Data Access Committee approval for higher-risk datasets and limits users' ability to download the highest-risk data, which can be used only in cloud environments. While further research is needed to determine whether certain types of biological data are more likely to enable high-risk applications,¹⁴⁸ the close relation between the capabilities of AI biological models and their underlying training data makes policies around biological data central to leveraging the benefits while mitigating the risks from AI models in this field.

API access provides some additional protections. When users access a tool via API, the host can monitor BDT use in real time and exercise oversight. Experts told the NTI report authors that even "the number or pattern of queries" by users could provide helpful red flags for the API hosts.¹⁴⁹ Upon detection of high-risk use of the system, the API host could require additional identity verification, restrict or refuse outputs, trigger human review, or even notify law enforcement.

Despite the utility of a tiered access approach, BDT developers lack ready access to resources that would help them create, implement, and maintain their own managed access platforms, especially with sufficient monitoring safeguards in place.¹⁵⁰ As a result, this is an area where the federal government can provide powerful assistance, potentially by providing a centralized API host for BDTs and collaborating with researchers on independent monitoring.

Science. A February 2026 proposal by a multi-institution group of university-based experts in *Science* magazine offers another detailed framework for restricting access to a small subset of newly generated, high-risk biological datasets, especially those that enable prediction or design of pathogen traits like transmissibility or immune evasion.¹⁵¹ The approach borrows from privacy governance, where sensitive datasets are widely used for research but only under structured access regimes.

147 Sarah Carter and Greg Butchello, "A Framework for Managed Access to Biological AI Tools," Nuclear Threat Initiative, January 2026, https://www.nti.org/wp-content/uploads/2026/01/NTIBio_Paper_A-Framework-for-Managed-Access-to-Biological-AI-Tools_FINAL.pdf.

148 For instance, studies suggest data on human-infecting pathogens significantly contributes to creating AI models that can design pathogens with enhanced pandemic potential. Doni Bloomfield et al., "Securing Dual-Use Pathogen Data of Concern," *Proceedings of the NeurIPS Workshop on Biosecurity Safeguards for Generative AI* (2026).

149 Ibid.

150 Stephen Turner, "Tiered Access for AlxBio Governance," Paired Ends, Feb. 9, 2026, <https://blog.stephenturner.us/p/tiered-access-AlxBio-governance>.

151 Doni Bloomfield et al., "Biological data governance in an age of AI," *Science* 391, Feb. 5, 2026, <https://www.science.org/doi/10.1126/science.aeb2689>.

The framework mirrors the Centers for Disease Control’s (CDC) Biosafety Level (BSL) laboratory classification system, with five escalating tiers governing access to biological datasets on the basis of corresponding risks. At the lowest level, data remain fully open. The higher the tier, the more restricted the access, starting with basic user registration and monitoring, then requiring institutional affiliation and identity verification, then formal approval of research use, and finally, at the highest level, government review of outputs and strict control over resulting models. The most sensitive data are not downloadable; rather they are accessed through secure environments where users can run analyses without directly handling the underlying data.

Liability Frameworks

Beyond access controls, liability regimes can shape incentives to promote investments in biosecurity. Tort liability creates legal accountability for harms caused by inadequate safeguards. While liability operates after harm occurs rather than preventing harm directly, the threat of liability, when paired with clear standards of care, can shift private sector behavior. Multiple liability models are available, ranging from negligence-based standards to statutory negligence and strict liability approaches. However, current tort liability alone is likely insufficient to address the risks posed by AI and biological design tools. More tailored governance would prioritize *ex ante* safety requirements and structured access controls, with tort law serving as a backstop and incentive for responsible development.

Liability options could include a negligence-based duty of care—modeled on California’s requirements for risk assessment, testing, and safety protocols—or, for certain inherently dangerous and nondeterministic AI systems, strict liability where deployment creates foreseeable risks of catastrophic harm. Some have suggested forms of strict liability may be more appropriate because of AI systems’ nondeterministic nature, in which they can produce different outputs in response to the same input. In this respect, certain high-risk AI systems may be analogized to *ferae naturae*, or wild animals, for which tort law has historically imposed heightened responsibility. Under this rationale, a person who keeps or seeks to domesticate a wild animal may be subject to strict liability for harm caused by the animal, even where the person exercised reasonable care in training or controlling it, because the owner assumes responsibility for a source of danger that cannot be fully controlled. **For a more detailed discussion of potential liability frameworks for biological AI tools, including negligence standards, strict liability models, including *ferae naturae*, Section 230 applicability, and constitutional considerations, see Appendix 3.**

Section C Findings and Recommendations

Finding 3.3.1: Current governance of biological AI tools lacks a standardized, government-supported framework for managing access and monitoring risks across different tool types. While tiered access models and API-based controls exist, they are inconsistently applied, vulnerable to circumvention, and under-resourced. Developers lack the infrastructure needed to implement robust access controls, particularly for high-risk biological design tools and datasets.

Recommendation 3.3.1.1: The federal government should establish a centralized infrastructure (e.g., a secure API hosting and monitoring platform) to support biological AI developers in

implementing standardized tiered access controls, including real-time monitoring, user verification, and risk-based usage restrictions.

Recommendation 3.3.1.2: Congress should mandate the development of unified binding national standards for tiered access and risk classification of biological AI tools, similar to biosafety levels, incorporating best practices from frameworks proposed by NTI and *Science*.

Recommendation 3.3.1.3: Biological AI systems operating under managed access regimes should maintain logging and monitoring systems capable of identifying anomalous or high-risk usage patterns. Developers should establish escalation procedures for suspicious activity, including human review, additional identity verification, temporary access suspension, and referral pathways to appropriate authorities where necessary.

Finding 3.3.2: Safeguards applied to open-source biological AI models are easily circumvented once model weights are publicly released. The Evo1 example demonstrates that withholding dangerous training data is insufficient if downstream users can fine-tune openly available weights.

Recommendation 3.3.2: Establish federal guidance governing release of open-weight biological AI models. Federal agencies should develop mandatory risk-based standards governing the public release of biological AI model weights, training code, and fine-tuning capabilities. Developers of frontier biological design tools should conduct documented dual-use risk assessments before releasing model weights or enabling unrestricted fine-tuning.

D. Nucleic Acid Synthesis Screening and Laboratory Regulation

The convergence of AI and biotechnology is amplified by rapid advancements in DNA synthesis and laboratory automation. DNA synthesis and the creation of synthetic DNA through modern sequencing techniques are a classic dual-use problem. There are numerous beneficial applications, including in therapeutics, vaccine development, diagnostics, and drug discovery research. But the lowering of synthesis costs in the early 2000s also lowered barriers for misuse in development of novel pathogenic agents.¹⁵²

In the United States, the regulation of DNA synthesis processes and products is both piecemeal and permissive. Regulation is principally built on voluntary industry standards. The few existing federal regulatory requirements apply only to exports and imports or in cases where entities are receiving federal funding. There is great opportunity and need for mandatory requirements to strengthen the screening of customers, synthesized sequences, and access to benchtop synthesizers.

Critical Gaps in Current Screening

Existing screening for harmful DNA sequences relies on comparing synthesized sequences against previously identified pathogens and toxins, called **list-based screening**. Two major frameworks

¹⁵² Jonathan B. Tucker, "Double-Edged DNA: Preventing the Misuse of Gene Synthesis," *Issues in Science and Technology* 26, no. 3 (Spring 2010), <https://issues.org/tucker-2/>.

guide sequence screening by synthesis providers: the International Gene Synthesis Consortium (IGSC) Harmonized Screening Protocol and federal guidance on synthetic nucleic acid screening.¹⁵³ Separately, the Federal Select Agent Program (FSAP) regulates the possession, use, and transfer of listed biological agents and toxins. This list-based model has three major vulnerabilities:

- Sequences that do not match a controlled agent do not receive human review, even under rigorous IGSC-compliant screening.
- Sequences modified sufficiently that they are not clearly associated with a pathogen on the screening list can avoid detection. For example, in 2017, Canadian researchers reconstructed an infectious horsepox virus using DNA fragments ordered from GeneArt, a commercial provider that screened orders against lists of regulated pathogens. Because horsepox shares substantial genetic material with other poxviruses, its sequences were not necessarily distinctive enough to trigger additional scrutiny.¹⁵⁴
- Currently, most list-based screening does not cover DNA segments below 200 base pairs. However, many dangerous biological functions including toxin production, antimicrobial resistance, and virulence factors can be encoded in sequences shorter than 200 base pairs. Additionally, sophisticated actors, aided by AI, can order fragments below 200 base pairs separately to evade screening, then assemble them into functional threats.

As a result of these weaknesses, list-based screening is not suited to emerging or poorly understood risks.

In addition to lax requirements for list-based screening, **domestic DNA synthesis providers** are not subject to mandatory requirements for customer or sequence screening. The IGSC Protocol and HHS guidance are both voluntary. FSAP, on the other hand, is mandatory. Managed jointly by CDC and USDA, FSAP comprises a list of biological agents and toxins that have the potential to pose a severe threat to animal or plant health and public safety.¹⁵⁵ Entities possessing select agents are required by law to develop explicit biosecurity and biosafety plans, which are reviewed and certified by FSAP.¹⁵⁶ Although FSAP is very stringent if an agent is covered, if the agent is not covered, there is almost no other federal regulation or oversight for the agent.

Benchtop DNA synthesizers also present a regulatory gap. Benchtop DNA synthesizers are compact, increasingly affordable devices that allow users to synthesize genetic sequences without relying on commercial synthesis providers. As a result, benchtop synthesizers avoid the customer screening and order verification steps that reduce risk at commercial providers. The few binding federal regulations that exist apply only to exports and imports (meaning products used within the country are out of reach) or to federally funded entities. The Export Administration Regulations (EAR) only cover a subset of very capable machines.¹⁵⁷ For example, EAR controls certain automated

153 International Gene Synthesis Consortium, "Harmonized Screening Protocol 3.0," <https://genesynthesisconsortium.org/wp-content/uploads/IGSC-Harmonized-Screening-Protocol-v3.0-1.pdf>; US Department of Health and Human Services, "Screening Framework Guidance for Providers of Synthetic Double-Stranded DNA," <https://aspr.hhs.gov/S3/Documents/syndna-guidance.pdf>.

154 Curtis Matthew Sharkey et al., "Enhancing Gene Synthesis Security: An Updated Framework for Synthetic Nucleic Acid Screening and the Responsible Use of Synthetic Biological Materials," *Applied Biosafety* (2024).

155 By statute, CDC and USDA review and republish the lists every two years. 42 U.S.C. § 262a; 7 U.S.C. § 8401.

156 Ibid.

157 Machines generating continuous nucleic acid sequences greater than 1.5 kilobases in length with an error rate less than 5 percent in a single run.

DNA assemblers and synthesizers and certain software designed for those systems.¹⁵⁸ However, these export regulations do not currently cover most benchtop DNA synthesizers on the market. As next-generation benchtop devices approach the capability thresholds currently covered by export controls, they pose an increasingly important threat.

From List-Based to Function-Based Screening

Transitioning from list-based to function-based screening would yield significant improvement in the ability to detect novel harmful DNA sequences. The National Security Commission on Emerging Biotechnology's September 2025 report to Congress noted that "[l]ist-based systems, such as the Federal Select Agent Program (FSAP), are most appropriate when researchers and innovators know how concerned they should be about a particular organism or type of experiment. List-based systems work less well for emerging or poorly understood risks."¹⁵⁹ The commission, however, did not address solutions to this issue.

Function-based screening provides one solution. A function-based approach looks at what a synthesized sequence *does*, not whether it matches an existing controlled agent on a list. Transitioning to function-based screening would close the most significant structural vulnerability in current screening: the ability of novel, modified, or emerging pathogens to evade detection by deviating from listed sequences.

A 2025 proposal by RAND is currently the most detailed approach to reforming the list-based model.¹⁶⁰ RAND focuses on the FSAP list itself because all individuals are legally required to comply with FSAP regulations, and FSAP coverage would permit criminal investigation. In the absence of reform, the incomplete treatment of bioconvergence will persist: The unauthorized development and possession of a controlled agent is currently a federal crime, but developing a novel pathogen that does not appear in the list is not.

Innovations in machine learning now better enable function-based screening. The Office of the Director of National Intelligence's Intelligence Advanced Research Projects Activity's (IARPA) Functional Genomic and Computational Assessment of Threats (FunGCAT) is developing tools to predict the function of DNA sequences,¹⁶¹ which could be incorporated into existing DNA sequence screening. A next step could entail increased funding into FunGCAT and/or tasking relevant US government agencies to develop a function-based tool that can be added to existing sequence screening. This would require ongoing innovation and adaptation as technologies evolve.

The US synthetic biology industry's long-term commercial legitimacy depends on the transition to function-based screening as much as national security does. Mandatory function-based screening provides industry-wide risk management that no individual firm can provide on its own. Governing synthesis does not constrain industry; rather it is the precondition for the sector's continued ability to operate without the kind of catastrophic incident that would invite far more restrictive intervention.

158 Export Control Classification Number (ECCN) 2B352.j; ECCN 2D352.

159 NSCEB, *Charting the Future of Biotechnology*.

160 Adeli Williams et al., "Identifying and Closing Gaps in the Federal Select Agent Program: Opportunities for Improvement in an Era of Emerging Biotechnologies," RAND, 2025, https://www.rand.org/pubs/research_reports/RRA3628-1.html.

161 Intelligence Advanced Research Projects Activity, "FunGCAT," <https://www.iarpa.gov/research-programs/fun-gcat>.

Laboratory Regulation: Cloud Labs and Contract Research Organizations

One of the traditional barriers to biological misuse is the physical need for sophisticated lab equipment and specialized process knowledge. Cloud labs and contract research organizations (CROs) overcome this barrier by allowing clients to design experiments digitally and have them executed by automated or semiautomated laboratory systems. These services lower barriers to frontier biotechnology while bypassing the tacit knowledge constraints that have historically limited biological weapons development.¹⁶² Current regulation does not extend BSL requirements to cloud labs or CROs unless they receive federal funding. This is a significant regulatory gap that will increase as the commercial cloud lab industry expands.¹⁶³

CROs and contract labs execute research on behalf of clients across discovery, protein engineering, and clinical trials. They have become key agents for pharmaceutical and biotechnology firms by making research and development more efficient through outsourcing work to specialized labs. CROs conducted \$92.27 billion of business in 2025 and are projected to double by 2034.¹⁶⁴ CROs are a critical component in the biotechnology ecosystem, and many are heavily regulated in specific contexts such as clinical trials.¹⁶⁵

As automation and AI increase the remote use of CROs, the safety protocol of those labs becomes more important. One trend among CROs is the rise of “cloud labs,” which allow customers to design experiments digitally while execution occurs in automated, provider-operated facilities,¹⁶⁶ including facilities in other countries. Cloud labs can be researcher-driven or “self-driving.” In the former, researchers actively control and conduct experiments remotely. In the latter, the labs operate autonomously, and the experiments are conducted and managed by AI.¹⁶⁷ Both types of cloud labs present additional risks compared with traditional laboratories. Although some labs implement voluntary safety measures, there are no uniform or mandatory screening protocols to distinguish legitimate researchers from malicious actors or to screen for experiments that could result in dangerous gains-of-function. This gap will become more important as LLMs and BDTs engage directly with laboratories, reducing or even eliminating the need for humans in the process. As autonomous labs develop in capability and prevalence, there are reduced incentives to keep humans in the loop.

The regulatory gap extends to foreign cloud labs and CROs serving US customers or accessing US-provided biological materials. A foreign cloud lab may be subject to US export controls if it receives controlled technology from the United States, but it does not face equivalent biosafety

162 Jeffrey Lee et al., “Documenting Cloud Labs and Examining How Remotely Operated Automated Laboratories Could Enable Bad Actors,” RAND, Apr. 24, 2025, <https://www.rand.org/pubs/perspectives/PEA3851-1.html>.

163 Congressional Research Service, “Oversight of Laboratory Biosafety and Biosecurity: Current Policies and Options for Congress,” R48155, Aug. 8, 2024, <https://www.congress.gov/crs-product/R48155>.

164 Fortune Business Insights, “Contract Research Organization (CRO) Services Market Size, Share & Industry Analysis, by Type,” Feb. 2, 2026, <https://www.fortunebusinessinsights.com/industry-reports/contract-research-organization-cro-services-market-100864>.

165 21 C.F.R. pt. 312 (Investigational New Drug); 21 C.F.R. pt. 812 (Investigational Devices); 21 C.F.R. pt. 56 (Institutional Review Boards); 21 C.F.R. pt. 11 (Electronic Records & Electronic Signatures).

166 Lee et al., “Documenting Cloud Labs”; Kolja Brockmann et al., “Cloud Labs and Other Actors in the Biotechnology Ecosystem: Export Control Challenges and Good Practice Outreach,” Non-Proliferation and Disarmament Paper no. 98 (EU Non-Proliferation and Disarmament Consortium, May 2025), https://www.sipri.org/sites/default/files/2025-05/eunpdc_no_98.pdf.

167 Lee et al., Documenting Cloud Labs and Examining How Remotely Operated Automated Laboratories Could Enable Bad Actors, 2.

screening obligations.¹⁶⁸ Closing this gap requires both domestic laboratory regulation and international alignment on cloud lab biosecurity standards.

Regulating Risky Research

Federal oversight of high-risk biological research has historically relied on a fragmented set of policies applying primarily to federally funded entities, leaving significant portions of the private sector outside formal regulation. Prior to 2024, oversight operated through two Dual Use Research of Concern (DURC) policies and the Department of Health and Human Services' Potential Pandemic Pathogen Care and Oversight Framework, which together governed federally funded research involving select agents or experiments reasonably anticipated to create or enhance pandemic-capable pathogens. In 2024, the White House issued the US Government Policy for Oversight of DURC and Pathogens with Enhanced Pandemic Potential (PEPP)¹⁶⁹ to harmonize these frameworks and replace "gain-of-function" terminology with more specific outcome-based definitions.¹⁷⁰ This framework represented a shift away from regulating specific experimental techniques toward regulating the anticipated consequences of the research itself, reflecting broader recognition that rapidly evolving biotechnology and AI-assisted experimentation make technique-based definitions increasingly inadequate.

Before the policy could fully take effect, however, a May 2025 executive order on Improving the Safety and Security of Biological Research directed the Office of Science and Technology Policy (OSTP) to revise the framework and strengthen centralized oversight, auditing, and enforcement mechanisms. The executive order also called for a strategy to address regulatory gaps affecting non-federally funded research and ordered the suspension or termination of federal funding for certain forms of "dangerous gain-of-function research," including research conducted in jurisdictions deemed to lack sufficient oversight, specifically naming China. OSTP missed the initial 120-day deadline set by the executive order and has yet to publish its revised policy on biological research or the required strategy for oversight of non-federally funded research. The outcome of these revisions will significantly shape the future governance of high-risk life sciences research and determine the degree of regulatory consistency facing both industry and academic institutions.

Section D Findings and Recommendations

Finding 3.4.1: List-based DNA synthesis screening cannot detect novel, modified, or emerging pathogens. List-based screening fails when sequences are modified enough to avoid matching a listed agent. AI-enabled pathogen design will compound this vulnerability by generating novel

168 Kolja Brockmann et al., *Cloud Labs and Other Actors in the Biotechnology Ecosystem: Export Controls and Biosecurity Considerations*.

169 OSTP, *US Government Policy for Oversight of DURC and PEPP*.

170 The policy distinguished between "Pathogens with Pandemic Potential" (PPP), defined as pathogens capable of widespread and uncontrollable transmission, and "Pathogens with Enhanced Pandemic Potential" (PEPP), referring to PPPs modified to increase transmissibility, virulence, or immune evasion. The 2024 policy established a two-tiered oversight structure based on the degree of biological risk posed by the research. Category 1 applied to research involving select agents and toxins reasonably anticipated to increase transmissibility, virulence, or resistance, and required formal institutional risk assessments and mitigation planning. Category 2 imposed heightened review requirements on research anticipated to create, transfer, or use PEPPs or eradicated pandemic pathogens, triggering agency-level risk-benefit evaluations and mandatory safeguards.

sequences specifically designed to evade list-based detection while maintaining their capability to do harm.

Recommendation 3.4.1.1: Mandate function-based, AI-enabled DNA synthesis screening for all domestic synthesis providers regardless of funding source, replacing the current voluntary and list-based framework. Impose import restrictions on foreign DNA synthesis providers that do not conform to US or equivalent screening and customer verification standards. Consider a licensing scheme for DNA synthesis organizations and/or customers.

Recommendation 3.4.1.2: Transition the FSAP system from list-based to function-based, allowing for broader and faster adaptation to novel pathogens of pandemic potential.

Recommendation 3.4.1.3: Consider arrangements to compensate companies for screening costs through tax refunds, subsidies, or other means, as the cost of sequencing falls and screening takes up larger proportions of budgets.

Finding 3.4.2: Mandatory customer-verification requirements do not apply comprehensively to domestic DNA synthesis orders, particularly orders placed outside federally funded research and those produced using benchtop synthesizers. Although federal funding conditions require covered researchers to use providers that follow federal screening standards, much of the broader domestic market continues to rely on voluntary industry practices. While know-your-customer requirements are standard in finance and export controls, they are absent from synthesis. Without knowing who is ordering what, synthesis providers cannot identify suspicious acquisition patterns.

Recommendation 3.4.2: Require customer verification for all DNA synthesis orders exceeding defined thresholds, imposing mandatory know-your-customer protocols. Extend these requirements to benchtop DNA synthesizers as next-generation devices approach commercially relevant synthesis thresholds.

Finding 3.4.3: Cloud labs and CROs are not subject to BSL requirements unless they receive federal funding. Cloud labs lower the tacit knowledge barrier that has historically limited biological weapons development by providing automated execution of experimental protocols. Foreign cloud labs serving US customers or operating with US-origin materials face additional regulatory gaps under current export control frameworks.

Recommendation 3.4.3: Extend BSL regulatory requirements to all commercial cloud labs and CROs operating domestically. Develop international standards for cloud lab biosecurity to address the gap for foreign providers.

Finding 3.4.4: AI-enabled research acceleration and cloud lab proliferation expand access to dangerous biological work. Current federal oversight gaps leave privately funded dangerous research unregulated.

Recommendation 3.4.4: Congress should establish a unified, permanent statutory framework governing high-risk biological research that applies to both federally funded and private-sector entities. Congress should also direct OSTP to finalize clear oversight standards for non-federally funded research, including research conducted through international collaborations or in foreign jurisdictions with inadequate biosafety controls.

Detection

The second pillar of the deterrence by resilience strategy is detection. Because prevention efforts can reduce risks but will never eliminate them entirely, there must be continual efforts to develop response capabilities that can limit impacts. The time between biological emergence and actionable awareness—which can be described as the “detection interval”—determines whether response capabilities can be activated early enough to reduce harms. In the current threat environment, where AI-enabled pathogens may be engineered for stealth and diagnostic evasion, the detection interval is simultaneously the most critical and the most poorly resourced element of US biodefense. This chapter assesses the current state of US bio-detection capabilities, diagnoses its structural failures, and proposes a new architecture built for continuous readiness rather than ramping up for crisis response.

Maintaining persistent bio-detection capabilities is the first step in achieving a “warm start” for responding to known and novel public health emergencies. The concept of the “warm start” refers to initiating use of a capability or capabilities from a partially prepared state rather than from a zero basis (a cold start). It confers a strategic advantage through preexisting biological data, validated platforms, operational infrastructure, and institutional preparedness, and allows detection and response systems to accelerate rapidly whenever a biological threat emerges, whether known or novel. Although most associated with bio-response (addressed in Chapter 6 below), a genuine and effective warm start necessarily begins at the detection stage.

This report employs the term **“biological monitoring system”** to describe an integrated system used to detect, monitor, and analyze biological threats such as infectious diseases, outbreaks, emerging pathogens, or environmental exposures. An often referenced alternative is **“biosurveillance,”** but due to misinformation and misperceptions associated with the term “surveillance,” this report uses “biological monitoring systems” to emphasize the requisite scientific and nonpartisan characteristics of these potentially life-saving capabilities.

Protecting the United States from biological threats requires detecting those threats early. A modern biological monitoring system should be recognized as the foundational infrastructure for national security, public health, and economic resilience. Such a system must be capable of detecting natural outbreaks, accidental releases, and deliberately engineered biological threats before they become crises.

The nature of the biosecurity threat has direct implications for the detection architecture necessary to meet it. An engineered pathogen may be optimized for stealth, delayed clinical presentation, or diagnostic evasion. These characteristics render the traditional symptom-based detection systems ineffective, preventing sufficiently early interventions. Despite broad agreement that early detection of biological threats is a priority, there is currently no operational system that provides continuous, real-time biological situational awareness across the United States.

We have already seen the consequences of being unprepared: The COVID-19 pandemic demonstrated the cost of a detection architecture built for documentation—tracking and confirming cases after they are already known—rather than building for early warning capacity that can potentially limit the spread of harms. Building the system we need for the future requires confronting that failure, understanding its structural causes, and committing to a better model that can respond to biosecurity challenges, now and in the future.

A. Diagnosing Failures in the Existing Biological Monitoring Systems: The COVID-19 Pandemic as a Stress Test

SARS-CoV-2, the virus that causes COVID-19, was likely circulating in Wuhan, China, for weeks before Chinese authorities reported the presence of the virus to the WHO on December 31, 2019. By the time the World Health Organization declared a Public Health Emergency of International Concern on January 30, 2020, the virus had already spread widely, causing outbreaks across multiple continents.¹⁷¹ In the United States, retrospective genomic analysis suggests community transmission was underway in multiple cities by late January 2020, weeks before any public health action was taken.¹⁷² This gap in the United States and elsewhere between community transmission and public health acknowledgment was driven by inadequate detection capabilities and had devastating consequences for controlling the virus and saving lives.

The lag in detecting the spread of SARS-CoV-2 was the result of biological monitoring systems designed around a series of assumptions that were proved wrong: 1) that threats would present clinically before they spread widely; 2) that existing pathogen-specific diagnostic panels would be sufficient; and 3) that international reporting systems would provide timely early warning. These failures provide compelling lessons for policymakers today.

Structural Failures of the Existing System

In 2020, the US biological monitoring architecture that confronted COVID-19 was a patchwork of systems with significant gaps, limited interoperability, and no real-time capability. Its failures were structural. Made of a series of different biological monitoring systems, all with discrete mandates, data architectures, and reporting chains, the US biological monitoring ecosystem had no mechanism for integrating its components into a coherent national picture. Each system could see part of the threat, but none could see all of it. This can be analogized to the ways in which counterterrorism architecture failed us before 9/11: There was plenty of disaggregated information and because it was not integrated and analyzed in its totality and in a relevant timeframe, our warning system failed to “connect the dots” to prevent catastrophic attacks on the homeland.¹⁷³

171 World Health Organization, “Statement on the Second Meeting of the International Health Regulations (2005) Emergency Committee Regarding the Outbreak of Novel Coronavirus (2019-nCoV),” Jan. 30, 2020.

172 Michelle A. Jorden et al., “Evidence for Limited Early Spread of COVID-19 Within the United States, January–February 2020,” *Morbidity and Mortality Weekly Report* 69, no. 22 (2020): 680–84.

173 National Commission on Terrorist Attacks upon the United States, *The 9/11 Commission Report* (Government Printing Office, 2004), 408.

Table 4.1 summarizes the principal components of the current biological monitoring architecture in early 2020.

Type	System / Example	How It Works
Syndromic surveillance	CDC BioSense; ESSENCE (emergency department network)	Aggregates emergency department visit data and clinical signals to flag anomalous patterns in real time
Environmental / wastewater	CDC National Wastewater Surveillance System (NWSS); SecureBio Nucleic Acid Observatory	Samples wastewater at treatment plants to detect pathogen RNA/DNA across large populations; currently PCR-based, transitioning to metagenomic sequencing
Genomic surveillance	GISAID; CDC genomic sequencing networks; SARS-CoV-2 variant tracking	Sequences pathogen genomes from clinical and environmental samples to track variant emergence, transmission lineages, and novel features
Air / environmental sampling	BioWatch (DOD/DHS)	Deploys air samplers in high-risk urban locations to detect aerosolized biological agents; samples sent to laboratories for analysis
Laboratory / diagnostic reporting	State and local public health labs; clinical laboratory networks	Clinical laboratories report confirmed diagnoses of notifiable diseases and novel agents to public health authorities via mandatory reporting chains
Behavioral / ancillary signals	Prescription monitoring; absenteeism data; internet search trends; over-the-counter sales	Tracks proxy indicators of population illness—pharmacy sales, school/workplace absences, online symptom searches—as leading signals ahead of clinical presentation
AI-enabled analytics	BioRadar (Lawrence Livermore National Laboratory model); contracted analytic services	Integrates heterogeneous data streams (environmental, clinical, genomic, behavioral) and applies machine learning to detect anomalies against modeled baselines
International reporting	WHO International Health Regulations (IHR); Global Outbreak Alert and Response Network (GOARN)	Treaty-based obligations requiring member states to report public health emergency of international concern to the World Health Organization (WHO) within specified timeframes

Syndromic Surveillance

Syndromic surveillance systems, including the CDC’s BioSense platform and the ESSENCE network, aggregate hospital emergency department visit data and other clinical signals to detect anomalous patterns. These systems provide meaningful situational awareness but only *after* an illness has become widespread. Syndromic surveillance systems respond when enough symptomatic patients go to healthcare facilities to generate a statistically detectable signal. Further, these systems miss the spread of illnesses in populations with limited access to care. For

pathogens capable of wide asymptomatic spread, syndromic surveillance is a lagging indicator of infection. During COVID-19, these systems confirmed what clinicians already knew (people were getting sick) and they alerted long after early warning would have mattered.

Because they rely on symptom reporting, these systems also have trouble detecting new illnesses with symptoms similar to common illnesses. For instance, SARS-CoV-2 arrived in the United States at the height of influenza season. Early COVID-19 patients presented with respiratory symptoms indistinguishable from the flu. Before COVID, it was not routine to do tests to confirm that influenza-like symptoms were caused by the flu, so the initial surge in cases was absorbed into the background noise of a normal flu season. The syndromic surveillance systems saw elevated respiratory illness signals, which were indistinguishable from the normal seasonal influenza patterns.

Syndromic surveillance detects signal against a background of known disease patterns. But a pathogen that mimics an existing disease category can evade detection indefinitely, until case counts overwhelm the background or clinicians happen to test for the right thing. Any sufficiently sophisticated pathogen, whether naturally occurring or engineered, can exploit this structural vulnerability. A novel agent that produces respiratory symptoms is likely to be attributed as this flu season data. One that causes gastrointestinal illness may disappear into the background of foodborne disease. One engineered for neurological presentation may be attributed to any number of common causes. In the context of deliberately engineered biological threats where evasion of clinical detection may be an explicit design objective, this is a predictable attack surface.

Laboratory Reporting

The US laboratory reporting infrastructure, which includes public health laboratories, clinical laboratories, and the Laboratory Response Network, is oriented toward confirming diagnoses of known pathogens rather than identifying new ones. Pathogens that fall outside existing diagnostic panels pass through these systems undetected.

The COVID-19 pandemic exposed a more acute regulatory failure within this system. Under initial FDA policy, clinical and academic laboratories could not use their own SARS-CoV-2 diagnostic tests for patient testing without emergency authorization, delaying the expansion of testing until the agency relaxed this policy in late February 2020. This remained true even as the only authorized test, developed by the CDC, was producing invalid results because of a manufacturing contamination problem. For critical weeks in February and early March 2020, U.S. testing capacity remained extremely limited. Academic medical centers and commercial laboratories had the scientific capability to develop and run accurate COVID-19 tests but could not deploy them for clinical use. The result was a catastrophic diagnostic blackout at the precise moment when broad-scale testing could have better defined the outbreak and enabled targeted containment.

The contrast with the Republic of Korea (ROK) is instructive and sobering. The ROK detected its first COVID-19 case on January 20, 2020, the same day as the United States. The ROK rapidly authorized and deployed diagnostic testing through public and private laboratories; the government reported conducting approximately 20,000 tests per day by mid-February through a network of private laboratories, hospitals, and drive-through test centers. Its broader test–trace–treat strategy helped

contain transmission while preserving substantial freedom of movement.^{174, 175, 176} The technology was the same as what US laboratories possessed, but the difference was institutional: The ROK had preestablished emergency authorization frameworks that allowed rapid deployment of the tests, while the United States had a regulatory posture that concentrated authority and created a single point of failure. The ROK contained its initial outbreak without a national lockdown. The United States did not.

BioWatch

The BioWatch program is the nation's primary environmental biological monitoring infrastructure. Begun in 2003, it is a federally managed, locally operated system designed to provide early warning of an aerosolized biological attack. It functions in more than thirty jurisdictions nationwide using approximately 600 aerosol collectors, primarily in outdoor locations. Air samples are collected and transported to laboratories, where PCR-based tests look for the genetic signatures of specific biological threat agents. The full process, from the initial capture of an agent to determining whether a public health threat exists, can take between twelve and thirty-six hours, limiting the system's ability to provide rapid warning. BioWatch also depends on tests designed to recognize known genetic signatures. As a result, it may fail to identify novel threat agents for which no test has been developed or to identify agents that have been altered so that existing tests do not recognize them.¹⁷⁷

BioWatch's limitations were demonstrated by COVID-19, not because the program failed, but because the pandemic was not a threat that it was designed to detect. BioWatch monitors fixed locations for aerosolized threat agents associated with a deliberate attack. SARS-CoV-2 spread through sustained person-to-person respiratory transmission across the entire country, in settings BioWatch did not monitor, through a mechanism it was not built to discover. As a result, the BioWatch system played no role in the largest biological emergency in a century because it was designed to detect a specific threat: a bioterrorism event using a known biological agent. This approach no longer reliably maps onto the current threat landscape.¹⁷⁸

International Reporting

The COVID-19 pandemic also exposed the limits of the current international bio-detection system, which depends on voluntary national reporting, particularly when early disclosure carries significant economic and political costs. Under international health regulations, countries must assess public health events according to specific guidelines and notify the World Health Organization (WHO) within twenty-four hours of assessing an event that may constitute a "public

174 World Health Organization, "COVID-19—Republic of Korea," Disease Outbreak News, Jan. 21, 2020, <https://www.who.int/emergencies/disease-outbreak-news/item/2020-DON238>

175 U.S. Government Accountability Office, COVID-19: Continued Attention Needed to Enhance Federal Preparedness, Response, Service Delivery, and Program Integrity, GAO-21-551 (July 2021), <https://www.gao.gov/products/gao-21-551>

176 World Health Organization, "Republic of Korea: Success Against COVID-19 Based on Innovation and Public Trust," Dec. 2, 2020, <https://www.who.int/news-room/feature-stories/detail/republic-of-korea-success-against-covid-19-based-on-innovation-and-public-trust>

177 GAO, "Biodefense: DHS Exploring New Methods to Replace BioWatch and Could Benefit from Additional Guidance," GAO-21-292 (May 20, 2021), <https://www.gao.gov/products/gao-21-292>.

178 US Department of Homeland Security, Office of Inspector General, "Biological Threat Detection and Response Challenges Remain for BioWatch," OIG-21-27 (March 2021), <https://www.oversight.gov/reports/audit/biological-threat-detection-and-response-challenges-remain-biowatch-redacted>.

health emergency of international concern.” However, there is no enforcement mechanism, and compliance is in practice voluntary. China’s delayed reporting on COVID-19 was the predictable outcome of a global system built on voluntary compliance, with inadequate enforcement mechanisms and limited independent detection capacity. A detection architecture that depends heavily on the timely national disclosure despite strong political and economic incentives to delay reporting further exacerbates this vulnerability.

Wastewater Monitoring

Perhaps the most important bio-detection lesson from the COVID-19 pandemic was the value of wastewater monitoring and analysis. Studies showed that SARS-CoV-2 RNA could be detected in wastewater four to seven days before a corresponding surge in clinical cases, providing the early warning syndromic systems were unable to deliver.¹⁷⁹ The CDC’s National Wastewater Surveillance System, which monitors sewage for viral genetic material to detect and track community-level disease transmission, expanded rapidly during the pandemic, ultimately covering a significant portion of the US population. But the program relied primarily on targeted PCR assays for known pathogens. It has great, but untapped, potential as a pathogen-agnostic early-warning platform to detect unexpected biological activity.

B. Building Bio-Monitoring Systems for Continuous Readiness

Like clinical trials, discussed in Chapter 6 of this report, public health detection systems must be built in advance, rather than improvised under pressure. Detection systems assembled in response to a known threat arrive too late. The window for early, low-cost public health interventions such as isolation, targeted testing, and accelerated countermeasure development sometimes closes within days of a pathogen’s initial emergence. **The central challenge is to maintain a continuously operating early-warning system capable of detecting the unexpected.** This would contribute an essential element to the “warm start” capacity that our nation requires to enhance biosecurity resilience.

Foundational Principles for Modern Bio-Detection

Effective bio-detection for the current threat environment requires a fundamentally different architecture than what has existed to date. The new architecture must be centered on early warning with speed and breadth. In 2016, the Department of Energy’s Lawrence Livermore National Laboratory (LLNL) identified several design principles for such systems that remain valid today:¹⁸⁰

179 Jordan Peccia et al., “Measurement of SARS-CoV-2 RNA in Wastewater Tracks Community Infection Dynamics,” *Nature Biotechnology* 38, no. 10 (2020): 1164–67.

180 Stephan Velsko and Thomas Bates, *A Conceptual Architecture for National Biosurveillance: Moving beyond Situational Awareness to Enable Digital Detection of Emerging Threats*, LLNL-JRNL-680437 (Lawrence Livermore National Laboratory, June 17, 2016), <https://www.osti.gov/biblio/1262173>.

Framework for Effective Bio-Detection

- **Comparative baseline monitoring.** Effective bio-detection is defined by detecting abnormal biological activity. Knowing that a signal is anomalous, however, requires comparison against a well-characterized baseline of normal activity so that missing pathogens can be detected through deviations from the baseline.
- **Continuous operation.** Uninterrupted analysis of a community's biological activity is the only means of identifying deviations from the baseline efficiently and deploying targeted interventions before a threat becomes clinically apparent.
- **Population-scale.** Environmental biomonitoring, including wastewater and air sampling, captures population-level biological activity independent of individual clinical encounters. It can detect asymptomatic or low-level pathogen transmission that evades clinical systems.
- **Pathogen-agnostic.** Systems built around targeted assays can only detect known pathogens. In a threat environment shaped by rapid biological innovation and potential adversarial evasion, that is a structural weakness. Metagenomic sequencing, which characterizes all genetic material in a sample, including material with no match in any existing diagnostic panel, enables detection of novel variants, unexpected organisms, and engineered agents that targeted approaches miss.
- **AI-enabled.** The data volumes generated by continuous, population-scale, pathogen-agnostic bio-detection require automated analysis. AI is the only practical means of integrating large quantities of heterogeneous biological, environmental, clinical, and behavioral data into a coherent representation of biological normalcy. With this baseline, AI analysis can flag meaningful departures from baseline quickly enough for human decision-makers to act. AI enables scalable filtering, baseline learning, anomaly detection, and threat characterization across heterogeneous data streams. It can identify unusual genetic features, assess signal persistence and spatial coherence, and distinguish noise from meaningful results.
- **Decision-oriented.** Effective early-warning systems must be designed for decisions and built around the needs of the people who have to make them. They should clearly define what response options are available, at what confidence threshold, and for which decision-maker. Outputs must be calibrated to the needs of public health officials, military commanders, and national security leaders.

LLNL's framework has been reinforced by more recent studies by the National Academies and RAND Corporation.¹⁸¹ With these principles in mind, the following sections outline recommendations for how a stronger continuous bio-detection system could be built.

181 National Academies of Sciences, Engineering, and Medicine, *The Age of AI in the Life Sciences: Navigating Opportunities and Risks* (2025), <https://www.nationalacademies.org/read/28868>; RAND Corporation, "Bringing New Technologies to Bear for Biosurveillance," April 2024, <https://www.rand.org/pubs/commentary/2024/04/bringing-new-technologies-to-bear-for-biosurveillance.html>.

C. Building the System: Infrastructure and Integration

With the lessons from what failed during COVID-19 and an understanding of the necessary characteristics of a successful bio-detection system, we have the foundations for a new architecture for the United States. Built on environmental biomonitoring, genomic surveillance, and national security integration, the next generation system will need to evolve into one that is substantially more capable of identifying novel threats and providing early warnings.

Wastewater as Primary Environmental Monitoring Infrastructure

Wastewater monitoring and analysis has established itself as the most scalable, cost-effective, and analytically powerful environmental biomonitoring modality available. A single wastewater sampling site can provide biological signal from entire communities, capturing population-level data anonymously, and aggregating community-wide biological signal independent of individual healthcare encounters. Wastewater monitoring has repeatedly demonstrated the ability to provide days of advance warning ahead of clinical case surges.

The CDC's National Wastewater Surveillance System (NWSS) infrastructure was built during the COVID-19 pandemic and provides a foundation upon which to build. Its primary limitation is its reliance on targeted assays for known pathogens. Between 2021 and 2024, the federal government invested approximately \$500 million in NWSS to build a targeted-assay system covering approximately 47 percent of the US population. That coverage was achieved primarily through monitoring of a relatively small number of large urban treatment plants, each serving hundreds of thousands to millions of people.

Transitioning NWSS from a COVID-specific monitoring program to a pathogen-agnostic, continuously operating environmental early-warning network requires investment in metagenomic sequencing capacity, standardized pre-analytical protocols, baseline modeling infrastructure, and AI-enabled analytic capability. Expanding coverage beyond the current NWSS reach would require adding sites in smaller, more dispersed, and rural communities where per-site population yield is substantially lower. Adding pathogen-agnostic metagenomic sequencing, which is much more expensive than PCR testing, and AI-enabled analytics will also increase costs. To be most effective, new wastewater infrastructure must take the following issues into account:

Strategic site selection. Not all wastewater sites are equal as early-warning assets. Sites should be selected based on population coverage, flow dynamics, and likelihood of early signal emergence. Good locations include high-traffic transportation hubs, dense residential areas, and facilities serving vulnerable populations. For additional reasons, the system will also need to account for dispersed and rural sites where disease can also materialize and spread, as seen with the Avian flu.

Longitudinal consistency. Repeated measurement at the same locations over time is as important as site coverage. Repeated measurements allow the creation of a well-characterized baseline against which anomalies can be measured.

Integration of existing data streams. The most powerful early-warning architecture integrates multiple independent data sources: wastewater and air sampling, syndromic

surveillance data from emergency department networks like the Electronic Surveillance System for the Early Notification of Community-Based Epidemics (ESSENCE),¹⁸² genomic surveillance from clinical and laboratory sources, and behavioral signals including prescription patterns and absenteeism data. Integration across these streams is both a technical and a governance challenge. Technical interoperability of these diverse data streams requires standardized data formats, shared categories and labels, and secure data exchange infrastructure. Governance of this complex system requires clear authorities for data sharing across federal agencies and between federal and state systems—a domain where COVID-19 exposed significant gaps.

Genomic surveillance. Real-time genomic surveillance, including tracking the evolutionary trajectory of circulating pathogens, is a critical component of a complete bio-detection architecture. The importance of genomic surveillance was powerfully demonstrated during the COVID-19 pandemic: Genomic sequencing enabled identification of variants of concern, tracking of transmission chains, and assessment of vaccine effectiveness in near real time. Domestically, genomic surveillance capacity is concentrated in academic centers and large public health laboratories. Building distributed sequencing capacity, including at community hospitals and local public health laboratories, would substantially improve the timeliness and geographic representativeness of genomic data.

Global data infrastructure. Sharing genomic data across the globe depends on a robust data-sharing infrastructure with appropriate governance. During COVID-19, GISAID (the Global Initiative on Sharing All Influenza Data) demonstrated the importance of data sharing. On January 10, 2020, within forty-eight hours of obtaining the first sequences identifying an unknown betacoronavirus from patients in Wuhan, the China CDC made the first whole-genome sequences available via GISAID.¹⁸³ This early submission supported the rapid development of SARS-CoV-2 vaccines and diagnostic tests. GISAID's attribution-based model, under which governments and scientists receive credit for the sequences they submit, has been credited with encouraging broad participation in genomic data sharing.

Following on the success of sharing COVID data, GISAID added mpox and RSV to its supported pathogen list in 2022. GISAID's governance has drawn criticism from the European Commission, the Rockefeller Foundation, and others in the scientific community for its lack of transparency and its prohibition on resharing data.¹⁸⁴ Any successor framework for priority pathogens should preserve GISAID's attribution incentive structure while establishing more transparent governance and clearer mechanisms for dispute resolution.¹⁸⁵

182 ESSENCE is a syndromic surveillance platform originally developed by the Johns Hopkins University Applied Physics Laboratory in the late 1990s, and subsequently adopted by Walter Reed Army Institute for Research, CDC, and public health departments across the United States. It aggregates emergency department visit data, chief complaint information, and other clinical signals in near real time to detect anomalous patterns that may indicate disease outbreaks or bioterrorism events. CDC's BioSense platform runs on ESSENCE infrastructure. Howard Burkom et al., "Electronic Surveillance System for the Early Notification of Community-Based Epidemics (ESSENCE): Overview, Components, and Public Health Applications," *JMIR Public Health and Surveillance* 7, no. 6 (2021).

183 GISAID, "Data Availability," Mar. 21, 2023, <https://gisaid.org/statements-clarifications/data-availability/> (accessed May 7, 2026).

184 Martin Enserink, "Fresh Conflicts Erupt Around Giant Database for Flu and COVID-19 Sequences," *Science*, Jan. 6, 2026; Meredith Wadman and Jon Cohen, "The 'Invented Persona' Behind a Key Pandemic Database," *Science*, *Science* 380, no. 6643, April 2023.

185 Timothée Poisot and Colin J. Carlson, "To Finish the Pandemic Agreement, WHO Needs a Trustworthy Viral Database," *Think Global Health*, Nov. 5, 2025.

GISAID: Global Genomic Data Sharing

GISAID (Global Initiative on Sharing All Influenza Data) is an international public-private initiative launched in 2008 to enable rapid sharing of pathogen genomic sequence data. It was established as an alternative to traditional public-domain archives, where access is anonymous and submitters retain no rights over their data, conditions that deterred sharing, particularly by researchers in lower-income countries concerned about losing credit or intellectual property.

GISAID is structured as a horizontal, collaborative global platform rather than a centrally managed organization. Its administrative functions are handled by Freunde von GISAID e.V., a registered nonprofit association based in Germany. It operates as an independent public-private partnership with government contributions from Brazil, China, France, Germany, Senegal, Singapore, the United Kingdom, and the United States, and receives grants from WHO and private donors. Scientific oversight is provided by a Scientific Advisory Council drawn from leading public health laboratories, including WHO Collaborating Centres.

GISAID's core mechanism is attribution-based access: Contributors share sequence data and receive credit when it is used, with access controlled through registered credentials and terms of use governing how data may be cited and shared. This model has enabled collaboration among more than 42,000 researchers from 198 countries across more than 3,500 institutions. During the COVID-19 pandemic, GISAID's EpiCoV database became the primary global repository for SARS-CoV-2 sequences, supporting near-real-time variant tracking and vaccine development.

National security integration. Early detection of an unusual biological signal is simultaneously a public health and a national security event. The biological monitoring architecture that the United States builds must be designed with this dual character in mind. This means integrating biomonitoring outputs into national security information flows and decision-making structures, ensuring that early biological signals reach decision-makers with national security responsibilities quickly enough to be actionable. It also means building attribution capability alongside detection capability: The ability to determine whether an anomalous biological event is natural, accidental, or deliberate shapes the appropriate response (addressed in Chapter 5 below). Finally, this means engaging with the complex need to treat early-warning data as intelligence, with the classification, handling, and dissemination protocols that implies, while simultaneously maintaining the open data flows necessary for scientific analysis and response. The military's biological defense programs, including laboratory networks and environmental monitoring assets, represent significant untapped capacity for civil biological monitoring. Greater integration between DOD biological defense activities and civilian disease monitoring infrastructure would strengthen both.

Building community trust and support. A comprehensive biological monitoring system will be effective only if it is accompanied by sustained public trust, transparency, and community engagement.¹⁸⁶ Public trust is not just a communications consideration but

186 National Academies of Sciences, Engineering, and Medicine. Wastewater-Based Disease Surveillance for Public Health Action. Washington, DC: National Academies Press (2023). <https://doi.org/10.17226/26767>; Nikki Romanik and Ashish K. Jha, "Biological Threats and Community-Level Approach to Early Detection," *BMJ* 392 (2026).

an operational requirement: A monitoring system that communities do not trust will fail in practice.¹⁸⁷ Jurisdictions will decline to participate, local officials will resist site placement, and political support for stable funding will erode. Wastewater surveillance, despite its demonstrated value during the COVID-19 pandemic, has generated public concern that governments may use sewage sampling to identify or monitor individuals or stigmatize particular groups, even though public health systems generally analyze aggregated community-level samples rather than personal biological data.¹⁸⁸ The same boom-and-bust dynamic that has undermined past investments in biomonitoring infrastructure will recur if the public perceives the system as a surveillance apparatus rather than a public health asset.

Privacy protections are therefore a prerequisite for a functioning early-warning system. The most significant technical privacy concern is genomic: The transition to pathogen-agnostic metagenomic sequencing, while essential for detecting novel and engineered agents, raises important questions about human genetic privacy, data stewardship, and permissible use.¹⁸⁹ These concerns are manageable, but only if addressed through binding technical standards and governance rules from the outset.

Pathogen-agnostic metagenomic sequencing of wastewater can be implemented in a manner that is highly protective of individual genomic privacy. Environmental wastewater samples capture community-level biological signals from large populations, not data from individual patients or households, and therefore carry lower individual privacy risk than clinical genomic monitoring data drawn from individual patients.¹⁹⁰ However, privacy risk is not uniform: It increases when sampling occurs at smaller geographic scales, such as buildings, campuses, correctional facilities, military installations, tribal communities, ships, ports of entry, or other facility-specific sites. A national framework should recognize that privacy risks vary by sampling scale and use case. Community-level sewershed monitoring should be subject to baseline privacy and data-use protections, while more granular sampling should include proportionate safeguards—including clear purpose, data minimization, controlled access, and limits on public reporting where appropriate, and engagement with institutional leadership as appropriate.

As a mandatory condition of federal funding, metagenomic biomonitoring should use validated protocols to minimize, remove, or exclude human genetic material from analysis. These safeguards should include pre-sequencing host-DNA depletion where feasible, post-sequencing computational removal of human reads, residual human-sequence quality-control thresholds, independent auditability, and strict limits on retention or sharing of raw reads that may contain human genomic material. The public message should be clear and accurate: The system is designed to measure community-level pathogen signals, not to sequence or identify individuals.

187 Kata Chillag and Elizabeth Fenton. “Framework for Addressing Ethical Considerations in Infectious Diseases Public Health Wastewater Surveillance.” Association of State and Territorial Health Officials. Mar. 10, 2025. <https://www.astho.org/topic/report/addressing-ethical-considerations-in-wastewater-surveillance/>.

188 Dhiraj Nainani et al., “Balancing Public Health and Group Privacy: Ethics, Rights, and Obligations for Wastewater Surveillance Systems,” *Water Research* 258 (2024).

189 Hunter Jackson Smith, Richard T. Agans, and William J. Kowallis. “Ethical Considerations for Wastewater Surveillance Conducted by the US Department of Defense.” *JMIR Public Health and Surveillance* 11: e67145 (2025). <https://doi.org/10.2196/67145>.

190 Ibid.

Beyond that technical safeguard, biological monitoring data may be used only for purposes directly related to public health, biosecurity, and clearly defined biomedical research and countermeasure development. Use for purposes unrelated to this mission, such as for law enforcement, immigration enforcement, employment or insurance decision-making, commercial resale, non-public-health profiling, or other punitive uses, should be prohibited absent explicit statutory authorization. Secondary research uses should require independent review, data minimization, and public reporting sufficient to preserve public trust.

These protections require structure to be meaningful: An independent advisory body with community, public health, ethics, statistical, tribal, and data-governance expertise should review data-use practices, requests for access to more granular data or samples, proposed secondary uses, and any expansion of the system's purposes. This board should be supported by local community input mechanisms, with defined retention limits, data-use agreements, access logs, and auditable enforcement. The governance framework should also distinguish raw metagenomic reads, processed pathogen calls, aggregate trend data, and public dashboards, because each carries different privacy, security, and communication risks.

This governance framework must also explicitly address any circumstances under which biological monitoring data may be shared with national security entities or protected through classified channels. These circumstances should be governed by clear criteria, publicly disclosed whenever feasible, including the conditions under which data may be classified, what minimization rules apply, and how long data may be retained. Classification decisions and national-security uses should be subject to independent oversight, with public reporting wherever possible, so that the system's security functions do not inadvertently undermine the public trust on which its detection functions depend.

D. Financing and Governing a Biological Monitoring System

US biological monitoring investment follows a predictable and dysfunctional pattern: It surges in response to crises and contracts in their aftermath. For example, the NWSS, built at significant expense during COVID-19, faces uncertain long-term funding. Similarly, the CDC's biomonitoring programs are subject to appropriations cycles that create operational instability incompatible with the requirements of continuous early-warning systems. Supplemental emergency funding often arrives too late to build the infrastructure needed to respond to the emergency that triggered it.

This funding pattern reflects deeper governance challenges. Federal agencies face structural incentives that make sustained investment in continuous early-warning systems difficult even when resources are available. Priorities shift with changes of administration. Reorganizations disrupt institutional memory. Congressional appropriations cycles reward visible emergency spending over unglamorous continuous operations. The result is a system perpetually preparing for the previous pandemic rather than optimizing for future crises. This need mirrors the broader imperative to develop meaningful, sustainable, and ongoing readiness and resilience against biological attacks.

The Limits of Government as Builder

The responsibility to build biological monitoring capabilities is currently distributed across multiple federal agencies, each operating independently. These include, *inter alia*, the Centers for Disease Control (CDC), Administration for Strategic Preparedness and Response (ASPR); Department of Homeland Security (DHS); Department of Defense (DOD); the intelligence community; the US Department of Agriculture (USDA); the Federal Bureau of Investigation (FBI); and the Environmental Protection Agency (EPA).¹⁹¹ Each agency brings important capabilities and data streams. **What are lacking, however, are the authority, mandate, and organizational incentives to integrate these disparate organizations into a coherent national early-warning system.** The National Biosurveillance Integration Center (NBIC) within DHS has had a biological monitoring coordination mandate since 2007, but that mandate has been ineffective. The persistent failure to integrate agency data streams across that period reflects a structural problem that additional coordination authority alone will not solve. The restructured center must have the ability to set binding data-sharing standards and condition agency participation.

The creation of a state-of-the-art, government-led, integrated national early warning system is hampered by various obstacles. Government procurement cycles run years behind the pace of technical innovation. Hiring constraints limit access to the scientists, engineers, and data analysts these systems require. Interagency coordination is slow. Systems built inside agencies are perpetually vulnerable to the next budget cycle or reorganization. Although the government needs to set requirements, establish standards, provide oversight, and fund public health and security priorities, building and operating cutting-edge technical systems at speed should largely be done by the private sector.

A Private-sector-led, Government-funded Model

The central organizing principle of the recommendations below is a deliberate and sustained public-private partnership to build a comprehensive biological monitoring system. Recognizing the respective strengths and limitations of government and the private sector, it makes sense to create a system where government defines the requirements, sets performance standards, provides stable funding through contracts and grants, and maintains oversight and data access. Private sector companies such as specialized biological monitoring firms, national laboratories, genomic sequencing companies, and AI developers would compete to build and operate the capability within the guidelines defined by the government. This is the application of a proven model for building technical capability and one that the federal government has used for decades in defense, space, and intelligence. For example, the National Reconnaissance Office (NRO) has accomplished the vast majority of its technical work through private contractors, an approach that its former director Christopher Scolese described as “buy what we can, build what we must.”¹⁹² The intelligence community’s geospatial and analytic platforms are similarly built and operated primarily by contracted firms.

191 Institute for Progress, “Mapping America’s Biosurveillance,” November 2024, <https://ifp.org/mapping-americas-biosurveillance/>.

192 Naomi Cooper, “Christopher Scolese: NRO Looks to Expand Commercial Satellite Capabilities,” ExecutiveGov, Aug. 5, 2022, <https://www.executivegov.com/articles/christopher-scolese-nro-looks-to-expand-commercial-satellite-capabilities/>; L. Elaine Halchin, “The Intelligence Community and Its Use of Contractors: Congressional Oversight Issues,” Report R44157 (Congressional Research Service, Aug. 18, 2015), <https://www.congress.gov/crs-product/R44157>.

E. Conclusion

The tools required for a modern biological monitoring system (e.g., metagenomic sequencing, AI-enabled anomaly detection, environmental sampling) are available and, in important domains, already deployed, although not at scale. The urgent work to be done is largely in the domain of governance, financing, and institutional design: creating the authorities, the stable resources, and the integrated decision-making pathways that enable these tools to be widely deployed and to function as a coherent national early-warning system.

COVID-19 demonstrated, at enormous cost, what the absence of such a system means in practice. The window between biological emergence and the closure of early intervention options is measured in days and weeks. Building the system before the next pandemic arrives is the defining preparedness challenge of this moment.

In the next chapter, we will explore attribution. Early detection and accurate attribution are complementary capabilities. Detection establishes that an anomalous biological event is occurring; attribution determines whether it is natural, accidental, or deliberate, and if deliberate, by whom. That determination shapes every subsequent decision: the distinction between a public health response and a national security response, the evidence standard required for a diplomatic or military action, and the legal authorities that govern the government's next steps. A biological monitoring system that detects a novel pathogen but cannot characterize its origin leaves decision-makers with situational awareness but without the information they need to determine appropriate action. Furthermore, effective attribution can create a deterrent effect. For these reasons, the following chapter addresses the attribution capabilities the United States must build as quickly as possible.

Findings and Recommendations

Finding 4.1: The United States has no unified, authoritative strategy for national biological monitoring requirements. Coverage thresholds, detection lag standards, data interoperability obligations, and performance metrics are either undefined or fragmented across agencies. Without a common requirements framework, neither procurement nor accountability is possible.

Recommendation 4.1.1: Establish a national biological monitoring requirements framework. The White House and Congress should both take actions to establish a centralized and harmonized biological monitoring system within the US government. The White House can use executive actions based on existing authorities to establish a national biological monitoring requirements framework. The framework should be White House-led and designate the **DOD and HHS as joint lead federal agencies** responsible for maintaining the requirements framework and overseeing interagency coordination.

At a minimum, the framework must define:

- population coverage thresholds for environmental monitoring;
- maximum acceptable detection lag relative to clinical case emergence;
- data interoperability standards for federal, state, and private data streams; and,

- performance reporting requirements for all federally funded biomonitoring activities.

The framework must be **technology-neutral**, specifying outcomes rather than methods, to preserve room for private sector innovation and avoid being locked into approaches that will quickly become obsolete. Congress should further codify this structure and responsibilities in statute as well as provide any authorities, including for procurement, that the agencies identify as missing from their ability to fully launch a comprehensive biological monitoring system.

Recommendation 4.1.2: Incentivize nationwide investment in wastewater monitoring through competitive, performance-based contracts. These incentives and investments should include state and local public health agencies and hospital systems so that wastewater surveillance outputs will trigger specific actions from the public health and healthcare sector, such as increased testing for specific disease threats, deployment of new or existing medical countermeasures, and/or intelligence or law enforcement actions to find and stop malign actors. Require pathogen-agnostic metagenomic sequencing at all US sites within five years.

Finding 4.2: Public health agencies are trained to detect and respond to disease, not to assess threat intent, evaluate adversary capabilities, or integrate classified intelligence into operational decisions. As a result, biological signals that might warrant national security attention are routinely processed through an epidemiological lens alone, while intelligence assessments of biological threats rarely incorporate real-time public health monitoring data. This gap is not a failure of competence; it is a structural mismatch. Bridging requires a dedicated interface function, a standing team with the training clearances, and institutional authority to operate fluently in both domains.

Recommendation 4.2: Ensure HHS has a dedicated team of biosecurity experts capable of coordinating with the NSC, DOD, DOE, DHS, and the IC. Congress should require that HHS establish a new unit dedicated to the biosecurity staffed by security-clearance holding experts and appropriate dedicated no-year funding directly to that unit. Congress can require HHS to report on how it will implement this requirement, but at a minimum, the team must be capable of the following functions:

- **Receiving and acting on classified biological threat intelligence** from DOD, DOE, DHS, and the intelligence community (IC), and translating that intelligence into public health preparedness and response actions within HHS.
- **Providing public health expertise in national security processes**, including participation in interagency threat assessment, the National Security Council biological threat working group, and any classified biological early warning review.
- **Serving as the authoritative interface between civilian biological monitoring data and the national security community**—ensuring that anomalous signals detected through environmental or clinical monitoring are assessed in light of available intelligence, and vice versa.
- **Supporting attribution analysis** in coordination with the IC, DOD, and DOE when an anomalous biological event is detected, including assessing whether a signal is consistent with natural emergence, accidental release, or deliberate use.

- **Maintaining standing relationships and protocols** with DOD's biological defense programs, relevant DOE laboratories, DHS's CWMD office, the FBI's WMD Directorate, and relevant IC components, so that coordination in a crisis does not depend on building relationships from scratch.

Finding 4.3: The existing National Wastewater Surveillance System (NWSS), built during COVID-19, covers approximately 42 percent of the US population and relies primarily on targeted assays for known pathogens.¹⁹³ It has demonstrated the capacity to provide four to seven days' advance warning ahead of clinical surges, but its current architecture cannot detect novel or engineered agents. NWSS was established entirely through emergency supplemental funding and has never received a permanent, dedicated funding base. No stable, multiyear contracting framework exists to sustain or expand this infrastructure.

Recommendation 4.3: Establish a national environmental biological monitoring contracting framework.

- **Funding:** Congress must appropriate **stable (multiyear) and dedicated funding** for environmental biological monitoring services procured through competitive contracts with private sector firms.
- **Local governments:** States and municipalities must be permitted and incentivized to procure local environmental biological monitoring services from qualified private vendors, with federal matching funds available to jurisdictions that meet minimum coverage and data-sharing standards.
- **Data rights:** All contracts must include data rights provisions maintaining federal access to biological monitoring data and clear rules on downstream commercial use.
- **Advanced sequencing:** Contracts must require transition to pathogen-agnostic metagenomic sequencing at all sites, with targeted PCR assays retained for confirmed priority pathogens; technical requirements must be updated on a rolling basis as the state of the art advances.

Finding 4.4: No operational US system currently integrates environmental, clinical, genomic, and behavioral biological monitoring data streams in real time. AI-enabled analytics are the only practical means of processing the data volumes a continuous, pathogen-agnostic national biological monitoring system would generate, and the federal government needs sustained and dedicated procurement authorities to acquire these capabilities from the private sector.

Recommendation 4.4: Procure AI-enabled biological monitoring analytics as a competitive contracted service.

- **Analytics as a service.** The federal government should procure AI-enabled biological monitoring analytics as a service. The analytics should include baseline modeling, anomaly detection, multi-stream data integration, and threat characterization through competitive contracts with private sector firms, using the model established for intelligence community analytic services and DOD data platforms.

193 U.S. Centers for Disease Control and Prevention, "About CDC's Wastewater Monitoring Program," updated July 31, 2026, <https://www.cdc.gov/wastewater/about/index.html>

- **Human review.** All contracts must require human review of anomaly flags before consequential decisions are taken.
- **Prequalification.** The federal government should establish a prequalification process for biological monitoring analytic vendors, evaluating technical capability, data security, and interoperability before firms are eligible for federal contracts. This will create a qualified vendor pool available to cities, states, and federal agencies without duplicating procurement overhead.
- **Fund early-stage innovation.** Early-stage innovation in detection, sequencing, and integration must be supported through the Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) programs.

Finding 4.5: Real-time genomic surveillance capacity is concentrated in a small number of academic and public health laboratory centers, creating geographic gaps and single-point-failure fragility identical to the diagnostic testing failure exposed by COVID-19. A distributed private laboratory network with pre-negotiated surge contracting authority could provide both continuous coverage and rapid expansion capacity.

Recommendation 4.5.1: Expand genomic surveillance through distributed private laboratory networks.

- **Long-term contracts.** The federal government must establish long-term contracts with private genomic sequencing companies and laboratory networks to provide continuous, geographically distributed genomic surveillance coverage, with the goal of next-day sequencing turnaround for priority pathogen samples at sites covering 80 percent of the US population within five years.
- **Surge capacity.** The federal government must establish a standing emergency contracting mechanism with pre-negotiated surge contracts in place before the next outbreak that authorizes rapid expansion of genomic sequencing capacity upon declaration of a public health emergency. Congress can spur this action by mandating it specifically and requiring HHS periodically report on its development and status.

Recommendation 4.5.2: Provide HHS with other transaction authority to test, evaluate, and scale new capabilities quickly. DOD has shown that other transaction authority (OTA)—a contracting mechanism that allows the government to rapidly prototype, research, and acquire innovative technologies outside of normal, lengthy federal procurement timelines¹⁹⁴—can accelerate collaboration with nontraditional contractors and speed capability delivery for cutting edge technology. Congress should provide HHS with a similar authority to rapidly generate and scale the technology needed by the US government and private sector for real-time biological data collection, harmonization, and monitoring at scale.

Finding 4.6: US biological monitoring investment follows a boom-and-bust pattern driven by emergency supplemental appropriations that arrive too late to build the infrastructure needed to detect the emergency that triggered them. Annual discretionary appropriations are structurally incompatible with the continuous operations a national early-warning system requires. The

194 10 U.S.C. §§ 4021-4022.

private-sector-led model proposed in the preceding recommendations is only as durable as its funding mechanism; without stable, multiyear contracting authority, private firms cannot make the capital investments, hire the personnel, or build the systems that early warning demands. For other areas of critical national infrastructure, the US government has used trust funds—such as the Highway Trust Fund and the Airport and Airway Trust Fund—to ensure long-term, stable funding.

Recommendation 4.6: Establish a Biological Early Warning Trust Fund with mandatory multiyear contracting authority.

- **Early warning trust fund.** Congress should establish a Biological Early Warning Trust Fund modeled on other appropriations-funded trust funds, such as the Medicare Supplementary Medical Insurance Trust Fund. The Trust Fund should have dedicated, mandatory appropriations insulated from annual discretionary competition, with funding triggers and matching grants for state and local governments to support surveillance, data sharing, and surge capacity.
- **Emergency resources.** The Trust Fund must include automatic trigger provisions releasing additional emergency resources when prespecified epidemiologic thresholds are crossed, enabling rapid contract expansion without a new appropriation.
- **Long-term contracts.** Contracting agencies must be authorized to issue multiyear contracts of five to ten years for core biological monitoring services, providing the contractual stability necessary for private sector investment in infrastructure and personnel.
- **Matching grants.** Cities and states that meet federal coverage and data-sharing standards should be eligible for Trust Fund matching grants, creating a federated funding architecture that distributes both investment and resilience across levels of government.

Finding 4.7: The biological monitoring architecture proposed in this chapter will not succeed without a binding privacy and data-governance framework. Public trust is an operational requirement for any functioning early-warning system. Although community-level wastewater monitoring generally poses low individual-identification risk, the transition to pathogen-agnostic metagenomic sequencing, smaller-scale sampling, expanded data sharing, and potential national-security uses require enforceable standards to prevent misuse, function creep, stigmatization, and erosion of public trust.

Recommendation 4.7: Establish a binding privacy and data-governance framework for biological monitoring.

- **Technical standards.** Require federally funded metagenomic biomonitoring to use validated privacy-preserving protocols, including host-DNA depletion where feasible and computational removal of human reads, residual human sequence quality control thresholds, independent auditability, and strict limits on retention or sharing of raw reads that may contain human genomic material.
- **Tiered protections by sampling scale.** Require heightened public health justification, data minimization, access controls, and community engagement when monitoring is conducted at smaller scales. Distinguish large community-level sewershed monitoring

from small-site, facility-specific, tribal, correctional, military, campus, vessel, port-of-entry, or other targeted monitoring.

- **Use restrictions.** Biological monitoring data may be used only for purposes directly related to public health, biosecurity, and clearly defined biomedical research and medical countermeasure development. Use for purposes unrelated to these missions, including law enforcement, employment or insurance decision-making, non-public-health profiling, or other punitive uses, should be prohibited absent explicit statutory authorization. Secondary research uses should require independent review, data minimization, and public reporting sufficient to preserve public trust.
- **Data access and stewardship.** Distinguish raw metagenomic reads, processed pathogen calls, aggregate trend data, and public dashboards. Establish data-use agreements, access logs, retention limits, data-minimization rules, and review requirements for secondary research uses or requests for more granular data or samples.
- **Oversight.** Establish an independent advisory body with community, public health, ethics, statistical, tribal, and data-governance expertise to review data-use practices, proposed secondary uses, requests for granular data access, and any expansion of system purposes and report publicly, supported by local community input mechanisms at the point of implementation.
- **Tribal data sovereignty.** Require wastewater monitoring involving tribal communities to be governed by tribal approval, tribal data-stewardship expectations, data ownership and attribution rules, and participatory engagement consistent with tribal sovereignty.
- **National security interface.** Establish clear, publicly disclosed criteria governing the conditions under which biological monitoring data may transition from open public health infrastructure into classified national security channels, with classification decisions subject to independent oversight.

Finding 4.8: International biological monitoring depends on national reporting under the World Health Organization's International Health Regulations (IHR). The IHR obligates member states to report a public health emergency of international concern but provides no enforcement mechanism. China's delayed reporting on COVID-19 was the predictable result of a system in which economic and political incentives favor delayed disclosure. A detection architecture that depends on voluntary good-faith reporting by other nations is structurally vulnerable. **Building a meaningful international biological monitoring capacity requires independent detection capability that does not rely on host-country disclosure, as well as sustained investment in global data-sharing infrastructure that incentivizes voluntary participation.**

Recommendation 4.8: Build independent international biological monitoring capacity and integrate it with national security decision-making.

- **Partnerships.** The federal government should invest in bilateral and multilateral partnerships that build environmental monitoring and genomic surveillance capacity in high-risk regions, reducing dependence on national reporting systems subject to political and economic pressure to delay disclosure.

- **Funding for data sharing.** The federal government should provide contracts or grants to organizations with demonstrated operational capacity to build and sustain global genomic data-sharing infrastructure. This should include attribution-based sharing models that have proven most effective at incentivizing submission from scientists and governments in low- and middle-income countries.
- **Integrated national threat assessment capability.** The intelligence community's biological threat assessment capabilities must be integrated with international monitoring outputs. The government must establish protocols that ensure anomalous biological signals detected abroad reach national security decision-makers quickly and that attribution analysis (discussed in the next chapter) is initiated in parallel with public health response. With allies and partners, building a secure continuous biological monitoring data-sharing capacity will also enhance preparedness.

Attribution

Being able to credibly attribute responsibility for a deliberate biological attack or accidental release is the third component of building a layered biological deterrent, adding to prevention and detection. Accurate attribution is crucial for pinpointing the origin of a biological event, determining whether it is accidental or deliberate, countering additional attacks, and developing appropriate medical countermeasures. Historically, attribution has proven both time-consuming and difficult. It took nearly twenty years for the 1979 release of anthrax at the former Soviet Union biological weapons plant at Sverdlovsk to be attributed to a laboratory accident.¹⁹⁵ Nearly a decade transpired between the 2001 Amerithrax attacks and their attribution to Bruce Ivins in 2008.¹⁹⁶ Most recently, the origin of COVID-19 remains contested today, more than six years after its emergence.¹⁹⁷

Today, the same technology that has increased the threat of engineered biology is also improving our ability to detect and attribute those pathogens. Machine learning and AI hold enormous promise for enabling a more rapid and credible biological attribution system. The promise, however, has yet to be fulfilled. Although researchers are exploring the use of AI for attribution,¹⁹⁸ these efforts have not yet resulted in an effective national attribution capability. To achieve a bioresilient future, the United States must accelerate attribution science so that it can routinely and rapidly identify new pathogens amid a continuous stream of circulating disease threats, and determine whether and where a novel pathogen could have been engineered. Attribution can also illuminate where, whether, and how an emerging pathogen is affecting people, animals, and plants around the world, which can contribute to the development of more rapid and precise countermeasures. The expansion of AI capabilities places these goals within reach, but achieving them will require a focused national effort and renewed international partnerships.

A. Attribution as Deterrence

Robust attribution capabilities can play a powerful deterrent role in the emerging landscape of AI-enabled biotechnology. As discussed in previous chapters, advances in synthetic biology and AI-driven biological design tools may further reduce the technical barriers to engineering novel pathogens, increasing the risk that actors believe they can deploy novel pathogens covertly or through plausible deniability. Developing credible attribution systems changes this calculus. If

¹⁹⁵ Meselson et al., “Sverdlovsk Anthrax Outbreak.”

¹⁹⁶ FBI, “Amerithrax Investigation.”

¹⁹⁷ University of Nebraska Medical Center, “What We Know About the Origin of COVID-19 and What Remains a Mystery,” *Global Center for Health Security*, Apr. 11, 2023, <https://www.unmc.edu/healthsecurity/transmission/2023/04/11/what-we-know-about-the-origin-of-covid-19-and-what-remains-a-mystery/>.

¹⁹⁸ Oliver M. Crook, Anemone Franz, and Aaron Maiwald, “Tracing Engineered Biothreats with AI Forensics: Five Steps to Improve Attribution,” *Bulletin of the Atomic Scientists*, Dec. 1, 2025, <https://thebulletin.org/2025/12/tracing-engineered-biothreats-with-ai-forensics-five-steps-to-improve-attribution/>.

governments can reliably identify the origin of a pathogen, including the laboratory methods, synthesis platforms, computational models, or supply-chain signatures associated with its creation, then potential perpetrators must assume that deployment will carry a substantial risk of exposure and punishment. In this sense, attribution capability functions analogously to forensic capabilities in cyber and nuclear regimes¹⁹⁹: **Its value lies not only in post-incident investigation, but in preventing attacks from occurring in the first place.**

For the United States government, investment in attribution therefore yields both defensive and strategic returns. Domestically, attribution capabilities strengthen biodefense preparedness and accelerate incident response by helping distinguish accidents, natural outbreaks, and deliberate attacks. Internationally, they support norm enforcement and enable coalition-based responses to violations. Because deterrence depends on credibility, these systems must be operationally mature before a major incident occurs. Additionally, demonstrated capability matters as much as technical capability itself. Publicly signaling that the United States possesses sophisticated attribution mechanisms can shape adversary expectations and reinforce deterrence by generating uncertainty about intended effects and unintended consequences. Actors are less likely to undertake high-consequence biological attacks if they cannot confidently predict that attribution systems will fail. Investment in attribution science, international data-sharing frameworks, secure biological provenance systems, and AI-enabled forensic analysis is therefore not simply a research priority but a strategic national security imperative.

B. Current Challenges for Attribution

Although the United States has funded research on microbial forensics and attribution policy for decades,²⁰⁰ in practice it has proven challenging to determine the origin of pathogens and toxins. Attributing a specific biological incident to nature, a laboratory, or an actor requires at least three components:

- early samples of the emerging pathogen or toxin,
- a clear trajectory of the agent's emergence and spread, and
- reference data or materials that allow for the determination about whether a pathogen has been artificially created or engineered.

Each of these components requires intense and integrated cooperation among health, security, and emergency response personnel, all of whom compete for access to scarce samples and information to save lives, analyze origins, stop spread, deter a follow-on attack, and develop and deliver countermeasures.

Three issues compound these requirements: First, as discussed in the previous chapter, identifying key indicators that a biological incident is deliberate depends on capabilities the United States

199 National Academies of Sciences, Engineering, and Medicine, *Biodefense in the Age of Synthetic Biology* (National Academies Press, 2018); Jon R. Lindsay, "Tipping the Scales: The Attribution Problem and the Feasibility of Deterrence Against Cyberattack," *Journal of Cybersecurity* 1, no. 1 (2015): 53–67.

200 National Research Council, Committee on Research Standards and Practices to Prevent the Destructive Application of Biotechnology, *Biotechnology Research in an Age of Terrorism* (National Academies Press, 2004), <https://www.ncbi.nlm.nih.gov/books/NBK234883/>.

cannot yet reliably count on.²⁰¹ Indicators include: use of a novel or engineered agent; unusual patterns of disease and death across humans, animals, and plants; unusual disease characteristics; and the existence of a credible threat. Current systems fall short on each dimension. The US attribution system was originally geared toward anthrax and ricin, agents previously used by terrorists and states, and known threats. Systems like BioWatch, developed specifically for early detection of a deliberate biological attack, have not kept up with the range of pathogens, targets, or deployment methods that could be used now or in the future. Attacks conducted as part of hybrid warfare present a particular attribution challenge: They may be deliberately designed with false flags to obscure the adversary's identity, while pursuing objectives such as sowing fear, disrupting economies, or gaining battlefield advantage.

Addressing these issues requires a focused and synergized national effort across three central challenges:

1. Detecting and fingerprinting AI-designed pathogens;
2. Building a robust biological baseline; and,
3. Achieving effective interagency (federal, state, and local) coordination and decision speed.

Challenge 1: Detecting and Fingerprinting AI-designed Pathogens

AI is simultaneously the source of the attribution problem and its most promising solution. The same tools that enable the design of novel, unrecognizable pathogens are also generating new methods to detect genetic engineering signatures and trace them to a laboratory of origin. Experts have begun using AI tools to identify genetic changes such as insertions and subtle replacements of genetic information that can be matched to a specific laboratory of origin.²⁰² Design choices made by the laboratories enable attribution because they leave behind “genetic fingerprints.”²⁰³ These genetic fingerprints are readily identifiable: In a proof-of-concept challenge conducted in 2020, a machine-learning system identified the correct laboratory of origin for engineered plasmid sequences in 81.9 percent of cases.²⁰⁴ More recent work demonstrates how these capabilities have evolved. In 2026, researchers reported PlasmidGPT, a generative AI model trained on more than 153,000 engineered plasmids, that could recognize patterns associated with particular laboratories and also help analyze and design engineered DNA sequences.²⁰⁵

Using these novel AI tools, a modernized US attribution system must be able to determine whether a given sample is engineered and, where sufficient reference data exist, identify which laboratories

201 Tracey Rissman and Annette Prieto, *Attributing Biological Weapons Use: Strengthening Department of Defense Capabilities to Investigate Deliberate Biological Incidents* (RAND Corporation, 2024), https://www.rand.org/pubs/research_reports/RR2360-1.html.

202 Ethan C. Alley et al., “A Machine Learning Toolkit for Genetic Engineering Attribution to Facilitate Biosecurity,” *Nature Communications* 11 (2020): 6373, <https://doi.org/10.1038/s41467-020-19612-0>; William Mo, Christopher A. Vaiana, and Chris J. Myers, “The Need for Adaptability in Detection, Characterization, and Attribution of Biosecurity Threats,” *Nature Communications* 15 (2024): 10699, <https://doi.org/10.1038/s41467-024-55436-y>.

203 Crook, Franz, and Maiwald, “Tracing Engineered Biothreats with AI Forensics.”

204 Oliver M. Crook et al., “Analysis of the First Genetic Engineering Attribution Challenge,” *Nature Communications* 13 (2022): 7374, <https://doi.org/10.1038/s41467-022-35032-8>

205 Bin Shao et al., “PlasmidGPT: A Generative Framework for Plasmid Analysis and Generation,” *Science Advances* 12 (2026): eae6916, <https://doi.org/10.1126/sciadv.aee6916>

and/or engineering methods are associated with its creation. Pathogen sequences can contain both deliberate and unintentional markers that can indicate which laboratory engineered them. Laboratories purposefully include sequences that identify a sample as their own. This practice, called watermarking, allows responsible actors to track their own tools, techniques, and samples. Unintentional markers can also be found in the ways different laboratories stitch together their pathogen building blocks. The Finding Engineering-Linked Indicators (FELIX) program, created by the Intelligence Advanced Research Projects Activity (IARPA), pioneered research to match laboratories with genetic products that they had developed or engineered.^{206, 207, 208} Researchers have continued to sharpen these attribution tools, including by using AI to categorize publicly available engineered sequences and predict their origin.²⁰⁹

As AI-enabled biotechnology continues to evolve, it is not yet clear whether new AI techniques could also make it possible to mask these genetic identifiers. If the United States does not have the capacity in place to investigate the genetic origin of a sample, it will be far harder to evolve the attribution system alongside the advancing technology to hide sources. **The investments in developing advanced attribution capabilities need to be made now, before masking techniques mature.**

How Genetic Engineering Attribution Works

Genetic engineering leaves traces. When a laboratory modifies a pathogen's genetic sequence by inserting foreign genes or making subtle replacements of individual base pairs, the specific techniques it uses create identifiable patterns in the result. These patterns reflect the laboratory's standard toolkit and normal methods: which molecular building blocks it prefers, how it joins DNA fragments, and what shortcuts its researchers routinely take. AI systems can be trained to recognize these patterns, effectively reading a laboratory's fingerprint in an engineered sequence the way a forensic examiner reads handwriting.

The primary training data for these systems comes from plasmids, which are small circular bacterial DNA molecules widely used in genetic engineering research. The global plasmid repository Addgene currently hosts more than 1,000 viral vectors, the engineered viral particles most commonly modified for research and medical applications. The greatest limitation on AI capability is that list-based AI attribution systems can only identify laboratories whose sequences appear in their training data. The laboratories that pose the greatest biosecurity risk, those working with dangerous pathogens at Biosafety Level 3 and higher, are largely absent from public repositories like Addgene. Developing a capable attribution system for this tier will require scaling current datasets by one to two orders of magnitude, with strict access controls to prevent misuse of the resulting data.

206 Intelligence Advanced Research Projects Activity, "Finding Engineering-Linked Indicators (FELIX)," accessed Apr. 7, 2026, <https://www.iarpa.gov/research-programs/felix>.

207 Crook, Franz, and Maiwald, "Tracing Engineered Biothreats with AI Forensics."

208 Draper, "IARPA, Ginkgo Bioworks and Draper Announce New Technologies to Detect Engineered DNA," Oct. 17, 2022, <https://www.drapr.com/media-center/news-releases/detail/23491/iarpa-ginkgo-bioworks-and-drapr-announce-new-technologies-to-detect-engineered-dna>.

209 Mo, Vaiana, and Myers, "The Need for Adaptability in Detection, Characterization, and Attribution of Biosecurity Threats."

Challenge 2: Building a Robust Biological Baseline

Biological attribution requires a continuous, population-scale biological monitoring baseline against which anomalous signals can be identified.²¹⁰ Without this baseline, an engineered pathogen emerging in a poorly monitored population is difficult to distinguish from a gap in data collection. The design requirements for an attribution-capable bio-monitoring system are more demanding than those for public health response alone: Data must be generated with chain-of-custody protections and in formats that can be integrated with law enforcement and intelligence analyses. The US bio-monitoring architecture does not currently meet this standard.²¹¹ **This section addresses the specific gaps most consequential for attribution; the broader architecture of national bio-monitoring is addressed in Chapter 4.**

The United States has significantly reduced its baseline monitoring capacity. In 2025, HHS withdrew \$11.4 billion from state and local public health funding; though states have challenged the funding reductions in court, in many cases the funding has not been restored.²¹² This funding shortfall has substantially diminished the data flows that attribution depends upon. For example, \$8.9 billion of the funding withdrawn was tied to Epidemiology and Laboratory Capacity (ELC) grants for prevention and control of emerging infectious diseases. The ELC grants have provided the primary federal funding underwriting state and local disease monitoring, outbreak detection, and laboratory testing infrastructure.²¹³ Across the US, the terminated ELC awards were supporting data system modernization, laboratory capacity, electronic case reporting for time-sensitive infectious disease outbreaks, and improved H5N1 and measles testing capability. Health departments that documented the consequences of this disruption of federal funding reported losing people and capacity focused on COVID-19, RSV, and influenza monitoring, including syndromic surveillance—the system that gives health officials near-real-time data on why people are seeking care—as well as in the informatics staff needed to send and receive electronic laboratory and epidemiological data, and in the staff dedicated to outbreak response for threats like mpox and avian influenza.²¹⁴

Two of these bio-monitoring platforms are particularly consequential for attribution. First, national wastewater monitoring using **metagenomic sequencing** can detect novel biological signals that other targeted monitoring misses. Metagenomic sequencing generates sequence data capable of supporting attribution analysis. SecureBio's Nucleic Acid Observatory (NAO)²¹⁵ is scaling this type of capability, but the technology has not been deployed nationally nor funded as an integrated capability across state and local jurisdictions. Second, **diagnostic surge capacity** for testing during early outbreak phases creates the rich data environment from which unusual patterns can be distinguished. To ensure that these platforms are designed and implemented to support both

210 Ibid.; Crook, Franz, and Maiwald, "Tracing Engineered Biothreats."

211 Crook, Franz, and Maiwald, "Tracing Engineered Biothreats with AI Forensics."

212 Josh Michaud and Jennifer Kates, "Tracking Key HHS Public Health Policy Actions Under the Trump Administration," KFF, last modified Mar. 19, 2026, <https://www.kff.org/other-health/tracking-key-hhs-public-health-policy-actions-under-the-trump-administration/>.

213 A federal court blocked the terminations for twenty-three states and the District of Columbia, which sued, ordering the funds reinstated while litigation proceeds. That order does not extend to non-plaintiff jurisdictions, where the same category of ELC-funded surveillance and lab-capacity work remains cut or only partially restored. *Colorado v. US Department of Health and Human Services*, No. 1:25-cv-00121-MSM-LDA (D.R.I. 2025).

214 Network for Public Health Law, "Updates to HHS Restructuring and Funding Cuts: Impact on State and Local Public Health," Apr. 3, 2025, <https://www.networkforphl.org/news-insights/updates-to-hhs-restructuring-and-funding-cuts-impact-on-state-and-local-public-health/>; Connecticut Public Health, "DPH Federal Funding Updates," Mar. 27, 2025, <https://www.housedems.ct.gov/sites/default/files/Federal%20Funding%20Changes.pdf>.

215 Nuclear Threat Initiative, "NAO Observatory," accessed June 7, 2026, <https://naobservatory.org/>.

public health and biosecurity, **explicit attribution-design requirements need to be built in from the outset.**

Challenge 3: Achieving Interagency Coordination and Decision Speed

Successfully identifying the source of any biological event requires integrated information and cooperation among health, intelligence, defense, agricultural, and environmental agencies and communities. These partnerships and the data-sharing arrangements they enable take years to build. For a variety of reasons, they can be disrupted at the moment of crisis when they are most urgently needed.

Competing institutional priorities have made biological attribution difficult. During an emerging pandemic, law enforcement, public health, and the private sector have conflicting objectives. Law enforcement's imperative to preserve samples and data for evidentiary purposes is frequently pitted against the public health imperative to broadly share that same information to create countermeasures and ensure healthcare providers can act fast to save lives. Without close coordination among HHS, DHS, DOD, and FBI, these trade-offs can slow an investigation, impede a public health response, and delay attribution.

To overcome these challenges, the US government has spent decades developing and refining interagency protocols. One example is the Biological Incident Notification and Assessment Protocol outlined in the Biological Incident Annex to the Federal Emergency Management Agency's (FEMA) Response and Recovery Federal Interagency Operational Plan.²¹⁶ The protocol calls for a "lean forward" posture: when a new biological event emerges, health, security, and defense officials convene immediately and proactively seek information about whether an emerging biological event is deliberate or accidental. It also provides a venue for adjudicating competing interests, and enables the data sharing required to both protect and investigate evidence and to make information available to save lives.

No single federal authority is currently charged with overseeing these interagency protocols. Coordination has historically depended on senior White House leadership and career civil service experts with deep institutional knowledge of biological crisis response. Reduced National Security Council biosecurity focus and expertise, along with substantial reductions in the relevant federal workforce, have diminished the human capacity and infrastructure these response protocols need to function effectively.

C. Capability Inventory: Where the United States Stands Today

A robust attribution framework requires a range of specific capabilities, some of which exist in limited form today, and some of which will need to be built. A list of the most critical capabilities is below, along with a brief description of the status of the capability.

²¹⁶ Federal Emergency Management Agency [FEMA], *Biological Incident Annex to the Response and Recovery Federal Interagency Operational Plans* (US Department of Homeland Security, 2023), https://www.fema.gov/sites/default/files/documents/fema_incident-annex_biological.pdf.

Capability	Status
Rapidly detect any existing biological agent in samples from humans, animals, plants, wastewater, air, and soil. This will require national, state, and city-level programs that can routinely find circulating pathogens like measles, influenza, COVID, and RSV, as well as newly emerging and engineered threats.	This capability currently exists in fragmented forms. It is not well-funded, integrated, or inclusive of the state and local capacity needed to effectively use early warning data to take specific actions to save lives, stop disease spread, and attribute origin.
Identify signs of laboratory or AI-enabled creation by comparing samples from humans, animals, plants, wastewater, air, and soil with sequence data from laboratories around the world known to work with related samples.	This capability is in early stages of research and development, with limited geographic reach and open questions about capacity to attribute fully AI-generated agents.
Identify unusual signatures that do not match known pathogens in samples from humans, animals, plants, wastewater, air, and soil. This will require a greater understanding of pathogenic proteins that can be designed using AI.	This capability does not currently exist. SecureBio's Nucleic Acid Observatory has begun to utilize metagenomic sequencing for wastewater sampling to develop the tools necessary to find new pathogens quickly, but this technology has not been adopted nationally nor funded as an integrated capability across the nation.
Rapidly and securely share and integrate datasets from across public health, agriculture, law enforcement, and intelligence. To rapidly determine whether an outbreak of unknown origin was caused by a deliberate attack and to attribute its origin to a specific location, laboratory, or actor, the US government will require a system that can acquire, store, and analyze signatures of genetic engineering and compare them with known intelligence and law enforcement data on a regular basis.	This capability does not currently exist. There is no nationwide US biological alert system that can quickly integrate and analyze information from existing genetic engineering databases, clinical samples, wastewater sampling, global alerts, and law enforcement and intelligence data.
Partner with allies and institutions internationally to gain access to data, samples, and information, and to practice attribution at scale. The United States cannot develop this capacity alone. A robust biological attribution system will rely on datasets from partner laboratories around the world. To achieve this goal, the United States will require early warning capacity and regional information sharing across networks (including multilateral networks with the World Health Organization (WHO), the Food and Agriculture Organization (FAO), the West Africa Health Organization (WAHO), and the Africa Centres for Disease Control and Prevention (Africa CDC), and revitalized international investigative capacity through the United Nations Secretary-General's Mechanism for Investigation of Alleged Use of Chemical and Biological Weapons (UNSGM).	This capability is rapidly degrading. The withdrawal of CDC and USAID personnel from foreign countries and the US withdrawal from WHO has made it harder to access the information needed for rapid attribution. Bilateral global health security arrangements under development with lower- and middle-income countries have not been made public and do not appear to incentivize the kind of regional information sharing necessary to quickly identify patterns signifying a dangerous emerging pathogen, much less a deliberate biological weapons attack.

D. The US Policy Context

US biodefense policy has consistently prioritized attribution as a core component of the national biodefense posture. The 2018 National Strategy for Countering Weapons of Mass Destruction Terrorism asserted that deterring hostile states and individuals requires the means to quickly and accurately attribute an attack, achieved through traditional intelligence collection, law enforcement, and advanced scientific tools.²¹⁷ The 2022 National Biodefense Strategy reinforced this, prioritizing biological threat assessment and characterization—including enhancing national capacity to assess and characterize biological threats within one week of acquiring a suitable sample, and establishing foreign partnerships to recognize, interdict, and attribute responsibility for the use of biological weapons.²¹⁸

Within the current administration, the Department of Defense is leading the effort to pursue this critical objective. Its Biodeterrence Policy identifies attribution as a key element of a strategy of deterrence by denial and deterrence by punishment.²¹⁹ Elements of this work are funded via the 2026 National Defense Authorization Act and the Intelligence Authorization Act, including improving organizational and functional capabilities.²²⁰

Building on the Cold War mindset of deterring adversarial use of weapons of mass destruction, including bioweapons, the United States has traditionally led initiatives with international partners to improve global capacity to collect data and investigate the source of a deliberate biological weapons attack or outbreak of unknown origin. Relatedly, in 1987, the UN General Assembly established the United Nations Secretary-General's Mechanism for Investigation of Alleged Use of Chemical and Biological Weapons (UNSGM) to deter and attribute deliberate biological and chemical attacks. This initiative allows for a fact-finding mission to investigate the site of an alleged biological weapons incident. However, because the UNSGM is geared toward investigating events that appear deliberate, it has not dedicated equal attention to investigations of outbreaks of unknown origin or of a novel nature.²²¹ In practice, the UNSGM has been invoked only three times to investigate alleged chemical weapons attacks and has never been used to investigate an alleged biological attack.

217 US Department of Defense, *National Strategy for Countering Weapons of Mass Destruction Terrorism* (Department of Defense, 2018), https://www.acq.osd.mil/ncbdp/docs/20181210_National-Strategy-for-Countering-WMD-Terrorism.pdf.

218 The White House, *National Biodefense Strategy and Implementation Plan for Countering Biological Threats, Enhancing Pandemic Preparedness, and Achieving Global Health Security* (October 2022): §§ 1.1.4, 2.2.4, <https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/10/National-Biodefense-Strategy-and-Implementation-Plan-Final.pdf>.

219 Melissa Tune, “Dr. Robert Kadlec Unveils National Biodeterrence Efforts at DTRA-NDU Report Rollout,” Defense Visual Information Distribution Service, Mar. 12, 2026, <https://www.dvidshub.net/news/560774/dr-robert-kadlec-unveils-national-biodeterrence-efforts-dtra-ndu-report-rollout>; National Defense University Institute for National Strategic Studies, *A Framework for Biological Weapons Deterrence*, Mar. 16, 2026, <https://inss.ndu.edu/Media/News/Article/4435123/a-framework-for-biological-weapons-deterrence/>.

220 National Security Commission on Emerging Biotechnology, *Biotechnology Breaks Through in FY 2026 National Defense Authorization Act*, Dec. 9, 2025, <https://www.biotech.senate.gov/press-releases/biotechnology-breaks-through-in-fy-2026-national-defense-authorization-act/>.

221 In 2021, following largely unsuccessful attempts to obtain early data and samples from China regarding the origins of SARS-CoV-2, the World Health Organization convened the Scientific Advisory Group for the Origins of Novel Pathogens (SAGO). SAGO recommended studies necessary to gather evidence about the origins of SARS-CoV-2 and, more broadly, the origins of emerging and reemerging future epidemics and pandemics. Smriti Mallapaty, “WHO Abandons Plans for Crucial Second Phase of COVID-Origin Investigation,” *Nature*, Feb. 3, 2023, <https://doi.org/10.1038/d41586-023-00283-y>; Smriti Mallapaty, “The Search for the Origins of COVID-19 Continues amid Political Tensions,” *Nature*, Sept. 29, 2021, <https://doi.org/10.1038/d41586-021-02813-y>; World Health Organization, *Preliminary Report of the Scientific Advisory Group for the Origins of Novel Pathogens (SAGO)* (World Health Organization, 2022), <https://cdn.who.int/media/docs/default-source/scientific-advisory-group-on-the-origins-of-novel-pathogens/sago-report-09062022.pdf>.

E. Conclusion

Attribution is the bridge between detection and response. Knowing that a biological event has occurred is not enough; knowing whether it was deliberate, and who is responsible, determines what kind of response is possible and whether deterrence has any ongoing relevance. A nation that cannot attribute a pathogen to its source cannot credibly create meaningful consequences via law enforcement or military responses and cannot effectively adapt its defense and deterrence postures against future attacks.

Attribution is also where the deterrence by resilience cycle closes. Prevention constrains what adversaries, both natural and human, can do. Detection finds what prevention misses. Attribution identifies the actor, enabling accountability, informing countermeasures, and signaling to adversaries that using bioweapons (or mishandling pathogens) carries a price. **The response and recovery chapters that follow are interdependent with the credibility of that signaling to adversaries.** Building the attribution capabilities the United States currently lacks is not just a technical project; it is a strategic imperative.

Findings and Recommendations

Finding 5.1: The convergence of AI and biotechnology has outpaced US biological attribution science. Detection methods calibrated to known pathogens cannot identify agents designed to evade them. Research programs capable of matching engineered agents to their laboratory of origin—including IARPA's FELIX program—have demonstrated proof-of-concept but have not been transitioned into operational capabilities. If the United States does not build and sustain these capabilities now, it will be unable to adapt them as masking techniques develop alongside AI-enabled biodesign.

Recommendation 5.1: Congress should authorize and fund the acceleration and transition of IARPA FELIX and successor programs from research into a standing operational capability.

This should include: development of a national engineering-signature library covering laboratories known to work with dangerous pathogens; expansion of AI-based attribution tools capable of identifying fully AI-generated agents; and integration of these tools with existing intelligence and law enforcement databases to enable rapid comparison against known actors and methods.

Finding 5.2: Biological attribution requires a continuous, population-scale surveillance baseline against which anomalous signals can be identified. The United States currently lacks this baseline. Significant reductions in federal public health funding have degraded already limited state and local surveillance capacity. Existing surveillance platforms are not designed to generate data with the chain-of-custody and integration requirements necessary to support law enforcement and intelligence analysis. Without a surveillance architecture built for attribution as well as public health response, the United States cannot reliably distinguish an engineered pathogen from a data gap.

Recommendation 5.2: Congress should authorize and fund a national integrated biological attribution system. Core requirements of the system include: AI-enabled anomaly detection for both known pathogens and novel biological signals; and explicit chain-of-custody and data-sharing

protocols designed to support attribution as well as public health response. Diagnostic surge capacity for pandemic-capable pathogens should be a funded component, with federal guidance requiring states and cities to share de-identified healthcare data to support real-time national disease pattern analysis. This attribution system must be integrated into the national biological monitoring system recommendations identified in Chapter 4.

Finding 5.3: The interagency protocols designed to enable rapid biological incident response depend on White House leadership and a career civil service with institutional knowledge of biological crisis response. No single federal authority is currently charged with overseeing these protocols, and it is unclear whether they continue to be exercised. The elimination of dedicated White House biosecurity expertise and coordination capacity, and significant reductions in the federal public health workforce across multiple agencies, have materially reduced the US government's ability to execute rapid, coordinated attribution under crisis conditions.

Recommendation 5.3: Biosecurity coordination should be strengthened.

- Congress should codify requirements for a regularly updated Biological Incident Notification and Assessment Protocol in statute and designate a senior White House official with clear authority to convene and direct interagency biological incident response.
- The administration should restore dedicated biosecurity coordination capacity at the National Security Council and mandate regular interagency exercises with state and local partners to ensure the protocol can be executed under crisis conditions. These exercises should explicitly rehearse the adjudication of competing public health and law enforcement priorities to preserve both response speed and evidentiary integrity.

Finding 5.4: Rapid biological attribution at scale depends on international data-sharing arrangements, sample access, and investigative cooperation built over decades through CDC, USAID, and WHO engagement. The systematic withdrawal from these partnerships is eroding US access to the international biological data and cooperation that attribution requires. The UNSGM, the principal international instrument for investigating alleged biological weapons use, has never been invoked for a biological event and lacks standing technical capacity to do so effectively.

Recommendation 5.4: The administration should revitalize global health and global health security capabilities.

- CDC and US State Department global health security personnel (and USAID personnel, should the agency be reestablished by a future administration) should be deployed as an integral component of US biodefense capacity, and should reengage with WHO to rebuild data-sharing arrangements that support rapid attribution.
- Congress should fund revitalization of the UNSGM as an investigative mechanism for biological events of unknown origin, including building standing technical capacity for rapid biological investigation.
- The administration should negotiate bilateral and multilateral data-sharing agreements that explicitly incentivize building and strengthening regional information-sharing

networks capable of identifying patterns consistent with an emerging biological risk, whether deliberate, accidental, or naturally occurring.

Readiness and Response

Prevention reduces the probability of a biological incident. Detection and attribution ensure that when an event occurs, it can be identified and characterized in time to act. The next two chapters address what the United States must do before, during, and after an incident to ensure that it can respond effectively and recover fully: readiness, response, and resilience.

Together, readiness, response, and resilience form the operational core of deterrence by resilience. They are not sequential stages but mutually reinforcing layers. The stronger each layer, the more credible the deterrent and the more survivable the outcome when deterrence fails. This chapter will focus on the first two, readiness and response. The final chapter will address the path to resilience.

Readiness, Response, and Resilience

Readiness is the precondition that determines whether federal, state, and local governments can promptly and effectively mobilize capabilities when they are needed. A government that has not built these structures and processes before an incident will be very unlikely to have the time and capacity to build them during one.

Response is the action taken when a biological threat presents itself. Response encompasses the interagency decision architecture that translates detection into action, the clinical trial system that generates the treatment evidence guiding care, and the medical countermeasure infrastructure that saves lives and reduces harms.

Resilience is what determines whether the United States can withstand and recover from a biological incident that prevention did not stop and that response could not fully contain and/or effectively mitigate. Resilience is not a static condition to be achieved but a continuous investment that must be made.

A. Readiness

Readiness determines what happens once a threat is detected: how quickly the United States can field a response, at what scale, and at what degree of effectiveness. Readiness is the least visible and most chronically underfunded layer of deterrence by resilience. In an era of AI-enabled biology, it is urgently in need of structural overhaul.

The central challenge of readiness is institutional rather than technical. The United States has exceptional biotechnology firms, world-class research infrastructure, and substantial federal

health emergency authorities. What it lacks is the capacity to mobilize those tools rapidly, consistently, and at a nationwide scale.

The government levers that exist today were built for a time before COVID and before bioconvergence. We react to declared emergencies rather than preparing for anticipated ones. Today we stockpile finished pharmaceuticals rather than maintain surge capacity to produce new ones. Finally, we have yet to build a sustained framework for biosecurity collaboration with the frontier AI industry, whose systems are increasingly involved in both the biological threat and the response to biological incidents.

This section addresses three structural gaps in readiness to respond to biological threats:

- Biosecurity is not designated as a unified critical infrastructure category but is split across several sectors and miscategorized as a “support” capability;
- There is insufficient standing reserve of private-sector biological capacity; and
- There is no effective and durable public-private engagement model for the AI-enabled biotechnology industry.

Taking action to address each gap is critical to achieving national readiness for the landscape of current and emerging risks.

Establishing the Bio Ecosystem as Critical Infrastructure

Persistent readiness in an era of bioconvergence requires ongoing cooperation between the federal government and the private sector. Consistent public-private engagement ensures information is shared across different industries and the government on a continual basis in order to enable each side to have the full visibility required to inform mobilization in response to emerging threats. Today, coordination happens across sixteen sectors formally designated as “critical infrastructure,” which is defined by statute as the “systems and assets ... so vital to the United States” that their “incapacity or destruction ... would have a debilitating impact on security, national economic security, national public health or safety, or any combination of those matters.”²²² The bio ecosystem is not one of those sixteen; despite its vital importance to both economic and national security, it lacks the official critical infrastructure designation.

The critical infrastructure designation for biotechnology would enable classified public-private engagement, sector-wide threat assessments, and coordinated planning between government and industry for crisis prevention, management, and response. This process for the sixteen existing sectors formerly operated through the Critical Infrastructure Partnership Advisory Council (CIPAC), which was established in 2006 and provided designated public-private coordinating bodies exempt from Federal Advisory Committee Act provisions, allowing them to hold confidential, non-public discussions and to prepare for a wide range of risk scenarios, including those associated with foreign adversarial threats. In March 2025, the Department of Homeland Security announced

222 42 U.S.C. § 5195c(e); National Security Memorandum 22, “National Security Memorandum on Critical Infrastructure Security and Resilience,” Apr. 30, 2024 (replacing and enhancing the framework established in Presidential Policy Directive 21 (2013)). The sixteen current sectors include: chemical; commercial facilities; communications; critical manufacturing; dams; defense industrial base; emergency services; energy; financial services; food and agriculture; government services and facilities; healthcare and public health; information technology; nuclear reactors, materials, and waste; transportation systems; and water and wastewater systems.

the dissolution of CIPAC.²²³ Although this did not eliminate existing public-private mechanisms, it did remove the framework for confidential deliberation.²²⁴ According to public reports in 2026, DHS is in the process of replacing CIPAC with the Alliance of National Councils for Homeland Operational Resilience, although its structure and legal protections remain to be determined.²²⁵

Despite the elimination of CIPAC, critical infrastructure coordinating bodies remain. For each designated critical infrastructure, an assigned Sector Risk Management Agency oversees collaboration between the government and the sector. Furthermore, sector coordinating councils (SCC) composed of private sector owners and operators, their trade associations, and other industry representatives interact with a corresponding government coordinating council (GCC). Because the American biotechnology sector has not yet received the critical infrastructure designation, it does not have access to this framework.

Notably, since 2023, biotechnology has been classified as a “support” subsector, housed within the Healthcare and Public Health (HPH) SCC, led by HHS.²²⁶ **Biotechnology, however, is not only support infrastructure for healthcare delivery; it is a primary driver of biological innovation as well as a platform for and potential target of biosecurity risk.**

The bio sector’s absence from the critical infrastructure framework is not an oversight; it reflects an outdated governance architecture that was designed before AI-enabled biology emerged as a national security threat. Today’s critical infrastructure sectors that were designated in 2013 were largely organized around physical systems and traditional threat vectors. The emerging AI-enabled bio risk environment requires a coordination structure that bridges the life sciences industry, frontier AI developers, and the government agencies responsible for national defense and public health. Designating biotechnology as critical infrastructure would enable that to happen.

As a result of these structural issues, there is no existing sector coordination mechanism that captures the full AI-enabled bio risk picture. Because AI developers are not embedded in the HPH sector’s SCC structure, there is no standing mechanism through which government can share threat intelligence with the firms building the systems that adversaries are most likely to exploit.

Over the past several years, two unsuccessful attempts to reform the critical infrastructure system have been made. In 2021, the Department of Homeland Security (DHS) recommended designating

223 DHS, “Notice of Termination of Discretionary Federal Advisory Committees,” *Federal Register* 90, no. 48 (Mar. 13, 2025): 11995.

224 Brian E. Humphreys, *Critical Infrastructure: Emerging Trends and Policy Considerations for Congress*, CRS Report R48878 (Mar. 10, 2026), <https://www.congress.gov/crs-product/R48878>.

225 Justin Doubleday, “Five Updates on the Trump Admin’s Cybersecurity Agenda,” *Federal News Network*, Feb. 4, 2026, <https://federalnewsnetwork.com/cybersecurity/2026/02/five-updates-on-the-trump-admins-cybersecurity-agenda/>.

226 US Healthcare and Public Health Sector Coordinating Council, *Comprehensive Council Charter*, version 3.0 (2023), <https://www.cisa.gov/sites/default/files/2024-06/HPH-SCC-Charter-2023-508.pdf>. The reorganization lists twenty-five biotech and pharmaceutical companies in the support subsector, including Becton, Dickinson and Company; Boston Scientific; Merck & Co.; Thermo Fisher; and Twist Bioscience. Presidential Policy Directive 21 (PPD-21), issued in 2013, established US policy on critical infrastructure security and resilience, designating sixteen critical infrastructure sectors and assigning Sector Risk Management Agencies (SRMAs) to each. In 2024, PPD-21 was displaced by National Security Memorandum 22 (NSM-22), which directed federal agencies to play a more direct role in public-private partnerships while retaining the same sector structure. Under NSM-22, each sector’s private owners and operators interact with the relevant SRMA through a Sector Coordinating Council paired with a Government Coordinating Council under the CIPAC framework. This structure is exempt from the Federal Advisory Committee Act (FACA), enabling timely, consistent, and classified exchange with private sector actors without FACA’s transparency and consultation requirements. In March 2025, Executive Order 14239 directed the national security advisor to review NSM-22 within 180 days and to “recommend to the president the revisions, rescissions, and replacements necessary to achieve a more resilient posture.” As of the publication of this report, that review has not been publicly released.

the “bioeconomy” as a critical infrastructure sector.²²⁷ In 2025, the congressionally mandated NSCEB recommended that Congress ensure the biotechnology sector is categorized as critical infrastructure.²²⁸ Designation would have a host of benefits: It would authorize a designed-for-purpose SCC-GCC pairing (if CIPAC were to be restored); enable systematic classified engagement with industry; clarify risk management responsibilities currently diffused across HHS, DOD, and DHS with no single integrating body; and formally elevate biosecurity from a sub-element of healthcare delivery to a national resilience priority. Each of these elements would contribute to significantly improving the United States’ readiness posture.

Readiness Through a National Strategic Bio Readiness Reserve

Medical Countermeasures

Medical countermeasures are Food and Drug Administration (FDA)-regulated products that may be used during public health emergencies to diagnose, prevent, protect from, or treat conditions associated with emerging infectious diseases or chemical, biological, radiological, and nuclear (CBRN) threats.

Readiness requires the capability to rapidly access medical countermeasures to biological threats. The United States can substantially accelerate the speed with which it responds to biological threats by establishing a National Strategic Bio Readiness Reserve (NSBRR). A strategic reserve is a pre-negotiated commitment of private-sector capacity held in reserve for government activation during a national emergency. Importantly, the reserve model is neither stockpiling specific materials nor does it require government ownership or control of the underlying assets. Instead, it converts the government’s procurement power into a standing relationship with industry, structured through contracts executed in advance—before crises occur. In exchange for committing a defined share of their capacity in the event of a future need, participating firms receive guaranteed government procurements and priority contracting regardless of whether the reserve is ever activated. Critically, this readiness also contributes to deterrence.

Strategic reserves have proven valuable in a variety of sectors and circumstances by ensuring reliable US government access to materials and assets when urgently needed and providing a structured ongoing commercial benefit to participating entities in return. For example, the Civil Reserve Air Fleet (CRAF) has provided DOD guaranteed use of commercial airlift capacity since 1951. The Voluntary Intermodal Sealift Agreement does the same for commercial shipping. The Commercial Augmentation Space Reserve (CASR) is extending the model to commercial space capabilities.²²⁹

227 US Department of Homeland Security, *FY 2021 National Defense Authorization Act Section 9002(b) Report*, Nov. 21, 2021, [https://www.nextgov.com/media/section_9002_ndaa_report_final_11.15.21_1\).pdf](https://www.nextgov.com/media/section_9002_ndaa_report_final_11.15.21_1).pdf); Mariam Baksh, “CISA Highlights Space, Bioeconomy as Possible New Critical Infrastructure Sectors,” *Nextgov*, Nov. 16, 2022, <https://www.nextgov.com/cybersecurity/2022/11/cisa-highlights-space-bioeconomy-possible-new-critical-infrastructure-sectors/379818/>.

228 NSCEB, *Charting the Future of Biotechnology*, § 2.4.

229 Voluntary Intermodal Sealift Agreement, 46 U.S.C. § 57101 et seq.; Maritime Security Program, 46 U.S.C. §§ 53101–53111; Voluntary Tanker Agreement Program, 87 Fed. Reg. 67540 (Nov. 7, 2022), <https://www.federalregister.gov/documents/2022/11/07/2022-24184/renewal-of-the-voluntary-tanker-agreement-program-revised-form-of-the-voluntary-agreement>; “Commercial Augmentation Space Reserve,” discussed in *Defense News*, Nov. 22, 2024.

The CRAF Model

The Civil Reserve Air Fleet (CRAF), established in 1951 following the Berlin Airlift, demonstrates that a voluntary, pre-negotiated reserve of private-sector capacity can remain dormant for decades and be activated reliably in a crisis.²³⁰ Participating airlines commit specified aircraft and crews to DOD use when the program is activated, with response timelines defined by activation stage.²³¹ In exchange, participants receive minimum annual guaranteed airlift contracts and priority access to peacetime DOD contracts—benefits that flow regardless of whether CRAF is ever activated. CRAF has been activated three times: Desert Shield and Desert Storm (1990-1991), Operation Iraqi Freedom (2002-2003), and Operation Allies Refuge from Afghanistan (August 2021).²³²

Two features of the CRAF model apply directly to biodefense. First, participation is voluntary but structured through contracts executed before a crisis, converting the government's procurement power into a credible reserve without requiring permanent public ownership of the underlying assets. Second, incentives are ongoing. Firms perceive commercial value from participation, independent of whether a crisis emerges. This combination of obligation and commercial benefit sustains participation across the potentially long intervals between activations.

The need for a similar arrangement with the biotech and pharmaceutical industry was demonstrated during the COVID-19 pandemic. COVID-19 vaccine candidates were being designed in days and approved in months, but industrial production in the quantities needed to vaccinate the entire nation lagged for a year. Operation Warp Speed, drawing on Defense Production Act (DPA) authorities and Biomedical Advanced Research and Development Authority (BARDA) funding, dramatically accelerated development of countermeasures. But increasing production required ad hoc negotiations, improvisational contracting, and extensive federal coordination that measurably delayed the US response.

In a future where AI-enabled design compresses the interval between pathogen identification and countermeasure design, the time it takes to move from discovery to large scale production will be the major weakness in the ad hoc system. There are several existing programs to provide medical countermeasures, but none meet the current need. The Strategic National Stockpile (SNS) maintains government-owned inventories of finished medical countermeasures; it does not ensure that the United States can rapidly expand production of new countermeasures.²³³ BARDA funds advanced development and procurement but does not protect industrial capacity in advance for national security needs.²³⁴

230 10 U.S.C. §§ 9511 et seq. Carriers must maintain a minimum commitment of 30 percent of their CRAF-capable passenger fleet and 15 percent of their cargo fleet; aircraft must be US-registered; carriers must commit at least four crews per aircraft. Congressional Research Service [CRS], *Civil Reserve Air Fleet (CRAF)*, Report RL33692, Oct. 18, 2006, <https://sgp.fas.org/crs/weapons/RL33692.pdf>.

231 CRAF activation stages: Stage I (regional crisis, response within twenty-four hours); Stage II (major war not warranting full national mobilization, response within twenty-four hours); Stage III (multiple theater wars or national mobilization, total CRAF airlift within forty-eight hours). GAO, *Military Readiness: Civil Reserve Air Fleet Can Respond as Planned, but Incentives May Need Revamping*, GAO-03-278, December 2002, <https://www.gao.gov/assets/gao-03-278.pdf>.

232 Congressional Research Service, *Afghanistan Evacuation: The Civil Reserve Air Fleet (CRAF) and the Defense Production Act (DPA)*, IN11731, Aug. 25, 2021, <https://www.congress.gov/crs-product/IN11731>. Past annual government contracts to CRAF participants include: CY24: \$736.5 million; CY23: \$2.1 billion; CY22: \$2.1 billion; CY21: \$1.3 billion; CY20: \$1.3 billion.

233 42 U.S.C. § 247d-6b.

234 42 U.S.C. § 247d-7e.

A Promising DOD Model:

DOD's Capability Program Executive for Chemical, Biological, Radiological, and Nuclear Defense (now called CPE) provides rapid development and fielding capability for medical countermeasures—encompassing platform technologies for vaccine and therapeutic development, advanced manufacturing, and accelerated regulatory pathways. It already demonstrates the core mechanics that a national bioresilience architecture requires: pre-positioned development pipelines, regulatory acceleration, and the ability to move from target identification to fielded countermeasure on operational timelines. These mechanisms exist today at program scale, with documented outcomes against chemical, biological, radiological, and nuclear threats.²³⁵

Two laws provide liability protection for companies supporting public health emergencies. The Public Readiness and Emergency Preparedness (PREP) Act provides liability protections for countermeasure deployment but creates no production commitments.²³⁶ Section 708 of the DPA authorizes the president to bring industry together to voluntarily plan and coordinate actions to support national defense preparedness, including expanded production capacity and supply. Importantly, it also provides participants with a limited defense against antitrust liability for actions taken within an approved plan. During COVID-19, the federal government relied on this authority to align manufacturers and distributors of critical healthcare resources. Effective as that was in the moment, Section 708 agreements are episodic and reactive: They are stood up in response to an emergency, and not structured to create durable incentive frameworks or pre-negotiated surge relationships that persist between crises.²³⁷

This patchwork of programs and authorities covers many of the capabilities needed to respond to today's biothreats, but they are not coordinated nor synergized to provide the prompt and reliable readiness that would be created by a National Strategic Bio Readiness Reserve.

Designing a National Strategic Bio Readiness Reserve

A National Strategic Bio Readiness Reserve would establish a critical readiness framework for biological threats, natural and manmade. Under this framework **private firms would voluntarily commit defined portions of laboratory (such as assays and sequencing) and manufacturing capacity for federal activation during national security or public health emergencies.** Importantly, while existing DPA authorities provide a partial foundation for this, they are insufficient to establish a capability of this scope. Doing so would require express congressional authorization analogous to CRAF. A robust statutory framework ensures durability and predictability for participating firms and reliable readiness for the government.

235 Kelly Burkhalter and Daniel Critchfield, "JPEO-CBRND Embraces the Need for Speed in Biodefense: How JPEO Is Shaping Its Biological Defense and Medical Strategies," *US Army*, Aug. 9, 2023, https://www.army.mil/article/268954/jpeo_cbrnd_embraces_the_need_for_speed_in_biodefense_how_jpeo_is_shaping_its_biological_defense_and_medical_strategies.

236 42 U.S.C. § 247d-6d.

237 Defense Production Act of 1950, 50 U.S.C. §§ 4501 et seq.; 50 U.S.C. § 4558; 44 C.F.R. pt. 332; Congressional Research Service, "Defense Production Act (DPA): Recent Developments in Response to COVID-19," July 28, 2020, https://www.congress.gov/crs_external_products/IN/PDF/IN11470/IN11470.5.pdf; Congressional Research Service, "The Role of Section 708 of the Defense Production Act in the Federal Government's Response to COVID-19: Antitrust Considerations," LSB10534, Aug. 20, 2020, <https://www.congress.gov/crs-product/LSB10534>.

Under the NSBRR, the government acts as anchor tenant, coinvesting for capability and reserving a set percentage of partner capacity (e.g., ~10 percent) to reduce risk for private investment. The private sector builds and operates facilities for commercial purposes—drug development, agricultural biotech, diagnostics—while the government secures guaranteed access for crisis activation. Integrated government-private teams would exercise continually, including annual tabletop and field exercises simulating novel pathogen emergence. At the apex, NSBRR stands up Bio Emergency Services Teams (BEST) that function as a national biological “fire department”—relevant to localized public health threats and catastrophic disasters alike—mirroring the proven Nuclear Emergency Support Team (NEST) model that has provided the nation with rapid nuclear incident response for five decades. NEST is a specialized Department of Energy capability that deploys scientific, technical, and operational experts to detect, assess, and respond to nuclear and radiological incidents.

To respond to the complex biothreat landscape, the NSBRR will need to cover broad territory. This is best addressed through four distinct streams of national capability: design and discovery, compute power, advanced development, and biomanufacturing.²³⁸

Design and Discovery: Participating firms operating Biosafety Level 2 (BSL-2), BSL-3, or BSL-4 facilities would commit a defined percentage of laboratory throughput and trained personnel for surge pathogen characterization, assay validation, or countermeasure testing. Activation timelines would be tiered analogous to CRAF’s staged mobilization. Companies providing high-throughput sequencing, synthetic biology services, and AI-enabled protein design would reserve a portion of computational and wet-lab capacity for rapid genomic analysis and countermeasure design.

Compute Power: The government would have standing prompt access to high-performance computing resources, including classified systems, that are trained on biological data and which would be available for rapid deployment in emerging situations and throughout a crisis (described further below). This compute power is important for AI as well as for other biological response applications such as modelling and simulation.

Advanced Development: Participating organizations would commit a portion of the specialized facilities, personnel, and organizational capabilities needed to translate promising medical countermeasure candidates into deployable products, including preclinical testing, clinical trials, process and formulation development, and quality assurance.

Biomanufacturing: Vaccine, therapeutic, and advanced biologics manufacturers would commit a percentage of fill-finish lines, mRNA production suites, viral vector capacity, or other platforms. Pre-negotiated Emergency Use Authorization (EUA) pathways under Food and Drug Administration (FDA) authority could also be incorporated into NSBRR activation procedures.²³⁹

238 The three streams are modeled on CRAF’s organization. CRAF has three main segments: international, national, and aeromedical evacuation. The international segment is further divided into the long-range and short-range sections and the national segment into the domestic and Alaskan sections. CRS, *Civil Reserve Air Fleet*.

239 21 U.S.C. § 360bbb-3 (Emergency Use Authorization authority).

A National Compute Reserve

To maintain a strong readiness posture, the federal government needs immediate, guaranteed availability of AI-ready biological data and compute power to analyze dangerous pathogens, train and evaluate biologically trained models, run countermeasure design pipelines, and support real-time response. Without arranging for such access in advance, those AI resources will be subject to commercial demands and procurement delays.

The COVID-19 pandemic led to the creation of this type of consolidated compute access with the High Performance Computing Consortium. The Consortium brought together DOE National Laboratories, NSF, NASA, private technology companies, and international laboratories to offer computing capabilities for pandemic research.²⁴⁰ An October 2021 National Science and Technology Council blueprint proposed a National Strategic Computing Reserve explicitly modeled on lessons from that consortium.²⁴¹

The **Genesis Mission**, launched in November 2025, provides a partial foundation for a compute reserve capability.²⁴² The executive order establishing the Genesis Mission tasks the Department of Energy (DOE) with creating a new American Science Cloud (AmSC) and a Transformational AI Models Consortium (ModCon).²⁴³ Among the twenty-six Genesis Mission challenges announced in February 2026, “Scaling the Biotechnology Revolution” calls for using AI to accelerate bioproduct manufacturing and experiment with AI-enabled biological research tools.²⁴⁴ However, the defined challenges do not include a focused effort on combating bio threats and de-risking bioconvergence. This could be remedied by establishing a specific DOE challenge or workstream leveraging National Laboratories, the AmSC, and ModCon to research biological safeguards and build emergency response compute capacity to meet current and future needs. As with the NSBRR, this dedicated access to compute capacity would contribute to deterrence. Most recently, frontier AI labs have begun to recognize the need to provide compute capacity for public good. On May 14, 2026, Anthropic, in partnership with the Gates Foundation, announced that they would make available \$200 million over four years in “grant funding, API credits, and technical support to develop AI tools and shared public goods—freely available resources—across health, education, and agriculture.”²⁴⁵ The public health initiative, although directed at low- and

240 The COVID-19 High Performance Computing Consortium brought together DOE National Laboratories, NSF, NASA, private technology companies, academia, and international laboratories to offer computing capabilities for pandemic research. National Science and Technology Council, *National Strategic Computing Reserve: A Blueprint*, October 2021, <https://www.nitrd.gov/pubs/National-Strategic-Computing-Reserve-Blueprint-Oct2021.pdf>. It ultimately included forty-three participants and over 600 petaflops of supercomputing systems. Jim Brase et al., “The COVID-19 High-Performance Computing Consortium,” *Computing in Science & Engineering* 24, no. 1 (2022): 78–85, <https://doi.org/10.1109/MCSE.2022.3145608>.

241 Ibid. The NSCR has not been fully implemented. Following a 2024 tabletop exercise, the NSTC highlighted the need for formal authorities enabling agency participation and resource sharing. Dylan Cohen and Kush Patel, *National Strategic Computing Reserve Tabletop Exercise After Action Report*, IDA-3003599, December 2024, <https://www.ida.org/-/media/feature/publications/n/na/national-strategic-computing-reserve-tabletop-exercise-after-action-report/3003599.ashx>.

242 Exec. Order No. 21665, “Launching the Genesis Mission,” Nov. 24, 2025, <https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission/>.

243 US Department of Energy [DOE], *Genesis Mission and National Science and Technology Challenges*, Feb. 12, 2026, <https://www.energy.gov/documents/genesis-mission-science-and-technology-challenges>; DOE, “Energy Department Announces \$293 Million in Funding to Support Genesis Mission National Science and Technology Challenges,” Mar. 17, 2026, <https://www.energy.gov/articles/energy-department-announces-293-million-funding-support-genesis-mission-national-science>.

244 DOE, *Genesis Mission and National Science and Technology Challenges*.

245 Gates Foundation, “Making AI Work for More People,” press release, May 14, 2026, <https://www.gatesfoundation.org/ideas/media-center/press-releases/2026/05/ai-anthropic-partnership>; Anthropic, “Anthropic Forms \$200 Million Partnership with the Gates Foundation,” May 14, 2026, <https://www.anthropic.com/news/gates-foundation-partnership>.

middle-income countries, is an example of what is needed to ensure readiness in the United States: “accelerat[ing] the development of vaccines and other critical innovations” and “making large, complex datasets more accessible, interactive, and actionable for researchers and decision-makers.”²⁴⁶ However, this kind of philanthropic giving, while beneficial, is not sufficient to support national-scale readiness to respond to bio threats, and cannot be rapidly targeted on the areas most needed in a crisis. Creating readily accessible compute reserve, including classified capabilities within a secure enclave, will require sustained government action and enhanced, durable public-private partnerships.

Private-Sector Incentives for the NSBRR

Across the four pillars of the NSBRR, sustained participation in the reserve model requires private sector incentives structured to align commercial and national security interests between activations. This can be achieved in several ways. Minimum guaranteed government contracts can provide participating companies with a baseline of reliable annual federal procurement across DOD, HHS, and DHS. Federal procurements can also be streamlined for participants using priority and fast-track contracting, including accelerated review for BARDA awards and DOD research programs. In return for their participation, companies could be given priority access to government-controlled reagents, reference strains, and curated datasets. Building on the CASR example, structured participant engagement in federal response exercises would enhance both public and private sector interoperability, readiness, and resilience. Regular exercises also function as signals for deterrence. Finally, to protect participants from potential financial and legal risks, Congress could provide indemnification and liability protections modeled on CRAF as well as antitrust safe harbors when coordinated surge planning is required.²⁴⁷ The DPA provides vehicles for companies to coordinate with each other with oversight, but the NSBRR would benefit from such explicit authorization and protection for participating firms.

Understanding AI Safeguards for Bio Risks

Beyond establishing a reserve compute capacity, readiness requires understanding the safeguards in place on frontier AI systems. Public reviews of AI safety policies have found substantial variation in how leading AI laboratories structure and disclose their chemical, biological, radiological, and nuclear (CBRN) safeguards.²⁴⁸ Current industry practices are heterogeneous and, in many cases, only partially transparent to external stakeholders.²⁴⁹ As a result, the government lacks the information necessary to assess systemic risks, identify gaps, or design targeted interventions.

The DPA provides the government authority to obtain this information as part of an industrial base survey. The government has invoked this authority to require dual-use foundation model developers

²⁴⁶ Ibid.

²⁴⁷ Voluntary agreement under DPA Section 708, modeled on the FEMA pandemic-response committee, would allow private sector participants to coordinate across the three NSBRR segments, sharing technical information about capabilities, dividing roles for specific response scenarios, and establishing open channels with government, while operating under DOJ and Federal Trade Commission (FTC) oversight.

²⁴⁸ METR, *Common Elements of AI Safety Policies*, March 2025, <https://metr.org/assets/common-elements-mar-2025.pdf>.

²⁴⁹ EthicAI, *Frontier AI Safety Failures*, accessed 2026, <https://ethicai.net/frontier-ai-safety-failures> (“Even in high-attention domains such as biorisk, the report finds an overreliance on task-specific tests that do little to uncover latent dangerous capabilities or performance under adversarial conditions.”).

to report on physical and cybersecurity protections, NIST red-team results, and biological weapon capability red-team results.²⁵⁰ More recently, the CAISI issued voluntary sources-sought notices rather than compulsory surveys without focus on biorisks or CBRN safeguards specifically.²⁵¹ A more robust understanding of frontier AI safeguards built on mandatory reporting would strengthen US readiness to respond to biothreats. The most valuable improvements could be achieved with expanded insights into CBRN safeguards currently in place for both deployment and development, safeguards that companies have chosen not to implement; known gaps in existing safety mechanisms; and exposure to foreign distillation attacks on foundational models.²⁵²

Existing Authorities to Promote CBRN Safeguards

DPA Title I and Title III already provide authorities for both compelling and incentivizing CBRN safeguard deployment.²⁵³ On the incentive side, Title III allows the US government to financially incentivize domestic industrial base expansion for the production and supply of critical materials and goods for national defense purposes. Title III could therefore be used to co-fund or underwrite the development of AI models with built-in CBRN guardrails, including fine-tuning methods that reliably refuse to assist with dangerous wet-lab protocols or biological agent design. Title III allows government purchase guarantees or cost-sharing for secure, government-facing versions of frontier models incorporating CBRN safeguards.

When incentives are insufficient, DPA Title I allows the government to require private sector actors to both accept and prioritize contracts for materials and services and allocate materials, services, and facilities necessary to promote the national defense. For bioconvergence risks, Title I could be used to compel firms to prioritize contracts for CBRN safeguard development. However, as recent reports of the potential invocation of Title I against Anthropic demonstrated, this is only applicable where that capability already exists and only for clearly defined deliverables.²⁵⁴ The more precisely the government can define the required CBRN safeguard and explain the firm's capability to deliver it, the stronger the Title I footing. (For further discussion of DPA, see **Appendix 4.**)

Outside of DPA authorities, the Defense Innovation Unit (DIU) and DARPA challenge models offer a complementary approach to providing incentives for private sector innovation in CBRN safeguards. In 2022, the DIU launched a strategic initiative integrating Responsible AI (RAI)

250 Exec. Order No. 14110, "Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," §§ 4.2(a)-(b).

251 National Institute of Standards and Technology, Center for AI Standards and Innovation (CAISI), accessed June 18, 2026, <https://www.nist.gov/caisi>; National Institute of Standards and Technology, "Sources Sought Notice: Open-Weight AI Model Hosting and Inference Services for National Security Evaluations," SAM.gov, February 2026; "NIST Seeks Industry to Host AI Models for National Security Reviews," *MeriTalk*, Feb. 24, 2026, <https://www.meritalk.com/articles/nist-seeks-industry-to-host-ai-models-for-national-security-reviews/>. CAISI's sources-sought notices are voluntary RFIs; they do not carry penalties for nonresponse and have not focused on biorisk or CBRN safeguards specifically.

252 "Anthropic, OpenAI and China Firms Face Distillation Concerns," *CNBC*, Feb. 24, 2026, <https://www.cnbc.com/2026/02/24/anthropic-openai-china-firms-distillation-deepseek.html>; Anthropic, "Detecting and Preventing Distillation Attacks," <https://www.anthropic.com/news/detecting-and-preventing-distillation-attacks>.

253 50 U.S.C. §§ 4511–4517 (DPA Title I, priorities and allocations); 50 U.S.C. §§ 4531–4533 (DPA Title III, expansion of productive capacity and supply); Administration for Strategic Preparedness and Response [ASPR], "Defense Production Act Overview," <https://aspr.hhs.gov/legal/DPA/Pages/Overview.aspx>.

254 "Exclusive: Hegseth Gives Anthropic Until Friday to Back Down on AI Safeguards," *Axios*, Feb. 24, 2026, <https://www.axios.com/2026/02/24/anthropic-pentagon-claude-hegseth-dario>; "Anthropic Offered Pentagon the Ability to Use AI Systems for Missile Defense," *NBC News*, Feb. 25, 2026; "What the Defense Production Act Can and Can't Do to Anthropic," *Lawfare*, Feb. 25, 2026, <https://www.lawfaremedia.org/article/what-the-defense-production-act-can-and-can-t-do-to-anthropic>.

requirements into its commercial prototyping and acquisition program.²⁵⁵ DARPA's Bio-attribution Challenge, announced March 12, 2026, is focused on pathogen detection and attribution of pathogens to their engineering origin.²⁵⁶ Beyond bio-specific approaches, DARPA's AI Cyber Challenge (AlxCC) demonstrates the model's capacity to mobilize top private-sector talent on national security problems at significant scale.²⁵⁷ This challenge model can provide an excellent tool for providing low-cost incentives for private sector CBRN innovation. BARDA's Division of Research, Innovation, and Ventures (DRIVE) is already positioned to lead a challenge model for novel CBRN safeguards that can prevent, deter, or detect the use of AI models in the creation of bioweapons.²⁵⁸ DRIVE has the institutional expertise, resources, and track record in challenge-model design at the intersection of medical and biological national security needs.

Conclusion

The gaps addressed in this section share a common root: The government's tools for managing biological risk were built before AI-enabled biology became a national security risk. Designating biotechnology as critical infrastructure creates the coordination structure the current system lacks. The National Strategic Bio Readiness Reserve pre-positions the industrial capacity the government cannot improvise fast enough in a crisis. Engaging the frontier AI industry through mandatory surveys, DPA incentives, and targeted challenge models builds the government's capability to exercise and deploy the systems that will determine the preparations for and outcomes of a biological incident. Taken together, these steps convert latent American capability into reliable national readiness. This readiness provides the foundation for meaningful response to a biological event and contributes to creating a new model of deterrence against bio threats.

Section A Findings and Recommendations

Finding 6.1.1: The biotechnology ecosystem is fragmented across multiple critical infrastructure sectors with no standing mechanism for cross-sector coordination between government and the full range of relevant private actors—biotech firms, frontier AI companies, and pharmaceutical manufacturers. This structural gap cannot be closed within the current architecture.

Recommendation 6.1.1: Designate the biotechnology ecosystem, including AI-enabled technologies, as a critical infrastructure sector and establish a standing public-private coordination mechanism.

- Designate the biotechnology ecosystem as a new critical infrastructure sector within the national framework either by statute or through executive action. Designate DOD,

255 Defense Innovation Unit, *Responsible Artificial Intelligence—2022 in Review*, May 1, 2023, https://s3.us-gov-west-1.amazonaws.com/publicdocs.diu.mil/RAI_Guidelines_-_2022_in_Review.pdf.

256 DARPA, "Bio-Attribution Challenge," Mar. 12, 2026, <https://www.darpa.mil/news/2026/bio-attribution-challenge>; DARPA, "Bio-Attribution Challenge," accessed June 18, 2026, <https://www.darpa.mil/research/challenges/bio-attribution>. Two-round competition; data curated by Lawrence Livermore National Laboratory (LLNL); awards \$180,000 in total prizes; culminates June 30, 2026.

257 DARPA, "AI Cyber Challenge (AlxCC)," accessed June 18, 2026, <https://www.darpa.mil/research/programs/ai-cyber>; DARPA, "AI Cyber Challenge Marks Pivotal Inflection Point for Cyber Defense," Aug. 8, 2025, <https://www.darpa.mil/news/2025/aixcc-results>.

258 BARDA Division of Research, Innovation, and Ventures (DRIVE), established within the administration for Strategic Preparedness and Response, HHS. Patch Forward Prize: \$50 million total for revolutionary developments in vaccine-delivery technologies. ASPR, BARDA Division of Research, Innovation, and Ventures (DRIVE); "Patch Forward Prize," <https://www.patchforwardprize.com/>.

DHS, and HHS as co-Sector Risk Management Agencies (SRMAs) for the biosecurity and biotechnology sector.

- Organize a standing Biotechnology Sector Coordinating Council that includes: DNA synthesis providers and gene foundries; synthetic biology and advanced biologics firms; genomic data repositories and computational biology platforms; cloud-based bioinformatics infrastructure; AI model developers whose systems materially enable biological design; advanced laboratory automation and distributed biomanufacturing operators; high-containment laboratory operators; and critical biological supply chain nodes.
- A corresponding government coordinating council should include agencies capable of providing and examining threat information related to biosecurity, including, *inter alia*, DOD, DHS, HHS, DOE, NIST, CAISI, and components of the intelligence community.

Finding 6.1.2: No existing authority ensures that the United States can rapidly access private-sector industrial capacity—including laboratory throughput, sequencing pipelines, biomanufacturing lines, and personnel—to respond to a biological emergency. This is a major national readiness gap.

Recommendation 6.1.2: Establish a National Strategic Bio Readiness Reserve modeled on the Civil Reserve Air Fleet.

- Enact legislation authorizing the National Strategic Bio Readiness Reserve (NSBRR) as a standing, public-private reserve of design and discovery, compute power, advanced development, and biomanufacturing capacity, modeled on the Civil Reserve Air Fleet.
- Participating firms should receive guaranteed procurement preferences, standby payments, or other contractual incentives in exchange for committing a defined share of compute and production capacity during declared emergencies.
- Specify activation triggers, response timelines, incentive structures—including minimum government contract guarantees, liability and indemnification provisions—and annual audit requirements. Activation authority should be clearly delineated—potentially requiring a joint determination or consultation by the secretary of defense and the secretary of health and human services. Staged activation levels would distinguish between regional outbreaks and major public health emergencies.
- Conduct a DPA Section 705 Industrial Base Survey of each NSBRR segment and a parallel voluntary request for information on laboratory and manufacturing capability to properly scope the reserve.
- The federal government should conduct regular readiness exercises (e.g., annually or biennially) with participating firms to test surge production timelines, supply chain resilience, and distribution capabilities before future emergencies occur.
- Establish a specific Genesis Mission challenge to build bio emergency response compute capacity, including pre-negotiated access agreements for crisis activation.

Finding 6.1.3: The US government lacks a comprehensive, current assessment of frontier AI CBRN capabilities and safeguards. Although DPA authorities to compel or incentivize the deployment of safeguards exist, they are underutilized.

Recommendation 6.1.3.1: Strengthen CBRN safeguards on AI models.

- Perform a DPA Section 705 Industrial Base Survey of frontier AI developers' CBRN capabilities and safeguards including mandatory pre-deployment testing results, with specific focus on biosecurity and CBRN risks.
- Provide incentives for the development and procurement of frontier AI models with built-in CBRN safeguards using DPA Title III authorities. Use DPA Title I authorities to compel deployment of known, technically feasible safeguards when firms have chosen not to implement them.
- Design and fund a challenge under BARDA DRIVe to develop novel CBRN safeguards that can prevent, deter, or detect AI-enabled bioweapon development.

B. Response

In the previous section, we discussed generating the readiness that would enable the United States to respond quickly and effectively to biological threats. This section focuses on the response, including what happens when readiness capabilities must be activated. In some cases, response must begin before the full extent of a threat and the threat's origins are fully understood. As a result, response effectiveness may depend on decisions made under conditions of uncertainty, with incomplete information, limited time, and potential public panic. Adding to these challenges, in the age of bioconvergence, AI-enabled biology compresses every part of the response timeline: Novel pathogens may emerge faster, spread further before detection, and require countermeasures that existing systems were not built to rapidly produce or deploy.

This section addresses three major elements of response: the federal interagency decision architecture that enables expedited movement from detection of a threat into action; the medical countermeasure framework that delivers the response; and the clinical trial infrastructure that generates the evidence guiding treatment and delivers new countermeasures when needed. Building an effective clinical trial system requires investment before a crisis and activation during one. The reforms described in this section cover the infrastructure, financing, and governance changes that must be established in advance, and the trial networks, authorities, and data systems that enable a rapid response when a biological threat emerges.

Decision-Making

As discussed in Chapters 4 and 5, accurate and rapid detection and attribution contribute to the effective response to a biological incident, including promptly identifying whether the event is natural, accidental, or deliberate, and, where possible, its source. The capabilities required for attribution are addressed in Chapter 5, including AI-enabled genomic analysis, wastewater surveillance, the Finding Engineering-Linked Indicators (FELIX) program, and the broader

biosurveillance architecture. Understanding the nature and source of a biological threat informs the ensuing decision-making process about the level and character of response required.

Decision Architecture

The competing tensions between law enforcement and public health data sharing are addressed in Chapters 4 (Detection) and 5 (Attribution), which focus on the institutional challenge of coordinating across agencies to detect, characterize, and attribute a biological event of unknown origin. Response coordination poses a different challenge. Detection asks what is happening and where it originated. Incident response asks what we do once we have an answer or when we must act before we have one. The decision architecture for response differs from the coordination required for attribution. **Response involves determinations about countermeasure deployment, public health action, military posture, and international notification that must often occur with significant uncertainty, before attribution is complete.**

Advances in detection and attribution provide value only if they can be translated rapidly into operational decisions. The strategic advantage created by earlier warning is not only measured by how quickly an event is identified, but by how quickly leaders can decide on a course of action and implement it. Effective biological incident response requires a federally-led and coordinated decision architecture that is clearly understood, routinely exercised, and supported by experienced personnel at the federal, state, and local levels.

The US government's coordination across the complex federal system and among federal, state, and local authorities has relied upon the Biological Incident Notification and Assessment Protocol and the White House National Security Council-convened policy coordination and decision-making mechanism. Without such a structured process, interagency and intergovernmental coordination weaknesses can delay the first line of response, which is the rapid deployment of countermeasures to limit public harm once a threat is identified. A failure to respond effectively, particularly early in an emerging crisis, will undermine public confidence in government actions and can impact deterrence.

Although each step of response needs to be fully resourced, there are three current critical gaps in the US government's response architecture:

- 1. National Security Oversight.** Biological incidents of unknown or contested origin create policy questions that cannot be resolved within any single department or agency.²⁵⁹ Public health officials may prioritize rapid disease containment, law enforcement may seek to preserve investigative equities, intelligence agencies may protect sensitive sources and methods, and defense officials may assess implications for force protection.²⁶⁰ These interests are legitimate but can compete with one another and/or fragment information-sharing during the early stages of a crisis. Experienced senior-level White House coordination, particularly by the National Security Council,²⁶¹ is therefore necessary to establish priorities, share information, resolve disputes, and ensure that

259 GAO, *COVID-19: Lessons Can Help Agencies Better Prepare for Future Emergencies*, GAO-24-107175, Aug. 1, 2024.

260 FEMA, *Biological Incident Annex to the Response and Recovery Federal Interagency Operational Plans* (2017), https://www.fema.gov/sites/default/files/documents/fema_incident-annex_biological.pdf.

261 National Security Memorandum 15, "National Security Memorandum on Countering Biological Threats, Enhancing Pandemic Preparedness, and Achieving Global Health Security," Oct. 18, 2022, <https://www.govinfo.gov/content/pkg/DCPD-202200938/pdf/DCPD-202200938.pdf>.

response decisions reflect integrated national objectives rather than the perspective of any individual agency.

- 2. Regular Exercises.** Unlike many emergency response systems that are activated frequently, high-consequence biological incident mechanisms may go years without real-world use. Decision frameworks that exist only on paper often prove insufficient during a fast-moving crisis, particularly when multiple agencies possess overlapping authorities and competing priorities. Exercises allow agencies to test decision pathways, identify bottlenecks in information sharing, clarify authorities, and familiarize senior leaders with the difficult tradeoffs that characterize biological crises.²⁶² Although US government agencies regularly exercise for known biological risks, they still lack overall plans for novel threats, common data sharing mechanisms, and feedback loops to learn from exercises.²⁶³ Just as military readiness depends on repeated training and evaluation, preparedness for biological incidents depends on routine rehearsal of the decision-making system itself.
- 3. Data Sharing.** One of the most persistent challenges in biological incident response is the tension among public health, intelligence, and law-enforcement requirements.²⁶⁴ Information is often held by a range of organizations operating under different legal authorities and privacy constraints.²⁶⁵ During a crisis, uncertainty regarding what information can be shared, with whom, and under what conditions can delay decision-making at precisely the moment when speed matters most. Data-sharing agreements and associated legal authorities should therefore be established, reviewed, and updated regularly before an incident occurs rather than during a crisis.²⁶⁶

The faster the government can move from detection to decision, the faster it can authorize and deliver the diagnostics, therapeutics, and vaccines that form the operational core of a biological response. That delivery system is addressed in the following section.

Rapidly Financing, Developing, and Fielding Medical Countermeasures and PPE

Although this report focuses on innovations specific to bioconvergence, any effective response to biological threats must also address broader deficiencies in public health emergency preparedness. Bioconvergence therefore requires modernizing the US and global system for rapidly developing, manufacturing, and deploying vaccines, therapeutics, diagnostics, and personal protective equipment (PPE) in response to emerging biological threats.

262 World Health Organization, "Simulation Exercises," accessed June 18, 2026, <https://www.who.int/emergencies/operations/simulation-exercises>.

263 GAO, *Biodefense: After-Action Findings and COVID-19 Response Revealed Opportunities to Strengthen Preparedness*, GAO-21-513, Aug. 4, 2021; GAO, *Biodefense: Actions Needed to Address Long-Standing Challenges*, GAO-23-106476, Mar. 9, 2023.

264 Centers for Disease Control and Prevention, *Suspected Intentional Use of Biologic and Toxic Agents*, Aug. 4, 2024, <https://www.cdc.gov/field-epi-manual/php/chapters/biologic-toxic-agents.html>.

265 Federal Bureau of Investigation and Centers for Disease Control and Prevention, *Joint Criminal and Epidemiological Investigations Handbook* (2018).

266 World Health Organization, *WHO Benchmarks for Strengthening Health Emergency Capabilities*, (2023): Benchmark 13.1, <https://iris.who.int/server/api/core/bitstreams/8883ae92-ecbe-4cc6-8ad4-500297ef1df9/content>.

The COVID-19 pandemic demonstrated both the strengths and limitations of the nation's current response capabilities. Operation Warp Speed showed what is possible when financing, regulatory authorities, manufacturing capacity, and logistics are aligned behind a common mission. However, it also revealed that this alignment depended heavily on DOD authorities, contracting mechanisms, and logistical infrastructure that are not automatically available during civilian public health emergencies.

The United States needs a **standing national infrastructure for medical countermeasure development and deployment that does not rely on the availability of military resources**, which could be committed elsewhere in a time of crisis. Such a capability would improve routine public health preparedness while providing the “warm start” capacity needed to respond rapidly to emerging biological crises, whether natural, accidental, or deliberate. This capability is also critical globally. Investments in R&D and financing to support the 100 Days Mission globally not only save lives and stop spread around the world, they also protect Americans and allow the United States to focus on protecting its population.

A civilian MCM distribution architecture should build on existing assets while closing the gaps they were not designed to fill. The Strategic National Stockpile (SNS) remains the central component of federal medical response infrastructure, holding emergency medicines, vaccines, and supplies for CBRN threats, emerging infectious diseases, and pandemics.²⁶⁷ However, stockpiling and federal deployment are only part of the operational chain required during a biological emergency. The 2024 PHEMCE Strategy and Implementation Plan identifies “strengthening last-mile MCM efforts” as a goal, including lessons from COVID-19 and mpox related to MCM dispensing, distribution, and administration.²⁶⁸ Those goals should be converted into standing operational capacity backed by appropriations and statutory authorities.

The lesson from COVID-19 is not that the United States lacks emergency economic authorities. It is that those authorities are too often activated after a crisis has already exposed the fragility of the civilian distribution base. During COVID-19, federal agencies used DPA authorities to prioritize contracts for medical supplies, expand domestic production, and address shortages affecting respirators, ventilators, masks, testing supplies, and other critical resources.²⁶⁹ The next phase should be anticipatory. HHS should use DPA authorities during peacetime to map the MCM distribution industrial base; identify bottlenecks in cold-chain logistics, regional warehousing, ancillary supplies, diagnostic components, PPE replenishment, pharmacy distribution, and interoperable logistics data systems; and pre-negotiate the contracts and coordination mechanisms that would be activated during a biological emergency. DPA Title III can support domestic production and supply-chain resilience for critical health resources,²⁷⁰ while Title I priority-rated contracts and Section 708 voluntary agreements can support structured precrisis coordination among manufacturers, distributors, wholesalers, pharmacies, health systems, and

267 ASPR, “Center for the Strategic National Stockpile,” US Department of Health and Human Services, accessed June 28, 2026, <https://www.aspr.gov/about/centers/sns>.

268 ASPR, “2024 PHEMCE Goals,” *Public Health Emergency Medical Countermeasures Enterprise Strategy and Implementation Plan 2024*, US Department of Health and Human Services, accessed June 28, 2026, <https://www.aspr.gov/readiness-response/medical-countermeasures-biodefense/phecme/2024-SIP/goals#g1>

269 GAO, *Defense Production Act: Opportunities Exist to Increase Transparency and Identify Future Actions to Mitigate Medical Supply Chain Issues*, GAO-21-108 (Nov. 19, 2020), <https://www.gao.gov/products/gao-21-108>.

270 ASPR, “Title III – Expansion of Productive Capacity and Supply,” US Department of Health and Human Services, accessed June 28, 2026, <https://www.aspr.gov/about/law-policy-strategy/hhs-legal-authorities-related-disasters/dpa/title-iii>.

logistics providers.²⁷¹ The gap is therefore not the absence of DPA authority, but the absence of a standing civilian architecture that uses those authorities before a crisis begins. Without this capability, the United States may be able to develop and manufacture lifesaving countermeasures faster than ever while still failing to deliver them at the speed and scale a biological emergency requires.

Congress should provide HHS the statutory authority and sustained appropriations needed to build a civilian distribution architecture that can be activated rapidly during a biological emergency. That capability should include regional hubs, pre-negotiated commercial distribution contracts, routine exercises with state and local partners, and a national logistics data layer capable of tracking shipment, receipt, administration, wastage, adverse events, and geographic or demographic gaps.²⁷² Distribution analytics should be treated as core response infrastructure.

A key component of the United States' response capability is the Food and Drug Administration's Emergency Use Authorization (EUA) authorities, which provide an important mechanism for accelerating the deployment of medical countermeasures during public health emergencies. However, emergency regulatory flexibility alone is insufficient. Rapid response also requires surge-capacity biomedical research, accelerated evidence generation, and a far more robust clinical trial infrastructure. The United States should establish pre-negotiated manufacturing commitments, standing clinical trial networks, and regulatory fast-track pathways to ensure that diagnostics, therapeutics, and vaccines can be developed, tested, manufactured, and distributed at crisis speed.

Following COVID-19, numerous expert reviews outlined reforms needed to strengthen the physical foundations of biodefense. Among the most influential is the Coalition for Epidemic Preparedness Innovations' (CEPI) "100 Days Mission," which argues that future pandemic preparedness must shift from reactive emergency mobilization to a permanent preparedness infrastructure capable of developing and deploying vaccines, therapeutics, diagnostics, and protective measures within approximately 100 days of identifying a novel pathogen.²⁷³ Additionally, CEPI's capabilities globally allow the United States to focus at home during a bio crisis, while allies and other partners are covered globally.

Across its reports, CEPI consistently recommends sustained financing, public-private coordination, and international cooperation to maintain warm-base manufacturing, prototype pathogen research, and emergency response capabilities between crises. Importantly, in 2026, the US federal government began reengaging with CEPI, including through funding.²⁷⁴ Pandemic preparedness should function as a continuously maintained national and global security capability rather than an *ad hoc* public health response.

271 GAO, *Defense Production Act: Opportunities Exist*; Federal Emergency Management Agency, "Voluntary Agreement Under Section 708 of the Defense Production Act; Manufacture and Distribution of Critical Healthcare Resources Necessary to Respond to a Pandemic," *Federal Register* 85, no. 159 (Aug.17, 2020): 50035–50049.

272 ASPR, "2024 PHEMCE Goals."

273 Coalition for Epidemic Preparedness Innovations (CEPI), "The 100 Days Mission," <https://cepi.net/100-days-mission>; Coalition for Epidemic Preparedness Innovations, *Delivering Pandemic Vaccines in 100 Days: What Will It Take?* (2022), https://static.cepi.net/downloads/2024-02/CEPI-100-Days-Report-Digital-Version_29-11-22.pdf.

274 Coalition for Epidemic Preparedness Innovations, "US Department of State Commits \$50 Million to CEPI to Fast-Track Bundibugyo Virus Medical Countermeasures," accessed July 5, 2026, <https://cepi.net/us-department-state-commits-50-million-cepi-fast-track-bundibugyo-virus-medical-countermeasures>.

CEPI and other comprehensive assessments of US biodefense preparedness converge on four core priorities:

- 1. Invest in Adaptable Platform Technologies.** The federal government should sustain investment in technologies that can be rapidly adapted to novel pathogens,²⁷⁵ including:
 - mRNA vaccine platforms
 - Viral-vector platforms
 - Recombinant protein technologies
 - Nucleic-acid diagnostic platforms
- 2. Build Prepared Manufacturing Capacity.** Prepared manufacturing capacity is more valuable than emergency improvisation. Key investments include:
 - Warm-base domestic manufacturing
 - Flexible fill-finish facilities
 - Pre-negotiated surge manufacturing contracts
 - Strategic stockpiles of critical raw materials
 - Geographically distributed production capacity
- 3. Strengthen Diagnostic Readiness.** Diagnostics remain one of the most underinvested components of biodefense preparedness.²⁷⁶ Priority investments include:
 - Pathogen-agnostic testing platforms
 - Decentralized testing capabilities
 - Accelerated regulatory pathways
 - National diagnostic coordination mechanisms
 - Modernized public health laboratory networks
- 4. Treat PPE as an Industrial Base Issue.** PPE response depends on resilient supply chains rather than stockpiles alone.²⁷⁷ A durable PPE strategy would include:
 - Domestic and allied production incentives;
 - Reusable and reusable-certified PPE;

²⁷⁵ US Department of Defense, *Biodefense Posture Review* (2023), https://media.defense.gov/2023/Aug/17/2003282337/-1/-1/1/2023_BIODEFENSE_POSTURE_REVIEW.PDF; US Department of Health and Human Services, ASPR, *2024 Public Health Emergency Medical Countermeasures Enterprise Strategy and Implementation Plan* (September 2024), <https://aspr.hhs.gov/PHEMCE/2024-PHEMCE-SIP/Documents/2024-PEHMCE-SIP-508.pdf>; Center for Strategic and International Studies, “Building Apollo’s Arsenal: Acquiring COVID-19’s Lessons Learned for Government-Bioeconomy Partnership,” Dec. 14, 2021, <https://www.csis.org/analysis/building-apollos-arsenal-acquiring-covid-19s-lessons-learned-government-bioeconomy>.

²⁷⁶ Bipartisan Commission on Biodefense, *The National Blueprint for Biodefense* (May 2024), <https://biodefensecommission.org/reports/the-national-blueprint-for-biodefense/>; CSIS, “Building Apollo’s Arsenal.”

²⁷⁷ DOD, *Biodefense Posture Review*; HHS/ASPR, *PHEMCE Strategy and Implementation Plan*.

- Circular supply chains and recycling systems;
- Shelf-life management programs; and
- Diversified sourcing and supplier networks.

Together, these investments would establish a durable national capability for rapidly generating and deploying medical countermeasures against future biological threats, while strengthening routine public health preparedness and response.

Clinical Trials Infrastructure

The United States possesses world-leading biomedical research capabilities, yet it lacks a clinical trial enterprise capable of rapidly generating reliable clinical evidence that can support public health decision-making and patient care during a biological emergency. Today, the US clinical trial infrastructure is fragmented, slow, overconcentrated in large academic medical centers, and poorly incentivized for the collaboration that effective public health emergency response demands.²⁷⁸ This weakness is becoming increasingly consequential as advances in biotechnology and AI accelerate the pace of biological innovation, when the ability to rapidly generate reliable clinical evidence becomes a national security imperative. To meet this imperative, the United States must modernize and strengthen its clinical trial infrastructure so that it functions as a continuously operating national capability: embedded in everyday healthcare, sustained in peacetime, and ready to expand rapidly during future public health emergencies.

Recent public health emergencies exposed both the strengths and structural weaknesses of the US clinical trial system. Large, pragmatic platform trials during COVID-19 in the United Kingdom demonstrated that simple, adaptive trial designs can rapidly generate practice-changing evidence at national scale.²⁷⁹ In contrast, the United States launched hundreds of small, underpowered, and often duplicative studies, many of which failed to enroll adequately or produce definitive answers.²⁸⁰ Yet, the rapid development of mRNA vaccines through Operation Warp Speed demonstrated what is possible when financing, regulation, manufacturing, and clinical research are aligned.²⁸¹ Thus, the central challenge facing the United States is not whether it can develop new diagnostics, vaccines, or therapeutics. It is whether it can determine quickly enough which interventions work, for whom, and under what circumstances. This requires restructuring how clinical trials are designed, organized, financed, activated, governed, and integrated into care delivery by **shifting from facilitating a patchworked collection of studies to developing national assets integrated into everyday medicine.**

278 ACTIV Therapeutics–Clinical Working Groups, “The Future Is Now: Using the Lessons Learned from the ACTIV COVID-19 Therapeutics Trials to Create an Inclusive and Efficient Clinical Trials Enterprise,” *Open Forum Infectious Diseases* (2024), <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC11523011/>.

279 RECOVERY Collaborative Group, “Dexamethasone in Hospitalized Patients with Covid-19,” *New England Journal of Medicine* 384 (2021): 693–704, <https://doi.org/10.1056/NEJMoa2021436>; NHS England, “COVID Treatment Developed in the NHS Saves a Million Lives,” March 2021, <https://www.england.nhs.uk/2021/03/covid-treatment-developed-in-the-nhs-saves-a-million-lives/>.

280 Chana A. Sacks et al., “The Landscape of COVID-19 Research in the United States: A Cross-Sectional Study of Randomized Trials Registered on ClinicalTrials.gov,” *Journal of General Internal Medicine* 37 (2022): 154–61; Boghuma K. Titanji et al., “Strategies for Expediting Clinical Trials in the Next Public Health Emergency,” *JAMA Health Forum* 4 (2023): e233191, <https://doi.org/10.1001/jamahealthforum.2023.3191>.

281 Fernando P. Polack et al., “Safety and Efficacy of the BNT162b2 mRNA Covid-19 Vaccine,” *New England Journal of Medicine* 383 (2020): 2603–15, <https://doi.org/10.1056/NEJMoa2034577>; Lindsey R. Baden et al., “Efficacy and Safety of the mRNA-1273 SARS-CoV-2 Vaccine,” *New England Journal of Medicine* 384 (2021): 403–16, <https://doi.org/10.1056/NEJMoa2035389>.

The analysis of the US clinical trial infrastructure that follows is organized around four **strategic objectives**:

1. The United States must establish governance and financing mechanisms capable of supporting emergency clinical research.
2. The United States must build a continuously operating national clinical trial network capable of generating evidence at scale.
3. The United States must create the data and information infrastructure necessary to transform clinical research from a fragmented collection of studies into a durable national capability.
4. The United States must use steady-state resources so the private sector has incentives to invest in pandemic preparedness in peacetime, and so that the federal emergency clinical trial system is continuously tested and improved.

Together, these reforms would strengthen everyday healthcare while ensuring that the nation can generate reliable clinical evidence during the next biological emergency. In a biological crisis, delayed evidence is not a neutral administrative failure; it can lead to ineffective care, wasted clinical capacity, avoidable morbidity and mortality, and the erosion of public trust.

Strategic Challenge: Building a Responsive Standing National Learning Health System

Public health emergencies cannot be met with improvised research systems. Effective response depends on peacetime infrastructure in which clinical research is routine, efficient, and embedded within healthcare delivery. Nations that integrate trials into standard care achieve faster enrollment and clearer answers.²⁸² The experience of recent public health emergencies demonstrates that speed during crisis is largely determined by investments made before a crisis begins. Countries that maintained standing research networks, simplified administrative procedures, and integrated research into care delivery responded faster during the COVID-19 pandemic.²⁸³

Overall, the goal should be to reduce reliance on emergency improvisation and facilitate sustained readiness that improves everyday, peacetime care. Clinical trials should be treated as a public good and a **routine component of healthcare delivery**. In 2015, only 1 percent of overall cancer patients in the United States and 1.9 percent of those receiving active treatment participated in interventional cancer clinical trials.²⁸⁴ Yet a 2021 meta-analysis found that 55 percent of cancer patients offered participation agreed to enroll.²⁸⁵ The gap between public willingness and actual participation reflects structural barriers rather than patient reluctance.

282 Denise van Hout et al., “Hurdles for the Delivery of Multinational Randomized Clinical Trials,” *JAMA Network Open* 8, no. 7 (2025): e2518503, <https://doi.org/10.1001/jamanetworkopen.2025.18503>.

283 For example, the UK’s RECOVERY trial achieved a median time-to-enrollment approval of just eight days during COVID-19, compared to the 115-day median across European countries. UK contract completion times dropped from 196 days pre-pandemic to five days after trials were designated as priority and integrated into National Health Service (NHS) care delivery. The RECOVERY trial also benefited from a short, plain-language consent form. US consent forms, by contrast, typically range from ten to eighty or more pages.

284 Angela K. Green et al., “Clinical Trial Participation Among Older Adult Medicare Fee-for-Service Beneficiaries with Cancer,” *JAMA Oncology* 8 (2022): 1786–94, <https://doi.org/10.1001/jamaoncol.2022.5020>.

285 Joseph M. Unger et al., “‘When Offered to Participate’: A Systematic Review and Meta-Analysis of Patient Agreement to Participate in Cancer Clinical Trials,” *Journal of the National Cancer Institute* 113 (2021): 244–57, <https://doi.org/10.1093/jnci/djaa155>.

Breaking down these structural barriers is built on four foundational principles:

1. **Clinical research infrastructure must operate continuously.** Emergency research capacity cannot be created after a crisis begins. Site networks, contracts, institutional relationships, data systems, and trained personnel must be maintained in peacetime. Embedding research into routine clinical workflows improves safety, accelerates evidence generation, and enhances system learning. It also ensures that evidence reflects the patients who actually receive care, including older adults, individuals with multiple chronic conditions, rural populations, and patients treated outside academic medical centers.
2. **Clinical research must be designed for scale.** Overly complex protocols, excessive eligibility criteria, lengthy consent forms, and duplicative oversight procedures slow enrollment and dilute impact. Evidence from pragmatic trials demonstrates that streamlined eligibility criteria and simplified consent processes can maintain scientific validity while enhancing participation.²⁸⁶
3. **Clinical research must be nationally coordinated but locally executed.** Federal institutions should establish priorities, maintain readiness, and provide financing and governance. Community health systems, clinicians, and patients should be able to participate without navigating bespoke administrative barriers for every study.
4. **Trust must be treated as operational infrastructure.** Public participation depends on transparency, privacy protections, ethical governance, and confidence that clinical trials serve public health.

Taken together, these principles suggest a fundamental shift in how the United States approaches clinical trials. The objective should no longer be to build research capacity during emergencies. It should be to maintain a continuously operating learning health system capable of generating evidence during both routine care and crisis response.

Governance

Establishing Standing National Clinical Trial Emergency Authority

National clinical trial leadership requires clear authority and political backing to coordinate effectively across federal agencies, healthcare systems, industry partners, and research institutions. During a public health emergency, the government must be able to rapidly identify priority questions, align resources behind a limited number of high-value studies, and ensure that research efforts reinforce rather than compete with one another. The organizing principle should be national coordination with local execution: Federal institutions should set priorities, finance readiness, and simplify participation, while community health systems and clinicians execute trials within routine care.

The United States should establish a **permanent, federally supported clinical trial research infrastructure that functions as a standing national capability**. This infrastructure should maintain site networks, operational personnel, contracting mechanisms, and modern data systems

286 Ian Ford and John Norrie, "Pragmatic Trials," *New England Journal of Medicine* 375 (2016): 454–63, <https://doi.org/10.1056/NEJMr1510059>.

during peacetime. Continual investment serves two purposes. First, it preserves institutional expertise and operational readiness. Second, it eliminates the startup delays that occur when research systems must be reconstructed during emergencies. When a crisis emerges, the nation should not be scrambling to assemble research networks, negotiate contracts, hire staff, and build data systems from scratch.

Once national infrastructure exists, it must be activated through clear authority. The primary failure of emergency clinical research is often not lack of expertise but lack of organizational command. During a public health emergency, dozens of federal agencies, thousands of private institutions, product developers, and investigators must coordinate around a limited set of priority questions. Without a clear center of gravity, research fragments into duplicative and underpowered studies. For example, during the COVID-19 pandemic, US research was composed of hundreds of duplicative and underpowered trials. At the peak of the crisis, less than 1 percent of hospitalized patients in the US were enrolled in randomized trials.²⁸⁷ This occurred despite an unprecedented volume of cases that should have provided definitive answers in weeks.

A standing **National Clinical Trials Emergency Authority (NCTEA) within HHS** would ensure that the government, and the national infrastructure, can respond quickly in the event of a biological threat. The NCTEA should remain operationally capable during peacetime and receive expanded authority upon declaration of a public health emergency. In peacetime, this infrastructure should coordinate national clinical trial efforts around priority areas of biomedical innovation and public health needs, generating useful evidence while keeping site networks, protocols, data systems, and operational relationships warm. For example, standing platform trials for respiratory viral illness (further discussed below) could improve routine care during ordinary seasons while preserving the capacity to expand rapidly during a future pandemic. Its mission should be to coordinate federally supported emergency clinical research, designate priority national platform trials, activate pre-negotiated contracting and reimbursement terms, and align federal agencies around a unified evidence-generation strategy. To ensure the widest possible access to data collected by clinical trials, NCTEA should adopt an operational transparency standard for all federally designated emergency trials, requiring full public posting of all protocols and subsequent amendments, public dashboards showing real-time enrollment progress, and rapid dissemination of results.

The NCTEA should not replace NIH, FDA, BARDA, CDC, or CMS. Rather, it should provide the operational authority needed to align their activities during a crisis. Scientific review, regulatory decision-making, procurement, surveillance, and appropriations should remain with the agencies best suited to perform those functions. The missing capacity is not expertise—it is coordination backed by authority.

Streamlined and Predictable Paths for Medical Countermeasures

Effective emergency research requires clear authority and accountability. While the FDA demonstrated remarkable agility in vaccine review during the COVID-19 pandemic, the landscape for therapeutics was marred by broad and enduring uncertainty. Uncertainty regarding acceptable clinical endpoints, evidentiary standards, and the specific data required to support regulatory decisions increased investment risk, complicated trial design, and delayed the initiation of critical studies. This confusion was compounded by the inconsistent use of EUAs. Early in the pandemic,

287 Kamran Bugin and Janet Woodcock, "Trends in COVID-19 Therapeutic Clinical Trials," *Nature Reviews Drug Discovery* 20 (2021), <https://www.nature.com/articles/d41573-021-00037-3>.

several authorizations were issued for repurposed drugs like hydroxychloroquine based on weak, nonrandomized evidence. In contrast, when the FDA waited for high-quality data, such as the United Kingdom's RECOVERY trial's evidence for steroids, clinical practice changed within weeks.

Resolving these tensions requires the FDA to plan and standardize for crises:

- 1. Issue pathogen-class-specific roadmaps.** The FDA should publish clear, prespecified guidance for public health emergency regulations well in advance of the next biological threat. Guidance should focus on pathogen-class-specific roadmaps that define the evidentiary requirements for respiratory viruses, hemorrhagic fevers, and other priority pathogens. By clarifying the minimum evidentiary thresholds for novel antivirals and host-directed therapies, which are treatments that bolster the body's own immune response rather than attacking the virus directly, the FDA can provide the certainty needed to spur private investment. This guidance should serve as a public regulatory blueprint that replaces speculative trial design with a predictable pathway to authorization.
- 2. Pre-validate laboratories and assays.** A primary lesson from Operation Warp Speed was the value of processing clinical samples in a single, validated laboratory to generate cross-program comparative results. Instead of every sponsor developing a proprietary assay, a federally supported network should maintain standardized tests for priority pathogens. This would also eliminate duplicative validation efforts that currently add months to trial timelines.
- 3. Reform the EUA framework.** Emergency Use Authorizations provide an important mechanism for accelerating access to potentially beneficial interventions. However, premature authorization of inadequately studied products can undermine the very clinical trials needed to determine whether those interventions work. The EUA framework must prioritize the generation of randomized evidence over speculative use. Future authorizations should remain tethered to evidence-based hypotheses and rigorous data rather than the fluctuating demands of the political or social environment. By maintaining a high evidentiary bar, the government ensures that trial enrollment remains steady until researchers reach a definitive conclusion.

Reciprocity for clinical trials in allied countries that follow similarly rigorous standards could also break down barriers to quickly responding to public health emergencies. The benefits of conducting trials in allied jurisdictions with faster activation and simpler governance are limited by US regulations. Although the FDA has mutual recognition agreements with the European Union, Switzerland, and the United Kingdom²⁸⁸ for manufacturing inspections, there is not an equivalent structure in place for clinical trials. FDA regulations permit the acceptance of foreign clinical trial data under certain conditions, but they do not create a presumption of acceptance for data generated under allied regulatory systems, nor do they establish equivalency between FDA standards and those of peer regulators such as the European Medicines Agency, the UK Medicines and Healthcare products Regulatory Agency, Health Canada, or the Australian Therapeutic Goods Administration.²⁸⁹ Although the FDA may accept clinical data generated solely under allied oversight, such data do not receive a presumption of equivalence or automatic acceptance,

288 US Food and Drug Administration, "Mutual Recognition Agreements (MRA)," <https://www.fda.gov/international-programs/international-arrangements/mutual-recognition-agreements-mra>.

289 21 C.F.R. § 312.120 (conditions for acceptance of foreign clinical study data in IND applications).

which may create pressure on developers to include US sites or generate additional evidence even where the underlying standards are comparable.²⁹⁰ This gap contributes to duplicative trial requirements, delays in bringing effective treatments to US patients, and pressure on developers to conduct trials in lower-cost jurisdictions with weaker oversight, including China, rather than in allied countries whose regulatory systems and data standards are fully compatible with US requirements. **For bio emergencies, establishing prequalification for appropriately rigorous foreign approval processes would speed US access to important treatment options.**

Financing

The National Institutes of Health (NIH) remains a leading biomedical research institution, but it operates through a grant system built for deliberate research rather than rapid response. Because Congress generally provides separate appropriations to NIH institutes and centers, the NIH director has limited flexibility to transfer funds rapidly across accounts without specific statutory authority.²⁹¹ Research funding is distributed through a system optimized for hypothesis-driven science, extensive peer review, and annual appropriations cycles for grants. Those features, and the associated gold standard processes of evaluation and review, support scientific rigor under normal circumstances. During a fast-moving outbreak or biological emergency, however, they can impede the rapid deployment of clinical research.

HHS and NIH, therefore, need financing authorities specifically designed for public health emergencies. The centerpiece of this reform is an **Emergency Research Reserve Fund** that can be activated automatically through predefined statutory triggers modeled on the G7 100 Days Mission framework.²⁹² Rather than requiring new appropriations after a crisis begins, these triggers would release funding immediately upon identification of a significant biological threat. NIH should also receive limited emergency reprogramming authority that allows the agency to redirect resources across institutes during declared public health emergencies.

The United States should invest in global clinical trials networks, such as those provided by CEPI and other international organizations, particularly because many diseases that are critical for developing medical countermeasures and delivery do not regularly circulate in the United States, and pathogens know no borders.

Building a National Learning Clinical Trial System

A crisis-ready clinical trial system requires a healthcare system capable of generating evidence as a routine function of care. The United States needs a national learning health system in which community hospitals, clinical trial networks, platform studies, and patient participation mechanisms operate continuously during peacetime and can expand rapidly during public health emergencies.

290 Ibid.

291 Kavya Sekar, The National Institutes of Health (NIH): Background and Congressional Issues, CRS Report R41705 (Congressional Research Service, Jan. 13, 2025), <https://www.congress.gov/crs-product/R41705>; Dominick A. Fiorentino and Taylor N. Riccard, Transfer and Reprogramming of Appropriations: An Overview, CRS Report R47600 (Congressional Research Service, June 20, 2023), <https://www.congress.gov/crs-product/R47600>.

292 International Pandemic Preparedness Secretariat [IPPS], *100 Days Mission: Implementation Report—Progress in 2025 and Priorities for 2026*, Jan. 27, 2026, <https://ippsecretariat.org/publication/fifth-implementation-report/>.

Expanding Clinical Research Capacity in Community Health Systems

Clinical trials remain concentrated in academic medical centers despite the fact that most patients receive care in community settings.²⁹³ This concentration limits enrollment, slows evidence generation, and reduces participation by populations that are historically underrepresented in clinical research. COVID-19 demonstrated that community hospitals can successfully support enrollment and data collection at scale when provided appropriate infrastructure and support. Decentralizing infrastructure and funding to include community settings improves both the speed of a trial and the diversity of its participants,²⁹⁴ addressing persistent demographic disparities by meeting patients where they are.

First, the United States should build a national network around community systems and harmonize trial administration under a unified leadership structure. The federal government can leverage hospitals' participation in Centers for Medicare and Medicaid Services (CMS) by **requiring CMS participants above a defined capacity threshold (e.g., 100 beds) maintain standing clinical trial capability**. Participating institutions should maintain the personnel necessary to support participant enrollment, regulatory compliance, contracting, and data collection.

Second, scaling these capabilities would require federal funding and technical assistance. The current system creates a bottleneck by relying on fragmented reimbursement mechanisms, institution-specific contracting, and duplicative review processes. A **standardized national contracting and reimbursement** framework using pre-negotiated Indefinite Delivery/Indefinite Quantity (IDIQ) reimbursement contracts would reduce administrative friction and allow clinical trial infrastructure to operate at scale in community settings.

This infrastructure should operate continuously rather than being assembled only after a crisis emerges. Maintaining site networks and operational expertise during routine periods preserves institutional knowledge and eliminates the startup delays that routinely slow emergency research.

Streamlining Trial Activation and Oversight

The ability to activate a national clinical trial network depends on simplifying the administrative processes that routinely delay research. Federal policy already requires the use of a single Institutional Review Board (IRB)²⁹⁵ for multisite studies, yet implementation remains inconsistent,²⁹⁶ and in practice, many institutions still conduct their own independent ethics review, multiplying delays across every site in a multihospital trial.²⁹⁷ Federal agencies should **fully implement and**

293 Zachary A. Frosch, Nicole Illenberger, and Neha Mitra, "Trends in Patient Volume by Hospital Type and the Association of These Trends with Time to Cancer Treatment Initiation," *JAMA Network Open* 4 (2021).

294 Richard A. Oyer et al., "Increasing Racial and Ethnic Diversity in Cancer Clinical Trials: An American Society of Clinical Oncology and Association of Community Cancer Centers Joint Research Statement," *Journal of Clinical Oncology* 40 (2022): 2163–71.

295 An Institutional Review Board (IRB) is an independent committee that reviews and approves research involving human subjects and ensures compliance with federal law governing the ethics of human subject research. IRBs protect research participants by evaluating whether a study's risks are reasonable in relation to its anticipated benefits, whether participant selection is equitable, and whether informed consent procedures are adequate. The committee typically includes scientists, physicians, ethicists, and at least one community member who is not affiliated with the institution. Every institution conducting federally funded research, including hospitals, universities, and research centers, must maintain an IRB or use an approved external IRB. IRBs have authority to approve, require modifications to, or disapprove any research activity that falls within their scope, and they conduct ongoing review of approved studies to ensure participant protections remain in place throughout the life of a trial.

296 National Institutes of Health, "Final NIH Policy on the Use of a Single Institutional Review Board for Multi-Site Research," NOT-OD-16-094, June 21, 2016; 45 C.F.R. § 46.114(b) (Federal Policy for the Protection of Human Subjects).

297 Lauren W. Cohen et al., "Transitioning to the NIH Single IRB Model: Piloting the Use of SMART IRB," *Journal of Clinical and Translational Science* 3 (2019), <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7924135/>.

enforce existing single-IRB requirements for federally supported multisite trials. A designated central IRB should conduct protocol review on behalf of participating institutions while local sites retain responsibility for operational implementation and participant protection.

This streamlined environment must be overseen by a designated **emergency clinical trial officer at each participating institution**. These officers must have the explicit power to prioritize national protocols, coordinate emergency funding, and override nonessential administrative delays during a declared crisis.

Reforming Trial Design

Furthermore, a learning health system requires clinical trial infrastructure capable of answering multiple questions simultaneously and adapting as evidence evolves. **Platform trials** provide such a capability. A platform trial functions as a perpetual, multi-arm study conducted under a single overarching protocol. New interventions can be added and ineffective interventions can be removed without terminating the trial itself. This structure reduces duplication, improves statistical efficiency, and preserves institutional knowledge across studies.²⁹⁸

During the COVID-19 pandemic, large platform trials demonstrated that coordinated infrastructure can generate definitive answers faster than fragmented collections of individual studies.²⁹⁹ Their success reflected not only innovative trial design, but also the existence of standing operational systems capable of enrolling participants quickly and evaluating multiple interventions within a common framework.

However, the success of these platforms depends on a radical shift toward simplicity. Excessive eligibility criteria, expansive data collection requirements, and hundred-page consent forms undermine enrollment and delay answers. Evidence from pragmatic trials shows that streamlined protocols can maintain scientific validity while increasing both feasibility and representativeness.³⁰⁰ Three elements of simplicity are key here:

1. **Use master protocols for continuous disease-agnostic platform trials.** Rather than creating new studies for each emerging threat, federal agencies should maintain master protocols organized around priority syndromes such as viral pneumonia, severe respiratory illness, and other conditions likely to characterize future biological emergencies. A master protocol is a single, overarching framework that governs multiple sub-studies or treatment arms within a unified statistical and operational framework. They have already been recognized by the FDA and regularly used in oncology trials. The same model should be used for preapproving master protocols for priority respiratory pathogens.
2. **Build Research-Ready Populations.** Most Americans are willing to participate in clinical research when offered the opportunity.³⁰¹ Nevertheless, only a small percentage

298 Scott M. Berry, Connor J. Connor, and Roger J. Lewis, "The Platform Trial: An Efficient Strategy for Evaluating Multiple Treatments," *JAMA* 313 (2015): 1619–20.

299 RECOVERY Collaborative Group, "Dexamethasone."

300 Ford and Norrie, "Pragmatic Trials."

301 Allison Anderson, Deb Borfitt, and Ken Getz, "Global Public Attitudes About Clinical Research and Patient Experiences with Clinical Trials," *JAMA Network Open* 1 (2018): e182969.

of patients are ever invited to enroll in a study.³⁰² This gap reflects structural barriers rather than public reluctance. “Research-ready” cohorts should become a core component of national preparedness. These cohorts would consist of individuals who have agreed to be contacted regarding future studies and who reflect the geographic, demographic, and clinical diversity of the American population. Maintaining such cohorts would significantly reduce recruitment delays and allow priority studies to begin enrollment rapidly during emergencies.

- 3. Simplify Consent Documentation.** Participation should also be made easier once patients are identified. Evidence from behavioral science suggests that longer consent forms and increased cognitive burden reduce recruitment and compromise patient understanding.³⁰³ For minimal-risk emergency studies, consent procedures should focus on clarity and comprehension rather than legal complexity. Simplified consent forms—limited to two pages—can maintain scientific validity while improving participation.

Information Infrastructure

In responding to a crisis, consistent information matters both across the government agencies seeking to coordinate response and also to the public.

Data Infrastructure

The current US health data landscape is a fragmented collection of institutional silos and proprietary systems. Electronic health records, insurers, and clinical sites operate in isolation. Under this legacy model, researchers face months-long delays to obtain usable datasets due to complex contractual and governance barriers. During the COVID-19 pandemic, these delays meant that observational analyses took months to assemble data that could have informed national policy in days.

During the COVID pandemic, we saw a rare exception to these delays with the deployment of monoclonal antibodies. A specialized team utilized a direct data feed from several health systems and private pharmacies, moving from a doctor administering an infusion to an actionable analytical result in the hands of decision-makers in ten days. This success demonstrated that near-real-time outcome analysis is feasible when intermediaries are bypassed. The technical capacity for real-time data integration exists: The barriers are legal and political, not scientific.

To overcome these barriers, Congress should **designate clinical research data as a public utility and adopt a single, privacy-protective national patient identifier**. Such a tool would eliminate the constant duplication of records, enable enhanced pharmacovigilance, and vastly improve our ability to track long-term safety and efficacy across different care settings. To ensure the appropriate access to and use of the data, a member-governed, nonprofit data access broker could be created to function as a trusted intermediary that coordinates secure, standardized access to de-identified data without owning it. The secure data should be coupled with standardized, validated cybersecurity measures. The data broker could also establish emergency data authority,

302 Christopher P. Williams et al., “Demographic and Health Behavior Factors Associated with Clinical Trial Invitation and Participation in the United States,” *JAMA Network Open* 4 (2021).

303 Paul H. Beardsley et al., “Longer Consent Forms for Clinical Trials Compromise Patient Understanding: So Why Are They Lengthening?” *Journal of Clinical Oncology* 25 (2007): e13–e14.

modeled after the Defense Production Act, to compel temporary data sharing during a declared public health emergency. The United States should therefore treat clinical research data infrastructure as a public utility for public health preparedness.

In parallel, HHS should **establish a member-governed, nonprofit data access broker capable of coordinating secure, standardized access to research data**. This entity would not own patient information. Rather, it would serve as a trusted intermediary that reduces duplicative governance reviews, standardizes data-sharing agreements, and enables rapid analysis while preserving institutional autonomy. This should also be paired with appropriate cybersecurity and privacy protections.

Finally, Congress should establish **emergency data-sharing authorities** that can be activated during declared public health emergencies. Such authorities should be narrow, time-limited, transparent, and subject to oversight. Their purpose would be to ensure that critical evidence generation is not delayed by fragmented institutional decision-making during periods of national crisis. The objective is not unrestricted data sharing. It is the creation of a system capable of producing timely evidence while maintaining public trust.

Transparency Infrastructure for Public Trust

Public participation is essential to a learning health system. Patients cannot be expected to contribute data, enroll in studies, or trust research findings unless the system operates transparently. Transparency is particularly important when government agencies exercise emergency authorities. Public confidence depends on understanding both the evidence supporting decisions and the uncertainties that remain. This principle applies equally to clinical trial prioritization, data-sharing authorities, and Emergency Use Authorizations.

Transparency should therefore be treated as operational infrastructure rather than a communications exercise. Clinical trial systems function more effectively when participants, clinicians, and policymakers can understand what studies are being conducted, why they are being prioritized, and how decisions are being made.

All federally designated emergency trials should operate under a **radical transparency standard**. Protocols and protocol amendments should be publicly available. Enrollment progress should be reported through public dashboards. Results should be disseminated rapidly and in formats accessible to both researchers and the public. Participants should be informed not only about the risks of enrollment but also about the outcomes generated by the studies to which they contributed.

Transparency also requires strong privacy protections. Expanded data sharing, research-ready cohorts, and interoperable clinical trial systems must be accompanied by clear governance frameworks, cybersecurity standards, and meaningful accountability for misuse.

Steady-State Investment and Readiness

Industry engagement and incentives

To sustain a robust pipeline of new treatments, the federal government must address the unique economic challenges to innovation in the infectious disease sector. During peacetime, the market

for pandemic-related therapeutics is often nonexistent, which discourages private companies from investing the hundreds of millions of dollars required for drug development. When a crisis begins, companies must often pivot away from their primary commercial goals to address a public health emergency. Without clear financial and regulatory signals from the government, the most promising medical assets may never reach the clinical trial stage. The activation of federal economic levers can provide stability to the private sector long before a pathogen emerges. These incentives should be linked not only to product development and manufacturing readiness, but also to participation in prespecified clinical trial pathways, shared protocols, standardized assays, and standing platform studies.

The federal government should **reduce investor risk in pandemic-relevant research and development through advance market commitments and milestone-based payments**. Advance market commitments provide a guaranteed customer for successful products, which creates incentive for companies to enter the space. Similarly, milestone-based payments can support companies as they reach critical stages of development, such as the completion of Phase 2 safety trials. By sharing the burden of manufacturing capacity during peacetime, the government ensures facilities are ready to scale production the moment a treatment is authorized.

The federal government should also **establish structured, standing industry-government forums** to align government and industry expectations for the pace and process of trials. A formal industry-government forum should operate as a permanent bridge and prevent the improvised, crisis-driven negotiation that characterized the early COVID-19 response and wasted critical time.³⁰⁴ This body would clarify regulatory pathways and preclinical expectations during peacetime so that no time is wasted when an emergency is declared.³⁰⁵ By integrating these partnerships into the daily operations of our health system, we ensure that the private sector is a coordinated arm of the national response rather than a collection of independent actors.

Stress testing, exercising for emergencies, and government evaluation

Preparedness cannot be sustained through policy documents alone. Like military readiness, emergency management, or cybersecurity, clinical trial readiness must be exercised regularly if it is to function effectively during crisis. The rapid funding mechanisms, streamlined contracting, centralized review processes, data-sharing systems, and platform trial infrastructure detailed above all depend on relationships and procedures that can deteriorate if left unused.

The federal government should therefore conduct routine, **nationally coordinated simulation exercises designed to stress-test every component of the clinical trial enterprise**. These can serve as a critical diagnostic tool for the process itself. They allow institutions like the NCTEA (explained above) to identify bottlenecks in contracting, gaps in the centralized assay network, or delays in IRB approval before an actual pathogen emerges. The purpose of these exercises is not merely training. They should function as a mechanism for identifying weaknesses before a crisis occurs. The resulting findings should be used to improve performance and refine procedures.

304 For example, during the early months of 2020, communication between federal agencies and product developers was often fragmented and reactive. This led to confusion regarding which data elements the FDA would prioritize and which laboratory assays would be validated for official review.

305 This allows the government to provide informal, nonbinding guidance on evidentiary thresholds. This transparency allows companies to prepare their programs for emergency use long in advance. When the rules for licensure and the required data points are understood by all parties, the entire research ecosystem can move with synergy and efficiency.

To support accountability, the federal government should establish measurable readiness indicators and publicly report performance during these exercises. Relevant metrics could include time to protocol activation, time to first patient enrolled, time to contract execution, time to data availability, and demographic representativeness of enrolled participants. These indicators would provide a practical measure of preparedness while creating incentives for continuous improvement.

Section B Findings and Recommendations

Finding 6.2.1: The US government’s protocols for responding to a biological incident of unknown or contested origin, including the Biological Incident Notification and Assessment Protocol and the NSC-convened interagency response mechanism, depend on robust oversight and a well-trained civil service. Without a restored and regularly exercised decision architecture, the speed advantage that early detection and attribution can provide cannot be translated into effective response.

Recommendation 6.2.1.1: Restore and exercise the interagency biological incident decision architecture.

- The NSC should reestablish senior-level oversight of the Biological Incident Notification and Assessment Protocol and associated interagency response mechanisms.
- Congress should mandate and fund regular senior-level exercises of the biological incident response system across HHS, DOD, DHS, and the FBI, with performance metrics publicly reported.
- Data-sharing agreements between public health, law enforcement, and intelligence agencies that enable rapid response to an incident of unknown origin should be reviewed, updated, and confirmed in place before the next crisis.

Recommendation 6.2.1.2: Establish a National Clinical Trials Emergency Authority.

- Congress should authorize a standing National Clinical Trials Emergency Authority (NCTEA) within HHS with operational command over federally supported research during declared public health emergencies.
- The NCTEA should have authority to designate priority platform trials, mandate hospital participation, and enforce radical transparency standards for all federally designated emergency trials.

Recommendation 6.2.1.3: Establish a formal pathway for mutual recognition of clinical trial results from allied-country regulatory systems.

- The FDA should develop a formal framework for mutual recognition of clinical trial data generated by equivalently rigorous regulatory bodies belonging to US allies.
- Trials conducted to FDA-compatible standards should be eligible for FDA approval application without duplicative domestic trials.

Finding 6.2.2: The US clinical trial system cannot generate evidence at the speed of a biological crisis. COVID-19 demonstrated that a system optimized for hypothesis-driven, grant-funded research fails as a crisis response tool. The administrative, financing, and data infrastructure required for rapid, decisive evidence generation does not exist in peacetime and cannot be improvised under crisis conditions.

Recommendation 6.2.2.1: Build a crisis-ready learning clinical trial system as permanent national security infrastructure.

- Congress should provide the NIH with specific mandates for emergency surge capacity, including authority to deploy large-scale funding within days to weeks and a protected Emergency Research Reserve Fund with statutory triggers modeled on the G7 100-day mission frameworks.

Recommendation 6.2.2.2: Create a standing national clinical trial infrastructure.

- Congress should mandate that any hospital with over 100 beds participating in CMS maintain a standing clinical trial office, supported by pre-negotiated IDIQ reimbursement contracts; designate a permanent emergency clinical trial officer at each participating institution with authority to prioritize national protocols during a declared emergency.
- Congress should designate clinical research data as a public utility, adopt a privacy-protective national patient identifier, establish a member-governed nonprofit clinical trial data access broker with standardized and cyber-secure data, and authorize emergency data-sharing authority modeled on the Defense Production Act for use during declared public health emergencies.
- Congress and federal agencies should fund and maintain disease-agnostic platform trial infrastructures with preapproved master protocols for priority syndromes, and require the FDA to publish prespecified, pathogen-class guidance on evidentiary requirements and acceptable endpoints before the next biological emergency.

Finding 6.2.3: The private sector plays an indispensable role in developing and manufacturing pandemic-relevant therapeutics, but persistent market failures discourage sustained investment in infectious disease research and development during peacetime. Without predictable economic incentives and structured government engagement, promising medical countermeasures may never reach clinical testing, and critical manufacturing capacity may not be available when a biological emergency occurs.

Recommendation 6.2.3: Establish durable public-private mechanisms to sustain investment in pandemic-relevant therapeutics and manufacturing capacity.

- Congress should authorize and fund advance market commitments for priority pandemic countermeasures to provide predictable demand signals and reduce private-sector investment risk.
- HHS should establish a permanent industry-government coordination forum to align expectations regarding clinical trial requirements, regulatory pathways, manufacturing readiness, and emergency response planning before a crisis emerges.

Finding 6.2.4: Clinical trial readiness depends not only on infrastructure and authorities, but also on the routine exercise of the systems, relationships, and procedures that enable rapid research during a public health emergency. Without regular testing and evaluation, critical capabilities—rapid contracting, centralized review, data sharing, and trial activation—may degrade, creating delays when a biological crisis occurs.

Recommendation 6.2.4: HHS should conduct recurring national-scale exercises that stress-test the full clinical trial system. HHS should establish measurable readiness indicators—including time to protocol activation, time to first patient enrollment, time to contract execution, time to data availability, and participant representativeness—and publicly report performance results. The NCTEA (described above) should use these exercises to identify operational bottlenecks and produce after-action assessments with corrective action plans.

Resilience

A. The Interdependent Relationship Between Response and Resilience

Chapter 6 defined response as everything that is done in the immediate aftermath of a biological incident. Resilience describes whether and how well the United States can withstand an incident's consequences and, critically, how fast it can recover full functionality and reduce enduring harms. As a key element of the biosecurity equation, demonstrating national resilience also enhances deterrence.

Resilience is defined in many ways. For the purposes of this report, it has very specific meaning: **Resilience is the ability to withstand a significant biological incident and to recuperate from the effects as quickly and fully as possible and at the lowest societal cost.** Because the likelihood of preventing all naturally occurring, accidental, AI-enabled, and weaponized biological risks and threats is essentially zero, strengthening our national and international resilience is vital to enabling Americans and the global community to survive and thrive in a future in which bioconvergence is the norm.

Bioresilience is not a static state that is achieved and then maintained, but rather an ongoing, dynamic process. Resilience must evolve to keep pace with biotechnology innovation, private sector competitive motivations, democratization of access, and escalating concerns about the ways in which AI models can be misused. The United States must build a multidimensional ecosystem of initiatives and actions that, taken together, create a layered, continuous, multimodal system of resilience.

This report identifies simultaneous and complementary steps that need to be taken to rapidly and systematically build a system capable of withstanding and recovering from current and future biological threats. These actions must begin in the homeland but must necessarily extend overseas, especially with respect to surveillance and suppression of emerging threats. This is increasingly urgent given the pace of innovation, the economic and national security imperatives that are driving the competitive landscape of that innovation, and the growing and potentially existential risks that this ongoing innovation creates.

This section addresses the structural investments required to build durable national resilience against AI-enabled biological threats. It describes three elements that are necessary to achieve this objective:

- the data, continuous compute capacity, and operational response infrastructure that must be built;

- the international engagement frameworks that are needed to deal with threats that know no borders; and
- the conditions that would enable the US biotechnology industry to thrive in the face of its growing dependence on and fierce competition with China.

The Strategic Imperative for Investment in Biological AI

The United States faces serious and growing biological risks. As international coalitions fracture and great power competition intensifies in the AI-enabled biological landscape, the risk of falling behind at the intersection of artificial intelligence and biotechnology has concrete national security consequences. Past examples of technological watersheds provide excellent models for how to respond. The Manhattan Project and the creation of the National Laboratory science and technology innovation system secured US nuclear leadership. DARPA's investment in ARPANET created the internet.³⁰⁶ SEMATECH, the government-industry consortium, rebuilt American chip manufacturing competitiveness and ultimately enabled the extreme ultraviolet lithography breakthrough that today underpins the global semiconductor supply chain.³⁰⁷ The nation that builds the foundational data and model infrastructure for AI-enabled biology and maintains private sector innovation and resilience will generate a comparable advantage in this century's defining bioeconomy.

The US lead in biological AI is currently sustained almost entirely by private and philanthropically funded organizations such as Google DeepMind's AlphaFold series (recognized with the 2024 Nobel Prize in Chemistry), the Chan Zuckerberg Biohub's Virtual Biology Initiative (a \$500 million philanthropic commitment to AI-powered models of human cells), and the Arc Institute's Evo 2.12. This is a fragile foundation. Foreign state-funded competitors operate at comparable or greater scale—e.g., BioMap's trained xTrimoPGLM 100 billion parameter model, ByteDance's Protenix in molecular structure prediction—with the patience and concentrated resources that centrally directed national investment provides, while the US position depends on the continued willingness of select frontier firms and philanthropies to subsidize what the nation must depend upon.

Federal investment is not a substitute but a necessary complement to this private sector activity in order to maintain biotech leadership and protect national security. First, biodefense applications of biological AI will not be funded by private capital and philanthropy at the scale required because dual-use evaluation, classified threat characterization, and operational countermeasure design have no commercial market and no philanthropic mission fit; these are inherently federal responsibilities. Second, the generation of verified functional biological data at industrial scale—the rate-limiting input for every model described above—is a shared-infrastructure problem that no single company, philanthropic organization, or laboratory will solve alone, and which the federal government is uniquely positioned to coordinate.

306 GAO, *Federally Created Entities: An Overview of Key Attributes*, GAO-10-97, October 2009, <https://www.gao.gov/assets/gao-10-97.pdf>.

307 GAO, *Lessons Learned from SEMATECH*, GAO-92-283, September 1992, <https://www.gao.gov/assets/rced-92-283.pdf>.

Why Functional Data Drives Breakthroughs

Achieving defensive overmatch against AI-enabled biological threats depends on solving one problem: generating functional biological data at scale, paired with large-scale compute. AI in biology faces two distinct challenges. First, biological models are prone to confabulation—generating plausible but empirically false outputs. Better training data reduces this risk. Second, the field lacks wet-lab validation infrastructure to verify AI predictions at scale. Both problems must be solved in parallel.

Every major advance in biological AI has been built on publicly funded data. AlphaFold trained on the Protein Data Bank (PDB), which includes more than 200,000 experimentally determined structures curated across five decades of US and international investment. Arc Institute's Evo 2 trained on NCBI's Sequence Read Archive (SRA) and IMG/VR, both publicly funded US databases.³⁰⁸ The pattern replicates the dynamic that produced the large language model revolution: GPT was made possible by open internet text. In each case, whoever controlled the training data controlled the field's direction.

Sequence data, however, is not enough. Current models lack the functional annotations needed to understand what sequences do. This is the gap that separated the Human Genome Project's raw sequence from clinically actionable genomics over nearly two decades. The United States must expand the available data, secure who has access to the highest-risk data without stifling innovation, and compress the timeline to translate it into health interventions. One potential solution is the creation of a national public-private partnership for biodata[?].

B. Building Bioresilience Through AI, Functional Data, and Response Infrastructure

US government funding has historically been successful at driving progress in critical industries. The US funded the PDB, creating the data foundation that enabled AlphaFold, and, as noted previously, the US government-led SEMATECH restored US chip leadership in the 1980s and '90s,³⁰⁹ revitalizing innovation across fourteen competing chip makers to reduce manufacturing costs and product defects, and compress miniaturization cycles.

A similar effort for biotechnology would yield lasting benefits to US national security and economic strength, as well as the long-term viability of the US biotech sector. Using the model of SEMATECH, the US could construct a public-private partnership to build a strategic biological data and model commons, along with high-performance compute, for biological AI. The new entity could provide secure, large-scale compute; AI-ready biological data products; and AI-driven tools and testbeds for capability assessment, threat characterization, and countermeasure development and deployment necessary for long-term US bioresilience.

308 National Center for Biotechnology Information, "Sequence Read Archive (SRA)," accessed June 18, 2026, <https://catalog.data.gov/dataset/sequence-read-archive-sra>; Joint Genome Institute, "IMG/VR: Genomes of Cultivated and Uncultivated Viruses," accessed June 18, 2026, https://genome.jgi.doe.gov/portal/IMG_VR/IMG_VR.home.html.

309 Robert D. Hof, "Lessons from Sematech," *MIT Technology Review*, July 25, 2011, <https://www.technologyreview.com/2011/07/25/192832/lessons-from-sematech/>.

The need for this type of capability is well established. The September 2025 National Security Commission on Emerging Biotechnology called on Congress to authorize “a comprehensive central biological data infrastructure” with built-in security and access controls. NSCEB refers to this as a Web of Biological Data (WOBD).³¹⁰ As described in Chapter 6, the November 2025 Genesis Mission Executive Order similarly directs the Department of Energy to build an American Science and Security Platform integrating supercomputers, scientific datasets, foundation models, and autonomous experimentation tools.³¹¹ Positive early initiatives are underway: The Department of Defense’s Chemical and Biological Radiological Nuclear Defense (CBRND) program is establishing a Digital Biosecurity Forge as an initial implementation of the concept. Building on these forward-looking steps, a more comprehensive approach is needed.

Enabling End-to-End Resilience

The bioresilience challenge starts with data, which feeds almost all subsequent steps. The biodata scale challenge is immense. Frontier large language models train on 15 to 18 trillion tokens of text. Biology has no comparable reservoir. Functional annotation data such as protein interactions, immune evasion profiles, and drug responses represent a tiny fraction of what exists as raw sequence. Closing this gap requires generating orders of magnitude more verified functional data than current manual methods can produce. Self-driving laboratories, which combine AI-directed experimental design with robotic execution, have demonstrated up to ten times throughput gains over manual workflows with twenty-four-hour continuous operation. As a result, US national capability must include secure autonomous wet-lab infrastructure and cloud lab partnerships to industrialize functional data generation at the scale these models demand.

Access to compute resources, as discussed previously, is a limiting factor for data. Training Evo 2 required more than 2,000 NVIDIA H100 GPUs running for two to three months (cost in excess of \$30 million).³¹² Such resources are beyond the reach of most academic labs and small biotech firms. As a result, a government-led capability must include a dedicated compute allocation available to credentialed researchers and partners to bridge the gap between frontier model development and the broader innovation base.

Building American Resilience Infrastructure

Congress should establish a **federally chartered, public-private entity providing secure, large-scale compute, AI-ready biological data products, and AI-driven tools and testbeds focused on biodefense and bioresilience (including capability assessment, threat characterization and countermeasure development)**. This entity would be designed to work interoperably with the National Strategic Bio Readiness Reserve described in Chapter 6. For example, while the compute reserve described as part of the NSBRR would be a portion of compute made available during a crisis, the computing power needed to support national resilience, and described below, would be continuously maintained and used by the federal government. In a crisis, this continuous compute could be augmented quickly using the reserve capacity envisioned under the NSBRR.

310 NSCEB, *Charting the Future of Biotechnology*.

311 Exec. Order No. 21665, “Launching the Genesis Mission.”

312 ARC Institute, “Evo 2: A New Foundation Model for Genomic Discovery,” ARC Institute, Feb. 19, 2025, <https://arcinstitute.org/news/evo2>.

This broad capability is structured as a three-level pyramid, with each level providing progressively more restricted access to increasingly sensitive capabilities. The three tiers are designed to be mutually reinforcing: The broad innovation base generates the scientific advances that inform high-containment defense research, which in turn equips the operational teams that deploy when a biological incident occurs.

Base: Innovation Ecosystem Enablement. The broadest tier provides compute, data, and model access to universities, startups, and research institutions working on conventional, non-pandemic pathogens and capabilities broadly to bioresilience. Work at this level does not require elevated biosafety controls and is designed to be openly accessible to credentialed researchers. America's competitive advantage in biotechnology has always rested on its distributed innovation ecosystem in which thousands of independent laboratories pursue independent lines of inquiry. **This base tier converts that distributed energy into a coordinated national capability by giving every qualified researcher access to the foundational data, compute, and models needed to contribute to biological AI breakthroughs.** This is the tier characterized by the Genesis Mission's "Scaling the Biotechnology Revolution" challenge, which uses the DOE National Laboratory infrastructure as the enabling platform for broad-based biological AI research.


Middle: High-Containment Defense Research and Development. The middle tier supports federal research on high-risk pathogens, including BSL-3 and BSL-4 work, harmful biological design, gain-of-function research, and rapid countermeasure development. Access would be restricted to government laboratories, including DOE National Laboratories, USAMRIID, and NBACC, working alongside specialized private-sector partners with existing biodefense expertise, high-containment infrastructure, and appropriate security clearances. Sensitive functional data generated at this tier, including information on virulence, transmissibility, or immune evasion that could enable the design of more dangerous pathogens, is governed by a standing Dual-Use Biological Review Board and stored in secure enclaves with strictly controlled access. Export restrictions on sensitive functional datasets prevent adversarial acquisition through commercial channels.

This tier is also precisely the **environment in which dual-use risk from AI-bio convergence can be responsibly evaluated.** Benchmarking the biological capabilities of frontier AI models requires testing them against the specific tasks that define uplift risk: designing enhanced pathogens, identifying immune evasion strategies, optimizing transmissibility. Those evaluations cannot be performed safely at the open, base tier, and they cannot be performed credibly outside an environment with high-containment wet-lab validation, classified threat intelligence integration, and a standing Dual-Use Biological Review Board governing both the evaluations themselves and the resulting data. As discussed in Chapter 3, this sort of evaluation will require intensive collaboration between cleared experts across the US government, under either a central federal entity or a fusion-cell model in which experts representing their home agencies are co-located and tackling common issues together.

Peak: Operational Response. The apex of the pyramid is the operationalization of the biodata capabilities in teams modeled on the Department of Energy's Nuclear Emergency Support Team (NEST), which has provided rapid nuclear incident response to the nation

and the world for five decades.³¹³ The biodefense incarnation of this kind of capability, a proposed Bio Emergency Services Team (BEST), would rapidly deploy directly to the site of a biological incident, bringing advanced diagnostic tools, AI-enabled threat characterization, and reach-back capabilities to necessary technical expertise. Optimally these teams would serve as a national biological fire department for both localized public health threats and catastrophic incidents. A reserve network of precertified private-sector partners, as part of the NSBRR described in Chapter 6 above, would provide the surge capacity that BEST teams would draw on during large-scale activations.

Figure 1. Bioresilience Infrastructure Pyramid



Governance and Security

This proposed multitiered bioresilience infrastructure segregates sensitive functional data —such as information on virulence, transmissibility, or host range that could enable the design of more dangerous pathogens—into secure enclaves. Access to this sensitive data could be governed by specialized Dual-Use Biological Review Boards.³¹⁴

The tiered access controls must also include export restrictions on sensitive functional datasets to prevent adversarial acquisition of US biological data through commercial channels.³¹⁵ Current biosafety regulations, data security guidance, and collaboration models supply a partial regulatory framework. Department of Commerce and HHS regulations require registration, designated responsible officials, security barriers, and security-risk assessments for access to select agents

313 The CDC currently fields a Global Rapid Response Team in which personnel at CDC provide surge capacity and consultative support to host country in a crisis. It is primarily intended for public health outbreaks. NEST, in contrast, would provide unified incident response and consist of a dedicated cadre of personnel with integrated technical and reach-back capabilities. It is suited to address accidental and deliberate bio incidents in addition to natural ones.

314 The Review Board would apply current regulatory and controls-adapted criteria with structured appeals, consistent with NSCEB Recommendation 4.2a. NSCEB, *Charting the Future of Biotechnology*.

315 Ibid., § 4.2a.

and toxins.³¹⁶ NIST has provided guidance to manage cybersecurity and privacy risks that arise from processing genomic data.³¹⁷ These authorities provide a concrete basis for tiered access controls, misuse prevention measures, and high-risk technology collaborations between government labs and the private sector.

Internally, the governance framework ensures that the most sensitive biological data and capabilities are accessible only to credentialed actors with a legitimate research or defense purpose. Externally, it becomes a model for Biological Weapons Convention (BWC) Article IV national implementation and the basis for US Confidence-Building Measure submissions, setting an example of best practices and potentially creating pressure on China, Russia, Iran, and North Korea to match American transparency.

Legislative Gap Analysis for a Bioresilient Data and Deployment Infrastructure

A detailed legal gap analysis is provided in Appendix 5. A summary is provided here for reference. Several categories of legislative action are needed to implement the proposed infrastructure:

- Current law provides no express statutory authorization for a durable federally chartered entity of this kind, leaving the initiative dependent on executive action and vulnerable to changes in administration;
- Existing privacy law and confidentiality protections were not designed for a centralized biological AI database, creating uncertainty around data access, use, and intellectual property;
- Procurement and partnership authorities may require explicit adaptation for AI-bio infrastructure and rapidly scaling research programs;
- There is no biological equivalent of the mature statutory foundation that underpins NEST, meaning that absent legislative action the BEST teams would operate without a clear chain of command, standing readiness obligation, or federal-state coordination authority.
- The existing PREP Act liability protections are too narrow for the full range of peacetime and emergency response activities;
- No dedicated appropriations and funding mechanisms exist to sustain readiness between operational activations, undermining the goal of establishing a “warm start” foundation for incident response.

316 42 C.F.R. § 73.10 (restricting access to select agents and toxins; security risk assessments); 15 C.F.R. pt. 774 (Technology Export Controls on Instruments for the Automated Chemical Synthesis of Peptides). The Bureau of Industry and Security at the Department of Commerce has also relied on the Export Control Reform Act of 2018 to establish export controls on emerging biotechnologies.

317 National Institute of Standards and Technology, *Cybersecurity of Genomic Data*, NIST IR 8432, December 2023, <https://doi.org/10.6028/NIST.IR.8432>; National Institute of Standards and Technology, *Genomic Data Cybersecurity and Privacy Frameworks Community Profile*, NIST IR 8467 2pd, December 2024, <https://doi.org/10.6028/NIST.IR.8467.2pd>.

C. Industrial Resilience: Strengthening the Private Sector and Supply Chain

As stated at the outset of this report, the United States must lead in biotechnology to secure biotechnology. That is the virtuous cycle that can enable the United States to manage the risks of bioconvergence. The US private biotechnology sector is the key to bioresilience: Its strength and innovation are essential for every element of deterrence of the cycle. US national and economic security require that we lead in biotechnology.

Although some elements of the AI race and biotechnology competition resemble the technological rivalries of the past, it is clear that the deterrence frameworks of the past do not map neatly onto the challenges posed by bioconvergence. Increasingly, biological capabilities will be accessible to a broad range of state and nonstate actors. The ability to develop countermeasures, characterize threats, and respond rapidly, however, will reside primarily in the private sector. Unlike previous eras, the US government alone cannot maintain the constantly evolving technological edge required for biosecurity and bioresilience. The private sector is therefore not simply a supplier to government but an indispensable national security partner with literally vital responsibilities.

The United States possesses the scientific talent, innovation ecosystem, and entrepreneurial culture necessary to lead the world in biological AI. What it lacks is a coherent federal strategy that treats the domestic biotechnology industry as a national security partner. Today, federal biotechnology investments are substantial in aggregate but fragmented across agencies, focused on discrete threats, and often structured around short-term emergency responses rather than sustained platform development. Meanwhile, American firms increasingly compete not only against foreign companies but against foreign governments that have made deliberate, coordinated investments in biological infrastructure, manufacturing, and innovation ecosystems.

Strategic Competition and the China Challenge

China's growing dominance in global pharmaceutical research and manufacturing presents a significant risk to the US biotechnology sector, US national security, and homeland resilience. The shift to China as a source for active pharmaceutical ingredients, key starting materials, finished drugs, contract manufacturing services, and biotechnology research and development is the result not only of market forces responding to lower labor, production, and environmental compliance costs, but also of targeted long-term industrial policy choices made by the Chinese government intended to strengthen domestic pharmaceutical and biotechnology capabilities and increase Chinese firms' global competitiveness.^{318, 319, 320, 321}

318 U.S. Government Accountability Office, COVID-19: Critical Vaccine Distribution, Supply Chain, Program Integrity, and Other Challenges Require Focused Federal Attention, GAO-21-265 (January 2021), <https://www.gao.gov/products/gao-21-265>

319 U.S. Food and Drug Administration, "Safeguarding Pharmaceutical Supply Chains in a Global Economy," testimony before the House Committee on Energy and Commerce, Oct. 30, 2019, <https://www.fda.gov/news-events/congressional-testimony/safeguarding-pharmaceutical-supply-chains-global-economy-10302019>

320 Karen M. Sutter, Made in China 2025 and Industrial Policies: Issues for Congress, CRS Report IF10964 (Congressional Research Service, updated Dec. 12, 2024), <https://www.congress.gov/crs-product/IF10964>

321 U.S.-China Economic and Security Review Commission, Made in China 2025: Evaluating China's Performance (Nov. 14, 2025), <https://www.uscc.gov/research/made-china-2025-evaluating-chinas-performance>

China's pharmaceutical and biotechnology strategy has followed a recognizable pattern: Enter with generics, move to process chemistry, develop translational manufacturing capability, and then use that position to generate insight into innovators' proprietary methods and to move up the value chain. In the era of bioconvergence, this pattern is accelerating. China is no longer only a low-cost manufacturer. It is increasingly a competitor in biological research, sequencing platforms, contract manufacturing, clinical development, biological AI, and drug discovery. In a 2026 survey of senior US biotechnology and academic leaders, 85 percent of respondents believed the United States' lead in biotechnology quality and development would last ten years or less, with China already leading in two sectors: clinical development and supply chain.³²²

This creates a dual challenge for the United States. First, US reliance on Chinese materials and manufacturing poses direct national security and economic risks. Second, China's growing role in biological innovation threatens to erode the domestic discovery ecosystem that the United States will need in a future biological crisis. A biodefense posture that depends on private-sector speed and innovation cannot remain resilient if the underlying industrial base is hollowed out.

Supply Chain Vulnerabilities

Technical capability without supply chain security is incomplete and unsecure. The United States has learned this lesson twice in recent years: first with rare earth minerals, where Chinese dominance over mining and processing created leverage over US defense manufacturing, and again with semiconductors, where concentration of advanced chip fabrication exposed a critical single point of vulnerability in the supply chain.

Biotechnology is following the same trajectory. US pharmaceutical and biotechnology supply chains already have significant direct and indirect China-linked dependencies, including key starting materials, active pharmaceutical ingredients, antibiotics and other generic-drug inputs, contract development and manufacturing services, and genomic-sequencing products and services. Increasingly, China is also becoming an important source of drug candidates licensed by major pharmaceutical companies for further development. The extent of these dependencies varies substantially by product and technology, and federal agencies have noted that incomplete supply-chain data make precise national estimates difficult.^{323, 324, 325, 326}

A US biodefense posture that is technically excellent but physically dependent on an adversary's industrial base is not resilient. It is a capability that can be degraded or denied without a single biological event occurring. Supply chain security is therefore more than a logistics challenge; it is central to deterrence and resilience.

322 Deena Beasley, "China Closing in but US Still Leads in Biotech Innovation, Survey Finds," *Reuters*, June 22, 2026, <https://www.reuters.com/business/healthcare-pharmaceuticals/china-closing-us-still-leads-biotech-innovation-survey-finds-2026-06-22/>.

323 U.S. Food and Drug Administration, "Exploring the Growing U.S. Reliance on China's Biotech and Pharmaceutical Products," testimony before the U.S.-China Economic and Security Review Commission, July 31, 2019, <https://www.fda.gov/news-events/congressional-testimony/exploring-growing-us-reliance-chinas-biotech-and-pharmaceutical-products-07312019>

324 Michael D. Sutherland and Karen M. Sutter, COVID-19: China Medical Supply Chains and Broader Trade Issues, CRS Report R46304 (Congressional Research Service, updated Dec. 23, 2020), <https://www.congress.gov/crs-product/R46304>

325 National Security Commission on Emerging Biotechnology, Charting the Future of Biotechnology (Apr. 2025), <https://www.biotech.senate.gov/final-report/>

326 U.S.-China Economic and Security Review Commission, Made in China 2025: Evaluating China's Performance (Nov. 14, 2025), <https://www.uscc.gov/research/made-china-2025-evaluating-chinas-performance>

Bio supply chain vulnerabilities are currently concentrated in two categories. The first is synthetic chemistry and fine chemistry, which produce small-molecule pharmaceuticals and their precursor chemistry. Small-molecule drugs are generally chemically synthesized compounds, often made through multistep organic chemistry, that can usually be formulated as tablets, capsules, or injectable products.³²⁷ This capability was systematically offshored over the past two decades, driven by cost, environmental regulatory pressure, and the absence of a national strategy to retain it. The second is the emerging frontier of biologics: large, complex therapies derived from or manufactured using living systems, such as monoclonal antibodies, recombinant proteins, vaccines, cell therapies, and gene therapies.³²⁸ Unlike small molecules, biologics are generally harder to copy, harder to manufacture, and more dependent on specialized facilities, sterile production environments, and platform-specific expertise.³²⁹

China has not yet achieved dominance in biologic manufacturing but is closing in. This distinction matters because China's greatest present leverage is in small-molecule inputs and chemistry-based contract development, while the United States retains stronger positions in many frontier biologics and advanced therapeutic platforms. The strategic risk is that these categories are converging. As China expands from chemistry into biologics, the United States could become vulnerable across both major therapeutic domains. Small molecules and biologics matter for different but equally important reasons. Small molecules form the backbone of routine medical care and become especially critical during a public health emergency, when hospitals must manage large numbers of severely ill patients with supportive medications, antibiotics, antivirals, steroids, sedatives, vasopressors, and other essential drugs. Biologics, by contrast, support many of the advanced therapeutic platforms that may be needed to develop targeted countermeasures against novel threats, including next-generation vaccines and other precision biological interventions.

A resilient biodefense posture, therefore, requires domestic strength in both the medicines needed to manage patients during a crisis and the platforms needed to develop new interventions when existing tools are insufficient. Policy interventions that treat the biotechnology supply chain as uniform will misallocate resources. The near-term priority should be to invest in reshoring (and “friend-shoring”) fine chemistry and manufacturing of synthetic biologics; the medium-term priority should be ensuring that frontier therapeutic platforms do not follow the same trajectory before domestic alternatives are established.

APIs and KSMs

A finished drug contains two components: active pharmaceutical ingredients (APIs), which are derived from key starting materials (KSMs), and inactive ingredients that help deliver, stabilize, or formulate the medicine. The API is the biologically active compound that produces the therapeutic effect. APIs are rarely manufactured by the companies that sell the final drug. They are produced separately, often by specialized chemical manufacturers, and then formulated downstream into tablets, capsules, injectables, or other delivery solutions. KSMs are the chemical building blocks from which APIs are synthesized. From a supply chain perspective, a country can track

327 National Institutes of Health, Office of Extramural Research, Small Business Education and Entrepreneurial Development (SEED), *Regulatory Knowledge Guide for Small Molecules* (March 2024), <https://seed.nih.gov/sites/default/files/2024-03/Regulatory-Knowledge-Guide-for-Small-Molecules.pdf>.

328 US Food and Drug Administration, *Biological Product Definitions* (n.d.), <https://www.fda.gov/files/drugs/published/Biological-Product-Definitions.pdf>.

329 Andrew W. Lo, Robert F. Nelsen, and Bob More, “The Golden Age of Biotechnology,” *Nature Biotechnology* 42, no. 12 (2024): 1805–1808, <https://doi.org/10.1038/s41587-024-02401-3>.

API suppliers while remaining blind to the upstream KSM dependencies that make those APIs possible. For instance, the United States gets many APIs and drugs from India, but India gets much of its KSM supply from China.

This layered dependency creates serious vulnerabilities. In 2023, more than 90 percent of drugs prescribed in the United States were generics, and most were manufactured abroad.³³⁰ Roughly 80 percent of pharmaceutical APIs are sourced overseas, much of that directly or indirectly from China. These dependencies are not theoretical. Estimates of US reliance on Chinese-made APIs have varied widely, ranging from 8 percent to as much as 80 to 90 percent, with the Brookings Institution having estimated in 2025 that perhaps one quarter of all drug volume sold in the US includes Chinese API.³³¹ In 2024, the United States recorded record-high drug shortages affecting hundreds of medications, many tied to delays or interruptions in the API and KSM supply chain. Antibiotics were among the drugs most affected. These shortages threaten individual health, national preparedness, and strategic stockpiles.³³²

China is the United States' largest foreign supplier of critical pharmaceutical inputs by volume, accounting for approximately 40 percent of imports in 2024. Its dominance is most acute in generic drugs and upstream inputs: The US relies on China for 99 percent of imported prednisone, 92 percent of penicillin and streptomycin antibiotics, and 41 percent of KSMs used in API synthesis.³³³ China also controls the raw or starting materials for 94 percent of amoxicillin, 74 percent of heparin, and 70 percent of acetaminophen.³³⁴ For one in ten critical pharmaceutical inputs, China's market share exceeds 99 percent. These dependencies are concentrated most heavily in chemistry-based and small-molecule supply chains, which underscores why a strategy focused only on finished drugs or only on advanced biologics would fail to address the most immediate vulnerabilities. Simultaneously, biologics and advanced therapies should not be treated as safe simply because the United States remains comparatively stronger in those markets today. China's increasing role in biologics development, licensing, and contract manufacturing suggests that the same dependency pattern could extend into higher-value therapeutic platforms if the United States does not act early. This concentration creates vulnerability to deliberate coercion and to other types of market disruption caused by accidents, natural disasters, and manufacturing failures at facilities with outsized global market share.

Sequencing Platforms

DNA sequencing occupies a distinct position in the biotechnology supply chain. It is the foundational analytical layer on which drug discovery, pathogen characterization, biosurveillance, and countermeasure development depend. Controlling sequencing platform infrastructure means controlling the data generated by it. That data constitutes the training material for next-generation biological AI and the basis for biological research.

330 Dave Gryska, *America's National Security Vulnerability: Generic Drug Manufacturing and Biotechnology Innovation* (Business Executives for National Security, January 2026), <https://bens.org/wp-content/uploads/2026/01/BENS-Biotechnology-Report.pdf>.

331 Marta E. Wosińska and Yihan Shi, "US Drug Supply Chain Exposure to China: Myths, Omissions, and Related Insights," *Brookings Institution*, July 28, 2025, <https://www.brookings.edu/articles/us-drug-supply-chain-exposure-to-china/>.

332 Gryska, *America's National Security Vulnerability*.

333 Atlantic Council, "Pharmaceuticals Are China's Next Trade Weapon," Feb. 13, 2026, <https://www.atlanticcouncil.org/blogs/econographics/sinographs/pharmaceuticals-are-chinas-next-trade-weapon/>.

334 Thomas J. Bollyky et al., *The Pharma Choke Point: How to Reduce US Dependence on Chinese Pharmaceutical and Biotechnology Supply Chains* (Council on Foreign Relations, 2026), https://assets.cfr.org/images/The-Pharma-Choke-Point-by-Bollyky-Yadav-Kosloff-Yadav-and-Every_June-2026_final-6-4-1/The-Pharma-Choke-Point-by-Bollyky-Yadav-Kosloff-Yadav-and-Every_June-2026_final-6-4-1.pdf.

The United States dominated sequencing platforms for decades. Recognizing the strategic significance of sequencing platforms, China has deliberately undermined that position. Twelve years ago, Chinese company BGI acquired Complete Genomics, an American sequencing technology company, gaining access to foundational intellectual property that it used to develop its MGI sequencing platform. The resulting product line is now heavily subsidized by the Chinese government and priced below cost. This strategy replicates China's competitive approach to solar panels and electric vehicles.

BGI and MGI now face restrictions under the BIOSECURE Act (described below), but the broader sequencing market continues to face intense pressure from Chinese platforms whose low prices are based on government subsidies. Domestic genomic sequencing platform capacity should therefore be treated as a critical supply chain priority alongside APIs, KSMs, and contract manufacturing.

Contract Development and Manufacturing Organizations

China has developed substantial dominance in contract development and manufacturing organizations, particularly in fine chemistry and translational scale-up. One firm, WuXi AppTec, commands an outsized share of the process-development market for small molecules. That concentration creates several national security concerns: supply-chain leverage, visibility into proprietary processes across large portions of the global pharmaceutical pipeline, and potential coercive advantage in a geopolitical crisis. The BIOSECURE Act, enacted in 2025, restricts US federal contracts with certain Chinese biotechnology companies of concern, including WuXi AppTec, as a partial response to these risks.

The BIOSECURE Act demonstrates the impact of legislative action. Enacted in 2025, it is already having positive effects on the US biotechnology industry. BIOSECURE is accelerating a significant rebalancing of the geography of pharmaceutical manufacturing that has existed since China became a low-cost hub in the 2000s.³³⁵ Even before enactment, legislative pressure was sufficient to force structural change: WuXi AppTec announced the sale of its US and UK cell and gene therapy operations, WuXi Advanced Therapies, to New York-based Altaris in December 2024, a year before the act became law. Companies that have long relied on WuXi AppTec, WuXi Biologics, and other Chinese CDMOs are now securing alternative manufacturing capacity in the United States, Europe, South Korea, and India.³³⁶

Clinical Development and Trial Execution

Clinical trial execution is emerging as a separate strategic vulnerability. Clinical trials are not merely a regulatory step at the end of drug development; they are the translational infrastructure that determines whether discovery can become deployable medical capability. The United States remains the world's most important commercialization market, but, as discussed in Chapter 6 of this report, its clinical trial system is slow, expensive, fragmented, and increasingly difficult

335 Vision Lifesciences Strategy Team, "CDMO Market Analysis 2026: Size, Trends and Key Players," *Vision Lifesciences*, Feb. 28, 2026, <https://visionlifesciences.com/insights/cdmo-market-analysis>.

336 Ibid.

to navigate. China has built a competing model that is faster, cheaper, and more centralized, with one 2024 estimate suggesting clinical trial costs are generally 25 to 40 percent cheaper in China than in the United States. Its advantages include large hospital networks, rapid patient recruitment, expanding investigator capacity, lower direct trial costs, and a regulatory system that has become more efficient over the past decade.³³⁷ These advantages are already changing global investment behavior. China now rivals or exceeds the United States in important measures of clinical trial volume,³³⁸ and US and European pharmaceutical companies are increasingly licensing Chinese-developed drug candidates or relying on Chinese development infrastructure to accelerate pipelines.³³⁹

This creates a national security problem that is distinct from, but connected to, manufacturing dependence. A biodefense strategy that depends on rapid private-sector innovation also depends on the ability to test, validate, and iterate medical countermeasures quickly. If clinical development capacity migrates offshore, the United States will lose more than individual trials. It will lose investigator networks, trial-site readiness, patient recruitment infrastructure, regulatory expertise, clinical operations talent, and the tacit knowledge that allows a biotechnology ecosystem to move quickly in a crisis. In a future public health emergency, especially one involving an adversarial state or a deliberately engineered biological threat, dependence on Chinese clinical trial infrastructure could create delay, denial, coercion, or data-access vulnerabilities at precisely the moment when speed would be most important. Even if no access were formally denied, geopolitical tension could slow data transfer, complicate regulatory review, restrict collaboration, or make US firms reluctant to depend on Chinese trial sites for countermeasure development.

The longer-term risk is innovation displacement. Clinical trials are where capital, data, patients, investigators, and intellectual property converge. If more early development, proof-of-concept testing, and first-in-human work occurs in China, then more of the economic value of biotechnology will accumulate there as well. Revenue, investor attention, licensing activity, clinical data, and scientific talent will follow the trial infrastructure. Over time, this could weaken US biotechnology clusters in the same way that offshoring manufacturing weakened domestic production capacity: slowly at first, then suddenly once the supporting ecosystem has moved elsewhere. The policy goal should not be to prevent all international clinical collaboration, which can accelerate access to important medicines, but to ensure that the United States retains enough domestic clinical trial capacity to develop, test, and deploy critical countermeasures without dependence on a strategic competitor.

AI-accelerated Intellectual Property Erosion

AI tools can now read a published patent and generate chemical structures that replicate the likely therapeutic mechanism of the patented molecule while altering features sufficient to avoid formal infringement. What previously required years of chemistry work can now be accomplished in a fraction of that time. The practical consequence is that the interval between a US innovator publishing a patent and the emergence of structurally similar Chinese competitors, which

337 NSCEB, *Charting the Future of Biotechnology*.

338 Beasley, "China Closing in but US Leads in Biotech Quality."

339 Michael Erman, "Pfizer CEO Says US Pharma Industry Needs to Collaborate with China," *Reuters*, Oct. 14, 2025, <https://www.reuters.com/business/healthcare-pharmaceuticals/pfizer-ceo-says-us-pharma-industry-needs-collaborate-with-china-2025-10-15/>; Andrew Silver and Yadarisa Shabong, "China's Innovent Biologics, Pfizer Strike Up \$1.05 Billion Cancer Drug Deal," *Reuters*, May 28, 2026, <https://www.reuters.com/legal/litigation/chinas-innovent-biologics-pfizer-strike-up-105-billion-cancer-drug-deal-2026-05-28/>.

previously provided several years of effective market exclusivity, is compressing significantly. This undermines the value that justifies the decade-long, billion-dollar investments required to develop novel therapeutics.

The problem is compounded by the structure of international patent disclosure requirements. Under the Patent Cooperation Treaty,³⁴⁰ patents are typically published within eighteen months of the application being filed. This creates a window within which competitors can begin design-around work before the patent is granted. This timeline was designed before AI-assisted drug design made rapid chemical variation possible. It is now a vulnerability that AI-enabled competitors can exploit. The result, if unaddressed, is a progressive degradation of the invention incentives that sustain US pharmaceutical leadership.

Strategic Tradeoffs

The United States faces two parallel challenges: managing first-order biosecurity risks, where China is both a source of concern and a potential partner in reducing global biological threats, and determining how much economic interdependence in biotechnology is acceptable before it becomes a strategic vulnerability. Building resilience may require accepting higher financial costs to sustain domestic capabilities, while openness and collaboration can accelerate innovation and improve access to new technologies. The central question for policymakers is not whether to compete or cooperate, but to effectively manage the enduring tension in which resilience requires strengthening domestic capacity while carefully managed engagement has the potential to reduce shared risks.

This tension is illustrated by the rapid growth in Western licensing of Chinese-developed therapeutics. Supporters argue that China's biotechnology sector still depends heavily on access to the US market, including FDA approval, American commercialization partners, and US capital markets. Under this view, the United States benefits from faster access to innovative medicines while retaining leverage over commercialization and market access. Critics, however, argue that large-scale reliance on Chinese-developed assets gradually weakens the US discovery ecosystem by shifting investment, talent, and future revenue streams away from domestic firms and toward a strategic competitor.

Ultimately, resilience requires balancing the short-term benefits of access and efficiency against the long-term imperative of preserving domestic innovation and manufacturing capacity. The goal should not be economic isolation but ensuring that the United States retains the ability to discover, develop, manufacture, and deploy critical biological capabilities without dependence on strategic competitors. As bioconvergence accelerates, maintaining that capacity will become an increasingly important determinant of both national security and economic strength.

Policy Direction

Supporting the long-term strength and innovation of the US biotechnology sector requires a fundamental rethinking of public-private innovation partnerships. Today, the biotech industry is hampered by structural challenges: long and expensive development timeframes, high failure

340 Patent Cooperation Treaty art. 21(2)(a), June 19, 1970; 35 U.S.C. §§ 351–376; 35 U.S.C. § 374; 35 U.S.C. § 122(b); World Intellectual Property Organization, "Patent Cooperation Treaty, Article 21," <https://www.wipo.int/pct/en/texts/articles/a21.html>.

rates, and often unclear regulatory pathways to market. The current model funds discrete products against specific threats, producing assets that are commercially unsustainable once the government contract ends, leaving the underlying platform capacity to atrophy. A more durable approach would orient federal investment toward platform-based technologies capable of producing multiple products against multiple threats over time, with both civilian and military applications that sustain commercial viability between crises.

Buttressing the US biotech ecosystem will require sustained and long-term industrial policy to meet the nation's needs. This ecosystem is critical to American national security and deterrence strategies, and therefore cannot simply be left to market forces. Policy options include strategic stockpiles, production tax credits or price guarantees, allied manufacturing coalitions, and direct government investment in US firms. But no single action will reduce the risk of US reliance on Chinese biological products. Comprehensive action is needed to address the entire biotech ecosystem, similar to what the US Congress did to address US supply chain vulnerability in the CHIPS and Science Act of 2022. The act combined direct subsidies, tax incentives, and government investment in research and development to strengthen domestic semiconductor production. For example, it provided more than \$52 billion in grants and incentives for companies building chip fabrication plants in the United States and offered a 25 percent investment tax credit for semiconductor manufacturing facilities, helping reduce dependence on foreign suppliers and improve the resilience of critical US supply chains. A similar comprehensive effort could address key US biotechnology and pharmaceutical dependencies on China.

The United States should also consider a national security review process for pharmaceutical in-licensing, out-licensing, and data-sharing agreements involving China where such transactions create systemic vulnerabilities in critical sectors of the biotechnology ecosystem. Such a mechanism could be modeled on existing national security review frameworks but tailored to biotechnology, data, and pharmaceutical innovation.

Finally, the United States needs a federal workforce with cutting-edge biological and AI skills. Current federal workforce efforts are hampered by draconian reductions in staff in key science and technology agencies, policies that limit opportunities for scientists and technologists to move with agility between the private sector and government, and a dearth of experts who understand the risks of bioconvergence and are cleared for access to classified information and policy deliberations. Furthermore, massive cuts in federal funding for graduate student education in relevant scientific fields will cripple the talent pipeline of the future if not reversed.

D. Expanding Global Resilience

Meeting the challenge of bioconvergence requires more than domestic resilience—it also requires global engagement. In the short term, the United States must return to playing a significant international leadership role in promoting best practices and reducing risks, including significantly strengthening the capacity to detect threats early and deter and attribute biological attacks. In parallel, governments, companies, and NGOs must develop global modalities for surging and financing medical countermeasures (MCM) and for shoring up supply chains for vaccines, tests, treatments, and PPE. Finally, public and private sector leaders must develop new models of decision-making and resilience-building in the face of growing AI-enabled biological threats.

International resilience requires sustainably financed institutions that can advance national and regional preparedness and resilience, especially among the weakest links in the chain. Threats must be detected and attributed early, before they spread; biosecurity and biosafety innovations must be continually implemented; and medical countermeasures must be manufactured and distributed. International resilience will require the engagement of the US and international private sectors alongside the federal government and other governments, both strategically and operationally.

US Government Investment to Counter Biological Threats

Since 9/11, US national strategies for biodefense have underscored the importance of US global engagements and investments of resources to save lives, prevent economic catastrophe, and protect US and allied interests if a major biological event occurs.³⁴¹ From 2015 to 2025, the United States developed and accelerated the Global Health Security Agenda and bilateral commitments to advance health security, culminating in a commitment to assist at least fifty countries to be better prepared against biological threats, integrating funding and targets across the US Centers for Disease Control and Prevention, US Agency for International Development, Department of State, and Department of Defense.³⁴² This was designed to be a biological early warning network, akin to the missile defense early warning architecture that protects the US homeland.

Decadal US investments in global biosecurity have also included collaboratively eliminating biological weapons capacity across the former Soviet Union, employing scientists with biological weapons and biological weapons-related knowledge to work on civilian applications of their know-how, strengthening laboratory biosafety and biosecurity, building detection capacity for emerging human and animal diseases, supporting outbreak response capacity to stop disease threats at the source, developing and globally delivering medical countermeasures, and related activities. There is no question these investments have provided the United States with early insights into emerging threats and capacity that helps countries and regions respond more effectively to suppress disease spread at its origin rather than in the homeland.³⁴³

The United States has made these investments both bilaterally, through the departments of Defense and State, CDC, and USAID, and through multilateral and global and regional health and security partners, such as the World Health Organization, World Organization for Animal Health, Pandemic Fund, Coalition for Epidemic Preparedness Innovations, and others. The United States has also driven the creation of frameworks to incentivize information and resource collaboration to fight biological threats, including the Biological Weapons Convention (BWC) and its Implementation Support Unit, the G7 Global Partnership Against the Spread of Weapons and

341 The White House, *Biodefense for the 21st Century* (Homeland Security Presidential Directive 10), Apr. 30, 2004, <https://georgewbush-whitehouse.archives.gov/homeland/20040430.html>; The White House, *National Strategy for Countering Biological Threats* (2009), https://jobamawhitehouse.archives.gov/sites/default/files/National_Strategy_for_Countering_BioThreats.pdf; The White House, *National Biodefense Strategy*; The White House, *National Biodefense Strategy and Implementation Plan for Countering Biological Threats, Enhancing Pandemic Preparedness, and Achieving Global Health Security* (2022), <https://bidenwhitehouse.archives.gov/wp-content/uploads/2022/10/National-Biodefense-Strategy-and-Implementation-Plan-Final.pdf>.

342 The White House, "Fact Sheet: The Global Health Security Agenda: Getting Ahead of the Curve on Epidemic Threats," Sept. 26, 2014; The White House, "Fact Sheet: The Global Health Security Agenda," July 28, 2015; The White House, *United States Government Global Health Security Strategy* (2019); US Department of State, "Release of the 2024 United States Global Health Security Strategy," 2024.

343 National Academies of Sciences, Engineering, and Medicine, *Global Health and the Future Role of the United States* (National Academies Press, 2017); Rebecca Katz et al., "Enhancing Public Health Preparedness through the Global Health Security Agenda," *Health Security* 16, suppl. 1 (2018): S-28–S-33.

Materials of Mass Destruction, the Global Health Security Agenda, the Global Health Security Initiative, the Australia Group, UN Security Council Resolution 1540, the Pandemic Fund, and many more.

As a critical element of its sustained global health leadership, the United States has spearheaded and implemented guidance to reduce the risks associated with dual use biological research, including, but not limited to: the historic 1975 International Congress on Recombinant DNA Molecules at Asilomar³⁴⁴; the enactment of the US Select Agent Program in 2002; the Fink Report in 2004³⁴⁵; the 2012, 2014, 2017, and 2024 federal policies on oversight for dual use research of concern and research on pandemic pathogens; the 2010 and substantially upgraded 2024 guidelines for DNA synthesis screening; and the 2025 Executive Order on gain of function research.

Maintaining US Leadership

Since dangerous pathogens know no borders, even if the United States were to fully safeguard its AI models, biological tools, and biotech supply chains, those unilateral actions would not stop the development or deployment of dangerous pathogens overseas. Building global resilience to potential large-scale biological catastrophe requires the US to reinvigorate ties with international partners and to lead in forging a new international architecture that responds to the novel challenges created by bioconvergence. This work could build on some elements of the previous 2019 Global Health Security Strategy of President Trump³⁴⁶ and 2024 Global Health Security Strategy of President Biden,³⁴⁷ as well as the America First Global Health Strategy.

US investments are likely to be most impactful when focused on preventing, detecting, and responding to large-scale threats while also expanding local capacity and incentivizing cofinancing for regional and national biomonitoring and biosecurity capacity. US foreign assistance should provide a sustainable funding stream for biosecurity research, development, and investments in partner countries. Currently, despite strong US support for policies against dangerous gain of function research, US grants and aid have no standard requirements for laboratory biosecurity, biosafety innovations, and managed access to pathogens and biodesign tools.

Finally, the US must recommit to playing a leading role in the international institutions that are shaping the future of biosecurity. These organizations continue to perform important normative functions and provide critical support for countries to become better prepared against biological threats. This includes supporting BWC implementation and new tools for verification. It also means reengaging with the World Health Organization and elevating US leadership in Gavi, the Global Fund, CEPI, and the Pandemic Fund, and focusing attention and resources on biosecurity initiatives that catalyze domestic investments. The United States will lose access to emerging biosecurity threat intelligence and permanently abdicate its leadership position to its competitors if it does not systematically reengage.

344 "The Future of Biosecurity: How AI Is Reshaping Pathogen Surveillance," *Nature* 639 (2025): d41586-025-01009-y, <https://doi.org/10.1038/d41586-025-01009-y>.

345 National Academies of Sciences, Engineering, and Medicine, *Safeguarding the Bioeconomy: Finding the Balance Between Open Science and Security* (2025), <https://www.nationalacademies.org/projects/PGA-DSC-C-18-P-44/publication/10827>.

346 The White House, *United States Government Global Health Security Strategy* (May 2019), <https://trumpwhitehouse.archives.gov/wp-content/uploads/2019/05/GHSS.pdf>.

347 The White House, *US Government Global Health Security Strategy 2024* (April 2024), <https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/04/Global-Health-Security-Strategy-2024-1.pdf>.

Global Public-Private Partnerships

Global biosecurity network. As is the case within the United States, bioresilience requires cooperation among the private, nongovernmental, and public sectors. Improving communication and coordination could best be achieved through the establishment of **a new global biosecurity network**, led by top private and public sector leaders in their independent capacities. The network would optimally include leadership from AI frontier labs (including in the US and China), NGOs, government officials, and philanthropic actors. The network could be led by a small, philanthropically supported secretariat, whose mission would be to convene stakeholders, identify actions, and catalyze their implementation.

This group would be uniquely positioned to focus on practical, near-term private sector actions and incentives for implementing them. These might include:

- **Developing practical safety and security solutions** that can prevent catastrophic misuse of AI-driven bio tools;
- **Placing humans in the loop** for any tools and technologies that can result in the physical creation of biological agents that can cause pandemics or other large-scale biological crises; and
- **Financing bioresilience institutions** and approaches that expand protections against biological risks.

Global biosecurity initiative. This group could also form the basis for a new global initiative for private and public sector leaders to bring forward commitments and actions, which can be taken now to defend against these threats in parallel with key technological developments. The initiative would be modeled on other successful and practical initiatives, such as the Nuclear Security Summit (NSS), Global Health Security Agenda (GHSA), and Global Coalition to Address Synthetic Drug Threats. These initiatives solicited specific national commitments and included a series of convenings to catalyze their announcement and implementation. In the NSS process, these were called “gift baskets”; in the GHSA, they were called “action packages.” In all cases, the key goals for such a process would be to persistently advocate for major commitments, incentives for their implementation, and demonstrable actions. Similarly, in 2024 the United States, alongside the EU, India, Japan, and South Korea, formed the Bio-5 Coalition aimed at building resilience in biopharmaceutical supply chains.³⁴⁸

The network also should prioritize actions for senior US and Chinese public and private sector leaders to take, which would focus specifically on countering AI-enabled biological threats. China and the United States have a shared interest in preventing nonstate actors from using AI to perpetrate biological harm. Existing Track 1 and Track 2³⁴⁹ dialogues have demonstrated that the United States and China can reach shared understandings on some AI safety issues; for example,

348 The White House, “Joint Fact Sheet: The United States and India Continue to Chart an Ambitious Course for the Initiative on Critical and Emerging Technology,” June 17, 2024, archived at <https://bidenwhitehouse.archives.gov/briefing-room/statements-releases/2024/06/17/joint-fact-sheet-the-united-states-and-india-continue-to-chart-an-ambitious-course-for-the-initiative-on-critical-and-emerging-technology/>.

349 For example, the Brookings Institution and the Center for International Security and Strategy at Tsinghua University have convened a US-China AI dialogue since 2019. A separate process, the International Dialogues on AI Safety, convenes senior AI scientists and has produced four joint statements since 2023.

both sides agreed in November 2024 on “the need to maintain human control over the decision to use nuclear weapons.”³⁵⁰

A sustained Track 1.5 dialogue bringing together government officials, scientists, biotechnology firms, AI developers, and nongovernmental experts would provide a practical forum for advancing risk-reduction measures such as **DNA synthesis screening, safeguards for AI-enabled biological design tools, and strengthened laboratory biosafety and biosecurity practices.** Such a dialogue could also develop **confidence-building measures and dedicated crisis communication mechanisms** to address the risk that unusual biological incidents, AI-enabled biological breakthroughs, or suspected malicious activity could be misinterpreted, misattributed, or exploited by third parties seeking to provoke escalation between the two countries. By focusing on concrete technical and operational safeguards rather than broader geopolitical disagreements, a US-China Track 1.5 process could offer a pathway toward reducing risks arising from the intersection of AI and biotechnology.

Finally, building on the US domestic resilience investments, the network should focus significant attention on actions that enable bioresilience through diversified MCM and surge manufacturing and financing. As part of this effort, development banks and development finance institutions (DFIs) should create regional hubs that can surge, scale, and, as needed, stockpile MCMs.³⁵¹ These regional manufacturing hubs should be developed in concert with allies and in partnership with US and local manufacturers.

This initiative would require close collaboration with existing regional efforts focused on geographically diversified manufacturing, including African Union,³⁵² the Accra Reset,³⁵³ and the G20 Joint Finance and Health Task Force.³⁵⁴ It would build on the G7 DFI MCM Surge Financing Initiative,³⁵⁵ the 100 Days Mission, and CEPI 3.0.³⁵⁶ A key goal should be the establishment of regional standby manufacturing capacity that has a market in peacetime and can pivot to develop and distribute PPE, tests, treatments, and vaccines during significant outbreaks. Early focus areas could be ensuring global stores of long shelf-life PPE, supporting geographically distributed manufacturers for multi-pathogen lateral flow assays, and bolstering the new IPPS therapeutics coalition.³⁵⁷

350 The White House, “Readout of President Joe Biden’s Meeting with President Xi Jinping of the People’s Republic of China,” Nov. 16, 2024.

351 This should build on recommendations No. 2 and No. 4 of the 2025 High Level Independent Panel report. National Academy of Medicine, *Closing the Deal: A Strategic Framework for Pandemic Preparedness and Response* (2025).

352 Africa Centres for Disease Control and Prevention, *Africa’s Health Financing in a New Era: Strategic Priorities for April 2025 and Beyond* (2025), <https://africacdc.org/news-item/africas-health-financing-in-a-new-era-april-2025/>; Africa Centres for Disease Control and Prevention, *Presidential Declaration on Advancing Local Manufacturing of Health Products in Africa* (2025), <https://africacdc.org/news-item/presidential-declaration-on-advancing-local-manufacturing-of-health-products-in-africa/>.

353 Accra Reset Secretariat, *The Accra Reset: A Global Initiative for Health Equity and Resilience* (2026), <https://accrareset.org/>.

354 G20 Secretariat, *Joint Finance and Health Task Force (JFHTF): Institutional Arrangements and Operational Framework* (G20 South Africa, 2025), <https://www.g20.org.za/wp-content/uploads/2025/12/JFHTF-Institutional-arrangement.pdf>; G20 Secretariat, *Joint Task Force on Finance and Health: Strengthening the Global Health Architecture* (G20 Brasil, 2024), <https://www.gov.br/g20/en/documents/joint-task-force-on-finance-and-health>.

355 US International Development Finance Corporation, “DFC Announces Surge Financing Initiative and Strategic Investments to Bolster Global Health Security,” press release, 2025, <https://www.dfc.gov/media/press-releases/dfc-announces-surge-financing-initiative-and-strategic-investments-bolster>.

356 International Pandemic Preparedness Secretariat, *The 100 Days Mission: Progress and Strategic Priorities* (2026), <https://ippsecretariat.org/>; Coalition for Epidemic Preparedness Innovations, *CEPI 3.0 Strategy Report: Accelerating Vaccine Development for a Safer World* (2026), <https://static.cepi.net/downloads/2026-02/CEPI%203.0%20Strategy%20Report%20DIG%20ENG%20FINAL.pdf>.

357 International Pandemic Preparedness Secretariat, *The 2025 Diagnostic (Dx) Gap Assessment: Addressing Vulnerabilities in Pandemic Surveillance* (2025), <https://ippsecretariat.org/publication/2025-dx-gap-assessment/>; IPPS, *100 Days Mission*.

International Data Sharing

The United States' withdrawal from the WHO and cessation of official participation in WHO-sponsored committees and technical working groups have disrupted established channels for international monitoring, detection, and data sharing, including US engagement with influenza-surveillance activities associated with the Global Influenza Surveillance and Response System (GISRS).^{358, 359}

To address these gaps, which deprive the United States of vital information on emerging international bio incidents, bio-surveillance MOUs mandating real-time genomic data sharing with priority partners should be immediately negotiated and implemented. These bilateral agreements should include standardized data formats, minimum sequencing coverage requirements, and mutual notification timelines that replicate the core functions of the multilateral frameworks now unavailable to US agencies.

The nation's bioresilience infrastructure must also build independent global pathogen monitoring capacity using a distributed network of sequencers, baseline characterization, both open-source and sensitive intelligence, and commercial capabilities.³⁶⁰ This infrastructure ensures the United States detects emerging threats at least as early as any WHO member state, without dependence on WHO's reporting systems. (For further discussion, see Chapter 4, Recommendation 4.7.)

Findings and Recommendations

Finding 7.1: The United States lacks the functional biological data infrastructure, high-containment defense research and development capacity, and operational response teams needed to maintain deterrence by resilience against AI-enabled biological threats. Existing programs are fragmented across agencies without a unifying institutional structure, standing operational commitment, or the legal authorities necessary for durable public-private engagement.

Recommendation 7.1: Charter a National Institute for Bioresilience (NIBR) as federally authorized public-private national infrastructure for biological AI and data.

- Congress should enact legislation creating and funding the National Bioresilience Institute as a federally chartered, public-private consortium. The statute should define NIBR's mission; specify its governance structure; designate a lead agency and its relationship with DOE, DOD, HHS, NIST, and DHS; and direct the lead agency to specify data types, establish data-sharing processes, and clarify applicable privacy protections and legal safe harbors.
- Congress should authorize and appropriate dedicated compute and data infrastructure for NIBR sufficient to train, evaluate, and serve frontier biological AI models at national

358 U.S. Department of Health and Human Services and U.S. Department of State, "United States Completes WHO Withdrawal," Jan. 22, 2026, <https://www.hhs.gov/press-room/united-states-completes-who-withdrawal.html>.

359 World Health Organization, "Science, Surveillance and Solidarity: Strengthening Global Influenza Preparedness," Apr. 8, 2026, <https://www.who.int/news-room/feature-stories/detail/science--surveillance-and-solidarity--strengthening-global-influenza-preparedness>.

360

scale—on the order of thousands of frontier-class accelerators, petabyte-scale curated biological datasets, secure enclaves for sensitive functional data, and pre-negotiated access to current frontier general-purpose and biological foundation models (Evo, ESM, AlphaFold-class, and successors). Capacity should support the conduct of high-containment research on AI-enabled biothreats and operational reach-back by BEST teams during incident response. NIBR could also use its expertise to develop recommendations and guardrails mitigating the risks of autonomous labs (discussed in Chapter 3).

- Congress should direct NIBR to establish formal partnerships with self-driving laboratories and cloud-laboratory operators for industrialized functional data generation, and joint operating authority with DOE national laboratories as well as with DOD and DHS labs (USAMRIID and NBACC) for high-containment research and dual-use evaluation.
- DOE, in coordination with DOD, HHS, NIST, and DHS, should take initial steps to implement NIBR's base tier leveraging and aligning with the Genesis Mission where appropriate, including creating federated biological data and controlled workspaces on Genesis infrastructure, establishing initial partnerships with private firms and universities through CRADAs and MOUs, and curating AI-ready biological datasets and AI-enabled workflows tied to DOE's biotechnology challenge.
- OSTP should examine different models for incentivizing private entities to participate across NIBR's three tiers and in NSBRR's reserve capability model (see Chapter 6).
- Congress should enact legislation establishing Bio Emergency Services Teams (BEST) as a cross-cutting operational response initiative drawing on NIBR analytical reach-back and NSBRR surge capacity, defining mission triggers and BEST's relationship with states, FEMA, CDC, ASPR, DOD, and national laboratories, and providing for regular preparedness exercises.

Finding 7.2. National resilience requires a robust private biotechnology research and manufacturing sector. The US private biotech sector is a national security asset whose health directly determines US biodefense capacity. Supply chain dependencies in biomanufacturing, APIs, sequencing, and drug discovery create exploitable vulnerabilities, and the federal government has not systematically treated private-sector biotech strength as a national security objective. Major cuts in NIH, NSF, and university funding along with reductions in the S&T workforce in key agencies are having direct impacts on America's capacity to win the bio race. Sustained, multiyear investments in research and human capacity are urgently needed to grow American talent to power the US AI-enabled biotechnology ecosystem, both in government and the private sector. In addition, offshoring of bio manufacturing and reliance on China for the bio supply chain pose long-term threats to the sector's sustainability. Furthermore, the US must have a skilled workforce to ensure the vitality and innovative capacity of the sector.

Recommendation 7.2.1: Rebuild and protect US biotech innovation capacity for the long term.

- Fund relevant research at levels commensurate with the competitive challenge.

- Extend financial incentives to on-shore pharmaceutical discovery and production through NIH grant funding, ARPA-H project funding, enhanced R&D tax credits, favorable capital gains treatment for biotech investments, and BARDA where appropriate.
- Expand incentives for domestic biotechnology innovation through NIH, ARPA-H, BARDA, enhanced R&D tax credits, and long-term support for translational biotechnology infrastructure and clinical trial capacity.

Recommendation 7.2.2: Work to reduce current pharmaceutical supply chain dependencies on China.

- Congress should direct USTR, HHS, DOD, and the Department of Commerce to conduct a comprehensive assessment of US dependency on foreign sources for critical pharmaceutical ingredients, precursor chemicals, biologics manufacturing inputs (including antibiotics and antivirals) essential to national defense and pandemic response. This should include assessment of how the capabilities of US allies and partners could contribute to advancing the objective of reducing dependencies on China.
- Congress should establish a national security review process for pharmaceutical in-licensing and out-licensing agreements involving China, as well as data-sharing arrangements that could create systemic vulnerabilities in critical sectors of the US biotechnology ecosystem. This review could be modeled on the Committee on Foreign Investment in the United States (CFIUS).

Recommendation 7.2.3: Strengthen US industrial capacity and supply chain security.

- Build support for a Boost American Bio Manufacturing Act modeled on the CHIPS and Science Act to reduce growing dependencies on China via a comprehensive industrial policy supporting the US biotech and pharmaceutical industries. This could include targeted subsidies, public R&D investment, tax incentives, workforce development, place-based economic strategies, and collaborative production with aligned foreign partners.
- Establish federal incentives for strengthening US biotech and pharmaceutical domestic manufacturing capacity, reducing single-point-of-failure risks (such as dependency on Chinese contract bio manufacturers), and improving long-term supply chain resilience. This could include investment tax credits, guaranteed procurement floors, accelerated depreciation, and BARDA- and ARPA-H-supported manufacturing partnerships for strategically essential products.
- Designate domestic genomic sequencing platform capacity as a critical supply chain priority alongside API, KSM, and CDMO manufacturing. Direct the Department of Commerce and BARDA to assess current US market share in high-throughput sequencing platforms, identify domestic capacity gaps, and develop incentive programs for domestic sequencing platform development and manufacturing analogous to those proposed for pharmaceutical APIs.
- Extend the procurement restrictions established under the NDAA for named Chinese biotechnology entities to additional categories of Chinese-origin sequencing and biological AI infrastructure.

Recommendation 7.2.4: Ensure access to a skilled workforce necessary to support the biotechnology sector.

- Create targeted federally funded fellowships in AI and bioscience to create a pipeline of AI-bio scientists with government service experience, appropriate clearances, and institutional knowledge.
- Enable rotational exchanges between government laboratories and leading AI firms.
- Provide special compensation authority and competitive retention bonuses to reduce the federal-private wage differential for AI-biology researchers to make long-term government careers feasible for top-tier AI-biology experts.

Finding 7.3. US security and resilience are directly related to global biosecurity and resilience. Recent reductions in US support for global health institutions have reduced our insight into and ability to respond to global health threats. Strengthening international capacity and resilience through public-private partnerships is a sustainable way to address these challenges.

Recommendation 7.3.1: Expand US engagement and investment in global biosecurity. Redeploy US public health experts to countries globally.

- Expand foreign assistance for biosurveillance and attribution. Include biosecurity provisions in foreign assistance grants and contracts.
- Reengage with WHO, Gavi, Global Fund, CEPI, and the Pandemic Fund.

Recommendation 7.3.2: Establish new pathways for global public-private partnerships to support bioresilience. Create a global senior leaders network initiative with a philanthropically supported secretariat. The initiative should focus on near-term actions and incentives to create sustainable biosecurity and resilience globally. Build out a sustainable process that continually solicits national actions improving global prevention and resilience.

- Build MCM regional manufacturing hubs through development banks and G20 DFIs (building on the G7 DFI MCM Surge Financing Initiative, 100 Days Mission, and CEPI 3.0).

Recommendation 7.3.3: Establish a US-China Track 1.5 Dialogue on AI-Enabled Biological Risks.

- Convene a sustained Track 1.5 dialogue involving government officials, AI and biotechnology leaders, scientists, and nongovernmental experts from both countries focused on reducing catastrophic risks arising from the convergence of AI and biotechnology.
- Prioritize practical risk-reduction measures, including DNA synthesis screening, laboratory biosafety and biosecurity practices, and approaches to preventing the design or creation of potentially pandemic-capable biological agents without appropriate human oversight.
- Develop confidence-building measures and crisis communication mechanisms to reduce the risks of misperception, misattribution, or escalation following unusual biological incidents or advances in high-risk biological capabilities.

Finding 7.4: Following US withdrawal from WHO, the United States lacks adequate bilateral biosurveillance architecture to detect emerging biological threats as early as multilateral surveillance systems previously enabled it to do. This gap is not merely diplomatic; it creates potential for both intelligence loss and warning failure with direct national and economic security consequences.

Recommendation 7.4: Build independent biosurveillance capacity and bilateral biosecurity engagement infrastructure to replace multilateral coverage lost through WHO withdrawal.

- The State Department and HHS should negotiate and implement biomonitoring MOUs with priority partners mandating real-time genomic data sharing, standardized data formats, minimum sequencing coverage, and mutual notification timelines.
- NIBR should build and operate independent global pathogen monitoring capacity using federated sequencing nodes and baseline characterization to ensure the United States detects emerging threats without dependence on WHO reporting systems.
- The US government should convene an annual Biosecurity and Bioresilience Summit with senior officials, scientific leads, and private sector biosecurity actors from priority partner nations to institutionalize bilateral coordination and provide a forum for sensitive threat intelligence sharing.

CONCLUSION

A Call to Action

The convergence of AI, biotechnology, and bioengineering is not a challenge for the future. It is already here, and it requires immediate attention and decisive action. The United States has the opportunity to lead if we seize the initiative now by shaping the infrastructure, incentives, laws, and norms through which biological capabilities develop, and by ensuring that innovation reinforces security and that security reinforces innovation. Doing so will require moving beyond fragmented, piecemeal responses toward a coherent national strategy that recognizes resilience as a source of deterrence, public-private coordination as a strategic asset, and sustained investment as a prerequisite for the safety and security of the American people and the world.

The United States has tackled similarly daunting challenges before through deliberate and comprehensive action. In the early 1990s, Senators Richard Lugar and Sam Nunn worked across party lines to enact landmark legislation that enabled the United States to manage the potentially existential security threat posed by the prospect of proliferating nuclear weapons and materials following the collapse of the Soviet Union. Congress authorized and appropriated funds for a multidimensional program of consequential actions—some unilateral and some cooperative—at home and around the world.

Now, in a new era of high risk, a similarly bold nonpartisan effort is required to build security and reduce dangers to the American people and humanity. The path to doing so is a different one because it must harness synergies between the private sector and public interests. Achieving this vision will require technically informed coordinated actions by multiple entities and individuals with the capacity to shape the United States' bio-innovation and biosecurity ecosystem:

Congress: Congress is essential to providing the legal and financial foundation for bio resilience by deterrence. To date, enacted and proposed legislation has targeted key elements of the bioconvergence threat, but a holistic approach is needed to fully address each element of deterrence. Short-term initiatives cannot substitute for the sustained commitment required to build resilience across decades. This report has discussed key areas where Congress can transform the bioconvergence landscape and establish the United States as a leader in this domain: legislating evaluation and access regimes for AI-enabled biological models and tools; mandating KYC for DNA synthesis and reforming screening to a function-based instead of a list-based model; creating a national biological monitoring system; expanding the government's capability to achieve rapid biological agent attribution; establishing a National Strategic Bio Readiness Reserve; reforming the national clinical trials system; creating a National Institute for Bioresilience; and launching initiatives to support the American biotech and pharmaceutical industries modeled after the CHIPS and Science Act. Each step must also be paired with sustained appropriations and innovative spending and contracting authorities, allowing federal agencies to adapt at the speed of the bioconvergence threat.

Executive Branch: Addressing the risks of bioconvergence cannot happen without federal leadership and coordination. No single department or agency can manage this threat alone; instead, success will require White House leadership and synergistic efforts involving the departments of Commerce, Health and Human Services, Defense, Energy, Homeland Security, and the intelligence community, working collaboratively with state and local partners. The executive branch must use existing authorities wherever possible to improve AI and biotechnology governance, strengthen biological monitoring, and modernize procurement, while continuing to develop new capabilities for prevention, detection, attribution, and response. Equally important, the executive branch must transform how it engages the private sector, including creating processes to share threat information; building and sustaining AI and biotechnology expertise across the US government that can work effectively with the private sector on prevention, risk reduction, and resilience building; and establishing reliable pathways to leverage private sector capabilities and reserve capacity to respond effectively during crises.

Private Sector: The private sector will determine, to a significant extent, whether the age of bioconvergence strengthens American security or accelerates potentially catastrophic biological risks. Frontier AI developers, biotechnology companies, DNA synthesis firms, pharmaceutical manufacturers, and research organizations possess much of the expertise and infrastructure on which US resilience depends. These entities should continue to drive toward the frontiers of scientific advancement while incorporating security and resilience into the design, deployment, and governance of their technologies. Simultaneously, the risks created by bioconvergence require that these firms collaborate in novel ways with the US government, including: sharing information on technology and risks; agreeing to continual testing, red-teaming, and the establishment of adaptive guardrails; participating in mutually beneficial schemes for supplying reserve capacity and emergency response capabilities; and joining forces to manage supply chain risks.

NGOs and Philanthropies: Nongovernmental institutions and philanthropies can perform vital functions and can provide additional capacity to meet unprecedented challenges. Universities, think tanks, and independent research institutions can offer technical expertise, analysis, and laboratory capacity for testing and assessing risks and developing innovative means of reducing or countering them. Global partners can help build and participate in hybrid models of international collaboration that seat governments alongside the private sector. Funders who are willing to support what governments choose not to or cannot support can play key roles in advancing resilience-building initiatives, both nationally and internationally. For example, the proposed Senior Leaders Biosecurity Network described in this report could serve as the “secretariat” that coordinates biotechnology security and resilience work among governments and industry globally.

The technological capabilities and associated risks discussed in this report will continue to rapidly evolve, demanding immediate and comprehensive action. Looking to the future that is coming fast, the recommendations throughout this report are designed to remain relevant and adaptive to innovation at the intersection of advances in AI, biotechnology, and bioengineering.

As this report clarifies, an effective response to bioconvergence should support rather than inhibit American innovation and advance US private sector strengths in AI, biotech, and bioengineering. That advancement must include recognition that biosecurity is fundamental to success in this

era of both promise and peril. A balanced public policy approach to the AI-enabled bio ecosystem is one that synergizes government requirements and incentives for safety and security with the strategic commitment to promote US science and technology leadership. Achieving and sustaining the desired outcome will require deftly managing bioconvergence on a technically informed basis with agility and at scale. This comprehensive framework and set of actionable policy recommendations provide policymakers, industry leaders, civic and global partners, and philanthropies with a roadmap for responding with the speed and seriousness that this pivotal moment in human history demands.

Consolidated Findings and Recommendations

Chapter 3

Finding 3.1.1: The United States lacks comprehensive federal oversight for AI systems with biological design capabilities. State-level regulations create compliance fragmentation. Current federal authorities do not adequately address AI-enabled biological threats.

Recommendation 3.1.1.1: Adopt a federal, capability-triggered AI-bio regime. Impose binding obligations on frontier AI developers whose models cross defined capability and compute thresholds relevant to biological design, synthesis, or dual-use research. Require pre-deployment biosecurity evaluations and post-deployment monitoring. Require model development companies to test qualifying models before release for biological and chemical misuse potential, document results, retain records, and retest after major updates.

Recommendation 3.1.1.2: Mandate public safety frameworks. Require AI companies to publish a frontier safety protocol describing dangerous-capability thresholds, mitigation measures, intended-use restrictions, and cybersecurity controls. Require companies to report critical safety incidents on short timelines.

Recommendation 3.1.1.3: Tie federal money to compliance with safety requirements. Make eligibility for federal contracts, grants, cloud credits, and research partnerships contingent on meeting baseline AI-bio safety requirements.

Finding 3.1.2: The US government lacks a central entity empowered to regulate AI safety. NIST is empowered to issue AI safety standards but only through voluntary standards, and does not have authority to promulgate binding regulations or enforce them. Within NIST, the Center for AI Standards and Innovation (CAISI) has the capacity to conduct national security assessments of advanced AI models but only on a voluntary basis.

Recommendation 3.1.2.1: Congress should authorize a central federal government entity empowered to issue regulations on safety testing for advanced AI models. Such regulatory authority must include the power to:

- Define covered models;
- Define risk thresholds;
- Evaluate frontier models;

- Conduct red-teaming of AI models;
- Review safety cases and incidents;
- Review risks to third parties; and
- Review the internal governance processes of AI labs.

Recommendation 3.1.2.2: Congress should authorize the central federal government entity to promulgate binding regulations specific to biological threats and AI models. This should include the regulation of AI models falling below defined compute thresholds for general purpose AI models but exceeding capabilities specific to biology or having been trained on biological data. In developing the requirements, the entity should consult with biology and public health experts at HHS, DOE national laboratories, and DOD. These bio evaluations should assess whether a model can:

- Generate actionable protocols for pathogen synthesis or modification;
- Assist nonexperts in overcoming tacit biological knowledge barriers;
- Enable optimization of delivery mechanisms, virulence, or immune evasion;
- Circumvent existing biosafety and biosecurity controls.

Evaluations should be independently reproducible and subject to audit by the designated federal entity. This entity would also work closely with the National Institute of Bioresilience, discussed in Chapter 7.

Recommendation 3.1.2.3: In the absence of a newly created entity, Congress should authorize CAISI to undertake the actions outlined in Recommendations 3.1.2.1 and 3.1.2.2.

Recommendation 3.1.2.4: Congress should create a mandatory AI-bio incident reporting regime. Covered AI developers, cloud providers, and relevant research institutions should be required to report material biological safety incidents, including:

- Demonstrated dangerous model capabilities;
- Unauthorized access to sensitive models;
- Evidence of malicious biological misuse;
- Failures of safeguards or containment mechanisms;
- Confidential risk information sharing.

The designated federal entity should maintain a secure incident database and produce anonymized lessons-learned reports to Congress.

Recommendation 3.1.2.5: Congress should create enforcement power along with compliance incentives. The designated federal entity should be authorized to levy civil fines for failure to comply and to make referrals to other federal regulators, such as the FTC.

Finding 3.1.3: Any AI-specific entity in the US government must be able to compete with tech companies to attract top talent and coordinate quickly with relevant government experts.

Recommendation 3.1.3.1: Congress should authorize the designated federal entity for AI safety to have excepted-service hiring authority for faster hiring and individualized compensation. This could be modeled after DARPA hiring authorities, DOE national laboratory appointments, NSF rotator programs, and US Digital Service hiring authorities.

Recommendation 3.1.3.2: Congress should establish an AI Safety Fusion Cell within CAISI. This would enable rapid information-sharing and colocation of national security and AI experts from Commerce, DOE, DOD, HHS, and the intelligence community.

Finding 3.1.4: AI-enabled biotechnology safety can be encouraged through federal incentives. Currently, contractors and grantees have no biosecurity obligations tied to receipt of federal resources. Conditioning federal funding on biosecurity compliance converts an externality into a cost firms must manage, creating market-wide incentives without waiting for legislation. Under current incentive structures, safety investment is a competitive disadvantage rather than a commercial advantage. Firms that invest in biosecurity safeguards bear costs avoided by competitors who do not, creating a race to the bottom. Positive incentives can reverse this dynamic without mandates.

Recommendation 3.1.4.1: Make eligibility for federal contracts, grants, cloud credits, and research partnerships contingent on meeting baseline AI-bio safety requirements.

Recommendation 3.1.4.2: Create positive incentives for safety investment. Examples could include: tax credits, accelerated procurement preference, liability safe harbors, and priority access to federal testbeds. These incentives would be available for firms that demonstrate they exceed baseline requirements through independent audits, robust red-teaming, and integration with synthesis screening safeguards.

Finding 3.1.5: Federal preemption of state AI laws could eliminate existing safety requirements without replacing them. State laws currently provide the only mandatory AI safety obligations in the United States. Preemption without replacement would widen rather than close governance gaps.

Recommendation 3.1.5.1: Congress should not preempt state AI safety laws unless it simultaneously installs a federal standard at least as protective on testing, transparency, and incident reporting related to biological risks as laws passed in California and New York.

Finding 3.1.6: International regulatory frameworks are increasingly shaping global AI-bio governance, but the United States lacks a coordinated strategy to respond to extraterritorial regulation and emerging international biosecurity standards. The European Union has established the most comprehensive international framework for regulating AI systems through the EU AI Act and the proposed EU Biotech Act. Because these regimes apply extraterritorially to companies operating within European markets, US-based AI firms developing biological AI capabilities may effectively become subject to foreign regulatory standards even in the absence of comparable domestic oversight.

Recommendation 3.1.6.1: Harmonize with international standards, where feasible. US rules should be interoperable with the EU-style approach to risk management, documentation, testing, and incident reporting, so firms are not forced to build separate compliance systems

for different jurisdictions. The executive branch should create an interagency coordination body responsible for tracking international AI and biotechnology regulation, assessing implications for US competitiveness and national security, and coordinating engagement with allies on AI-bio governance issues. The mechanism should include participation from agencies responsible for commerce, national security, public health, biotechnology, and AI policy.

Recommendation 3.1.6.2: Promote international harmonization of AI-biosecurity safeguards through multilateral engagement. The United States should work through forums such as the OECD, G7, and other allied institutions to promote common standards for evaluating biological misuse risks in advanced AI systems. These efforts should prioritize shared approaches to model evaluations, incident reporting, biosecurity safeguards, and oversight of high-risk biological AI applications while reducing regulatory fragmentation across jurisdictions.

Finding 3.2.1: There is no government-wide framework to evaluate dual-use risks associated with biological training data. The capabilities of biological AI models depend directly on training data. Current data governance lacks risk assessment for dual-use applications. Commodification and restricted access trends require coordinated federal response.

Recommendation 3.2.1.1: NIH should expand its tiered access model for the Database of Genotypes and Phenotypes to cover a broader class of biological datasets involved in AI training.

Recommendation 3.2.1.2: Congress should require developers of frontier AI models to disclose categories of biological datasets used in training.

Recommendation 3.2.1.3: Congress should establish a central office within the US government responsible for creating harmonized standards for biological datasets, and storing and managing access to biological data. This should include coverage of synthetic biological datasets used in AI model training.

Recommendation 3.2.1.4: Establish standardized national risk-tiering and access-control protocols for biological datasets. The federal government should develop standardized criteria for categorizing biological datasets according to dual-use risk potential. NIH, NIST, and relevant national security agencies should jointly establish interoperable access-control standards across federally supported biological databases. Federal agencies should require standardized auditing, credentialing, and monitoring procedures for access to high-risk biological datasets.

Recommendation 3.2.1.5: Create trusted-access pathways for trusted institutions that preserve beneficial biotechnology research. Federal agencies should establish expedited access mechanisms for accredited academic, medical, and public-interest researchers working with higher-risk biological datasets. NIH should develop standardized trusted-researcher credentialing systems for controlled-access biological data environments.

Finding 3.2.2: Biological datasets have become strategic national security assets, and access to high-risk biological datasets by adversary nations is a risk to US national security.

Recommendation 3.2.2: Congress should authorize federal review mechanisms for foreign access to high-risk biological datasets with national security implications. The Department of Commerce should evaluate whether certain categories of biological training data warrant export-control restrictions.

Finding 3.3.1: Current governance of biological AI tools lacks a standardized, government-supported framework for managing access and monitoring risks across different tool types. While tiered access models and API-based controls exist, they are inconsistently applied, vulnerable to circumvention, and under-resourced. Developers lack the infrastructure needed to implement robust access controls, particularly for high-risk biological design tools and datasets.

Recommendation 3.3.1.1: The federal government should establish a centralized infrastructure (e.g., a secure API hosting and monitoring platform) to support biological AI developers in implementing standardized tiered access controls, including real-time monitoring, user verification, and risk-based usage restrictions.

Recommendation 3.3.1.2: Congress should mandate the development of unified binding national standards for tiered access and risk classification of biological AI tools, similar to biosafety levels, incorporating best practices from frameworks proposed by NTI and *Science*.

Recommendation 3.3.1.3: Biological AI systems operating under managed access regimes should maintain logging and monitoring systems capable of identifying anomalous or high-risk usage patterns. Developers should establish escalation procedures for suspicious activity, including human review, additional identity verification, temporary access suspension, and referral pathways to appropriate authorities where necessary.

Finding 3.3.2: Safeguards applied to open-source biological AI models are easily circumvented once model weights are publicly released. The Evo1 example demonstrates that withholding dangerous training data is insufficient if downstream users can fine-tune openly available weights.

Recommendation 3.3.2: Establish federal guidance governing release of open-weight biological AI models. Federal agencies should develop mandatory risk-based standards governing the public release of biological AI model weights, training code, and fine-tuning capabilities. Developers of frontier biological design tools should conduct documented dual-use risk assessments before releasing model weights or enabling unrestricted fine-tuning.

Finding 3.4.1: List-based DNA synthesis screening cannot detect novel, modified, or emerging pathogens. List-based screening fails when sequences are modified enough to avoid matching a listed agent. AI-enabled pathogen design will compound this vulnerability by generating novel sequences specifically designed to evade list-based detection while maintaining their capability to do harm.

Recommendation 3.4.1.1: Mandate function-based, AI-enabled DNA synthesis screening for all domestic synthesis providers regardless of funding source, replacing the current voluntary and list-based framework. Impose import restrictions on foreign DNA synthesis providers that do not conform to US or equivalent screening and customer verification standards. Consider a licensing scheme for DNA synthesis organizations and/or customers.

Recommendation 3.4.1.2: Transition the FSAP system from list-based to function-based, allowing for broader and faster adaptation to novel pathogens of pandemic potential.

Recommendation 3.4.1.3: Consider arrangements to compensate companies for screening costs through tax refunds, subsidies, or other means, as the cost of sequencing falls and screening takes up larger proportions of budgets.

Finding 3.4.2: Mandatory customer-verification requirements do not apply comprehensively to domestic DNA synthesis orders, particularly orders placed outside federally funded research and those produced using benchtop synthesizers. Although federal funding conditions require covered researchers to use providers that follow federal screening standards, much of the broader domestic market continues to rely on voluntary industry practices. While know-your-customer requirements are standard in finance and export controls, they are absent from synthesis. Without knowing who is ordering what, synthesis providers cannot identify suspicious acquisition patterns.

Recommendation 3.4.2: Require customer verification for all DNA synthesis orders exceeding defined thresholds, imposing mandatory know-your-customer protocols. Extend these requirements to benchtop DNA synthesizers as next-generation devices approach commercially relevant synthesis thresholds.

Finding 3.4.3: Cloud labs and CROs are not subject to BSL requirements unless they receive federal funding. Cloud labs lower the tacit knowledge barrier that has historically limited biological weapons development by providing automated execution of experimental protocols. Foreign cloud labs serving US customers or operating with US-origin materials face additional regulatory gaps under current export control frameworks.

Recommendation 3.4.3: Extend BSL regulatory requirements to all commercial cloud labs and CROs operating domestically. Develop international standards for cloud lab biosecurity to address the gap for foreign providers.

Finding 3.4.4: AI-enabled research acceleration and cloud lab proliferation expand access to dangerous biological work. Current federal oversight gaps leave privately funded dangerous research unregulated.

Recommendation 3.4.4: Congress should establish a unified, permanent statutory framework governing high-risk biological research that applies to both federally funded and private-sector entities. Congress should also direct OSTP to finalize clear oversight standards for non-federally funded research, including research conducted through international collaborations or in foreign jurisdictions with inadequate biosafety controls.

Chapter 4

Finding 4.1: The United States has no unified, authoritative strategy for national biological monitoring requirements. Coverage thresholds, detection lag standards, data interoperability obligations, and performance metrics are either undefined or fragmented across agencies. Without a common requirements framework, neither procurement nor accountability is possible.

Recommendation 4.1.1: Establish a national biological monitoring requirements framework. The White House and Congress should both take actions to establish a centralized and harmonized biological monitoring system within the US government. The White House can use executive actions based on existing authorities to establish a national biological monitoring requirements framework. The framework should be White House-led and designate the **DOD and HHS as joint lead federal agencies** responsible for maintaining the requirements framework and overseeing interagency coordination.

At a minimum, the framework must define:

- population coverage thresholds for environmental monitoring;
- maximum acceptable detection lag relative to clinical case emergence;
- data interoperability standards for federal, state, and private data streams; and,
- performance reporting requirements for all federally funded biomonitoring activities.

The framework must be **technology-neutral**, specifying outcomes rather than methods, to preserve room for private sector innovation and avoid being locked into approaches that will quickly become obsolete. Congress should further codify this structure and responsibilities in statute as well as provide any authorities, including for procurement, that the agencies identify as missing from their ability to fully launch a comprehensive biological monitoring system.

Recommendation 4.1.2: Incentivize nationwide investment in wastewater monitoring through competitive, performance-based contracts. These incentives and investments should include state and local public health agencies and hospital systems so that wastewater surveillance outputs will trigger specific actions from the public health and healthcare sector, such as increased testing for specific disease threats, deployment of new or existing medical countermeasures, and/or intelligence or law enforcement actions to find and stop malign actors. Require pathogen-agnostic metagenomic sequencing at all US sites within five years.

Finding 4.2: Public health agencies are trained to detect and respond to disease, not to assess threat intent, evaluate adversary capabilities, or integrate classified intelligence into operational decisions. As a result, biological signals that might warrant national security attention are routinely processed through an epidemiological lens alone, while intelligence assessments of biological threats rarely incorporate real-time public health monitoring data. This gap is not a failure of competence; it is a structural mismatch. Bridging requires a dedicated interface function, a standing team with the training clearances, and institutional authority to operate fluently in both domains.

Recommendation 4.2: Ensure HHS has a dedicated team of biosecurity experts capable of coordinating with the NSC, DOD, DOE, DHS, and the IC. Congress should require that HHS establish a new unit dedicated to the biosecurity staffed by security-clearance holding experts and appropriate dedicated no-year funding directly to that unit. Congress can require HHS to report on how it will implement this requirement, but at a minimum, the team must be capable of the following functions:

- **Receiving and acting on classified biological threat intelligence** from DOD, DOE, DHS, and the intelligence community (IC), and translating that intelligence into public health preparedness and response actions within HHS.
- **Providing public health expertise in national security processes**, including participation in interagency threat assessment, the National Security Council biological threat working group, and any classified biological early warning review.
- **Serving as the authoritative interface between civilian biological monitoring data and the national security community**—ensuring that anomalous signals

detected through environmental or clinical monitoring are assessed in light of available intelligence, and vice versa.

- **Supporting attribution analysis** in coordination with the IC, DOD, and DOE when an anomalous biological event is detected, including assessing whether a signal is consistent with natural emergence, accidental release, or deliberate use.
- **Maintaining standing relationships and protocols** with DOD's biological defense programs, relevant DOE laboratories, DHS's CWMD office, the FBI's WMD Directorate, and relevant IC components, so that coordination in a crisis does not depend on building relationships from scratch.

Finding 4.3: The existing National Wastewater Surveillance System (NWSS), built during COVID-19, covers approximately 42 percent of the US population and relies primarily on targeted assays for known pathogens.³⁶¹ It has demonstrated the capacity to provide four to seven days of advance warning ahead of clinical surges, but its current architecture cannot detect novel or engineered agents. NWSS was established entirely through emergency supplemental funding and has never received a permanent, dedicated funding base. No stable, multiyear contracting framework exists to sustain or expand this infrastructure.

Recommendation 4.3: Establish a national environmental biological monitoring contracting framework.

- **Funding:** Congress must appropriate **stable (multiyear) and dedicated funding** for environmental biological monitoring services procured through competitive contracts with private sector firms.
- **Local governments:** States and municipalities must be permitted and incentivized to procure local environmental biological monitoring services from qualified private vendors, with federal matching funds available to jurisdictions that meet minimum coverage and data-sharing standards.
- **Data rights:** All contracts must include data rights provisions maintaining federal access to biological monitoring data and clear rules on downstream commercial use.
- **Advanced sequencing:** Contracts must require transition to pathogen-agnostic metagenomic sequencing at all sites, with targeted PCR assays retained for confirmed priority pathogens; technical requirements must be updated on a rolling basis as the state of the art advances.

Finding 4.4: No operational US system currently integrates environmental, clinical, genomic, and behavioral biological monitoring data streams in real time. AI-enabled analytics are the only practical means of processing the data volumes a continuous, pathogen-agnostic national biological monitoring system would generate, and the federal government needs sustained and dedicated procurement authorities to acquire these capabilities from the private sector.

361 U.S. Centers for Disease Control and Prevention, "About CDC's Wastewater Monitoring Program," updated July 31, 2026, <https://www.cdc.gov/wastewater/about/index.html>

Recommendation 4.4: Procure AI-enabled biological monitoring analytics as a competitive contracted service.

- **Analytics as a service.** The federal government should procure AI-enabled biological monitoring analytics as a service. The analytics should include baseline modeling, anomaly detection, multi-stream data integration, and threat characterization through competitive contracts with private sector firms, using the model established for intelligence community analytic services and DOD data platforms.
- **Human review.** All contracts must require human review of anomaly flags before consequential decisions are taken.
- **Prequalification.** The federal government should establish a prequalification process for biological monitoring analytic vendors, evaluating technical capability, data security, and interoperability before firms are eligible for federal contracts. This will create a qualified vendor pool available to cities, states, and federal agencies without duplicating procurement overhead.
- **Fund early-stage innovation.** Early-stage innovation in detection, sequencing, and integration must be supported through the Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) programs.

Finding 4.5: Real-time genomic surveillance capacity is concentrated in a small number of academic and public health laboratory centers, creating geographic gaps and single-point-failure fragility identical to the diagnostic testing failure exposed by COVID-19. A distributed private laboratory network with pre-negotiated surge contracting authority could provide both continuous coverage and rapid expansion capacity.

Recommendation 4.5.1: Expand genomic surveillance through distributed private laboratory networks.

- **Long-term contracts.** The federal government must establish long-term contracts with private genomic sequencing companies and laboratory networks to provide continuous, geographically distributed genomic surveillance coverage, with the goal of next-day sequencing turnaround for priority pathogen samples at sites covering 80 percent of the US population within five years.
- **Surge capacity.** The federal government must establish a standing emergency contracting mechanism with pre-negotiated surge contracts in place before the next outbreak that authorizes rapid expansion of genomic sequencing capacity upon declaration of a public health emergency. Congress can spur this action by mandating it specifically and requiring HHS periodically report on its development and status.

Recommendation 4.5.2: Provide HHS with other transaction authority to test, evaluate, and scale new capabilities quickly. DOD has shown that other transaction authority (OTA)—a contracting mechanism that allows the government to rapidly prototype, research, and acquire

innovative technologies outside of normal, lengthy federal procurement timelines³⁶²—can accelerate collaboration with nontraditional contractors and speed capability delivery for cutting edge technology. Congress should provide HHS with a similar authority to rapidly generate and scale the technology needed by the US government and private sector for real-time biological data collection, harmonization, and monitoring at scale.

Finding 4.6: US biological monitoring investment follows a boom-and-bust pattern driven by emergency supplemental appropriations that arrive too late to build the infrastructure needed to detect the emergency that triggered them. Annual discretionary appropriations are structurally incompatible with the continuous operations a national early-warning system requires. The private-sector-led model proposed in the preceding recommendations is only as durable as its funding mechanism; without stable, multiyear contracting authority, private firms cannot make the capital investments, hire the personnel, or build the systems that early warning demands. For other areas of critical national infrastructure, the US government has used trust funds—such as the Highway Trust Fund and the Airport and Airway Trust Fund—to ensure long-term, stable funding.

Recommendation 4.6: Establish a Biological Early Warning Trust Fund with mandatory multiyear contracting authority.

- **Early warning trust fund.** Congress should establish a Biological Early Warning Trust Fund modeled on other appropriations-funded trust funds, such as the Medicare Supplementary Medical Insurance Trust Fund. The Trust Fund should have dedicated, mandatory appropriations insulated from annual discretionary competition, with funding triggers and matching grants for state and local governments to support surveillance, data sharing, and surge capacity.
- **Emergency resources.** The Trust Fund must include automatic trigger provisions releasing additional emergency resources when prespecified epidemiologic thresholds are crossed, enabling rapid contract expansion without a new appropriation.
- **Long-term contracts.** Contracting agencies must be authorized to issue multiyear contracts of five to ten years for core biological monitoring services, providing the contractual stability necessary for private sector investment in infrastructure and personnel.
- **Matching grants.** Cities and states that meet federal coverage and data-sharing standards should be eligible for Trust Fund matching grants, creating a federated funding architecture that distributes both investment and resilience across levels of government.

Finding 4.7: The biological monitoring architecture proposed in this chapter will not succeed without a binding privacy and data-governance framework. Public trust is an operational requirement for any functioning early-warning system. Although community-level wastewater monitoring generally poses low individual-identification risk, the transition to pathogen-agnostic metagenomic sequencing, smaller-scale sampling, expanded data sharing, and potential national-security uses

362 10 U.S.C. §§ 4021-4022.

require enforceable standards to prevent misuse, function creep, stigmatization, and erosion of public trust.

Recommendation 4.7: Establish a binding privacy and data-governance framework for biological monitoring.

- **Technical standards.** Require federally funded metagenomic biomonitoring to use validated privacy-preserving protocols, including host-DNA depletion where feasible and computational removal of human reads, residual human sequence quality control thresholds, independent auditability, and strict limits on retention or sharing of raw reads that may contain human genomic material.
- **Tiered protections by sampling scale.** Require heightened public health justification, data minimization, access controls, and community engagement when monitoring is conducted at smaller scales. Distinguish large community-level sewershed monitoring from small-site, facility-specific, tribal, correctional, military, campus, vessel, port-of-entry, or other targeted monitoring.
- **Use restrictions.** Biological monitoring data may be used only for purposes directly related to public health, biosecurity, and clearly defined biomedical research and medical countermeasure development. Use for purposes unrelated to these missions, including law enforcement, employment or insurance decision-making, non-public-health profiling, or other punitive uses, should be prohibited absent explicit statutory authorization. Secondary research uses should require independent review, data minimization, and public reporting sufficient to preserve public trust.
- **Data access and stewardship.** Distinguish raw metagenomic reads, processed pathogen calls, aggregate trend data, and public dashboards. Establish data-use agreements, access logs, retention limits, data-minimization rules, and review requirements for secondary research uses or requests for more granular data or samples.
- **Oversight.** Establish an independent advisory body with community, public health, ethics, statistical, tribal, and data-governance expertise to review data-use practices, proposed secondary uses, requests for granular data access, and any expansion of system purposes and report publicly, supported by local community input mechanisms at the point of implementation.
- **Tribal data sovereignty.** Require wastewater monitoring involving tribal communities to be governed by tribal approval, tribal data-stewardship expectations, data ownership and attribution rules, and participatory engagement consistent with tribal sovereignty.
- **National security interface.** Establish clear, publicly disclosed criteria governing the conditions under which biological monitoring data may transition from open public health infrastructure into classified national security channels, with classification decisions subject to independent oversight.

Finding 4.8: International biological monitoring depends on national reporting under the World Health Organization's International Health Regulations (IHR). The IHR obligates member states to report a public health emergency of international concern but provides no enforcement mechanism. China's delayed reporting on COVID-19 was the predictable result of a system in

which economic and political incentives favor delayed disclosure. A detection architecture that depends on voluntary good-faith reporting by other nations is structurally vulnerable. **Building a meaningful international biological monitoring capacity requires independent detection capability that does not rely on host-country disclosure, as well as sustained investment in global data-sharing infrastructure that incentivizes voluntary participation.**

Recommendation 4.8: Build independent international biological monitoring capacity and integrate it with national security decision-making.

- **Partnerships.** The federal government should invest in bilateral and multilateral partnerships that build environmental monitoring and genomic surveillance capacity in high-risk regions, reducing dependence on national reporting systems subject to political and economic pressure to delay disclosure.
- **Funding for data sharing.** The federal government should provide contracts or grants to organizations with demonstrated operational capacity to build and sustain global genomic data-sharing infrastructure. This should include attribution-based sharing models that have proven most effective at incentivizing submission from scientists and governments in low- and middle-income countries.
- **Integrated national threat assessment capability.** The intelligence community's biological threat assessment capabilities must be integrated with international monitoring outputs. The government must establish protocols that ensure anomalous biological signals detected abroad reach national security decision-makers quickly and that attribution analysis (discussed in the next chapter) is initiated in parallel with public health response. With allies and partners, building a secure continuous biological monitoring data-sharing capacity will also enhance preparedness.

Chapter 5

Finding 5.1: The convergence of AI and biotechnology has outpaced US biological attribution science. Detection methods calibrated to known pathogens cannot identify agents designed to evade them. Research programs capable of matching engineered agents to their laboratory of origin—including IARPA's FELIX program—have demonstrated proof-of-concept but have not been transitioned into operational capabilities. If the United States does not build and sustain these capabilities now, it will be unable to adapt them as masking techniques develop alongside AI-enabled biodesign.

Recommendation 5.1: Congress should authorize and fund the acceleration and transition of IARPA FELIX and successor programs from research into a standing operational capability. This should include: development of a national engineering-signature library covering laboratories known to work with dangerous pathogens; expansion of AI-based attribution tools capable of identifying fully AI-generated agents; and integration of these tools with existing intelligence and law enforcement databases to enable rapid comparison against known actors and methods.

Finding 5.2: Biological attribution requires a continuous, population-scale surveillance baseline against which anomalous signals can be identified. The United States currently lacks this baseline. Significant reductions in federal public health funding have degraded already limited state and local surveillance capacity. Existing surveillance platforms are not designed to generate data with the chain-of-custody and integration requirements necessary to support law enforcement and intelligence analysis. Without a surveillance architecture built for attribution as well as public health response, the United States cannot reliably distinguish an engineered pathogen from a data gap.

Recommendation 5.2: Congress should authorize and fund a national integrated biological attribution system. Core requirements of the system include: AI-enabled anomaly detection for both known pathogens and novel biological signals; and explicit chain-of-custody and data-sharing protocols designed to support attribution as well as public health response. Diagnostic surge capacity for pandemic-capable pathogens should be a funded component, with federal guidance requiring states and cities to share de-identified healthcare data to support real-time national disease pattern analysis. This attribution system must be integrated into the national biological monitoring system recommendations identified in Chapter 4.

Finding 5.3: The interagency protocols designed to enable rapid biological incident response depend on White House leadership and a career civil service with institutional knowledge of biological crisis response. No single federal authority is currently charged with overseeing these protocols, and it is unclear whether they continue to be exercised. The elimination of dedicated White House biosecurity expertise and coordination capacity, and significant reductions in the federal public health workforce across multiple agencies, have materially reduced the US government's ability to execute rapid, coordinated attribution under crisis conditions.

Recommendation 5.3: Biosecurity coordination should be strengthened.

- Congress should codify requirements for a regularly updated Biological Incident Notification and Assessment Protocol in statute and designate a senior White House official with clear authority to convene and direct interagency biological incident response.
 - The administration should restore dedicated biosecurity coordination capacity at the National Security Council and mandate regular interagency exercises with state and local partners to ensure the protocol can be executed under crisis conditions. These exercises should explicitly rehearse the adjudication of competing public health and law enforcement priorities to preserve both response speed and evidentiary integrity.
-

Finding 5.4: Rapid biological attribution at scale depends on international data-sharing arrangements, sample access, and investigative cooperation built over decades through CDC, USAID, and WHO engagement. The systematic withdrawal from these partnerships is eroding US access to the international biological data and cooperation that attribution requires. The UNSGM, the principal international instrument for investigating alleged biological weapons use, has never been invoked for a biological event and lacks standing technical capacity to do so effectively.

Recommendation 5.4: The administration should revitalize global health and global health security capabilities.

- CDC and US State Department global health security personnel (and USAID personnel should the agency be reestablished by a future administration) should be deployed as an integral component of US biodefense capacity, and should reengage with WHO to rebuild data-sharing arrangements that support rapid attribution.
- Congress should fund revitalization of the UNSGM as an investigative mechanism for biological events of unknown origin, including building standing technical capacity for rapid biological investigation.
- The administration should negotiate bilateral and multilateral data-sharing agreements that explicitly incentivize building and strengthening regional information-sharing networks capable of identifying patterns consistent with an emerging biological risk, whether deliberate, accidental, or naturally occurring.

Chapter 6

Finding 6.1.1: The biotechnology ecosystem is fragmented across multiple critical infrastructure sectors with no standing mechanism for cross-sector coordination between government and the full range of relevant private actors—biotech firms, frontier AI companies, and pharmaceutical manufacturers. This structural gap cannot be closed within the current architecture.

Recommendation 6.1.1: Designate the biotechnology ecosystem, including AI-enabled technologies, as a critical infrastructure sector and establish a standing public-private coordination mechanism.

- Designate the biotechnology ecosystem as a new critical infrastructure sector within the national framework either by statute or through executive action. Designate DOD, DHS, and HHS as co-Sector Risk Management Agencies (SRMAs) for the biosecurity and biotechnology sector.
- Organize a standing Biotechnology Sector Coordinating Council that includes: DNA synthesis providers and gene foundries; synthetic biology and advanced biologics firms; genomic data repositories and computational biology platforms; cloud-based bioinformatics infrastructure; AI model developers whose systems materially enable biological design; advanced laboratory automation and distributed biomanufacturing operators; high-containment laboratory operators; and critical biological supply chain nodes.
- A corresponding government coordinating council should include agencies capable of providing and examining threat information related to biosecurity, including, *inter alia*, DOD, DHS, HHS, DOE, NIST, CAISI, and components of the intelligence community.

Finding 6.1.2: No existing authority ensures that the United States can rapidly access private-sector industrial capacity—including laboratory throughput, sequencing pipelines, biomanufacturing lines, and personnel—to respond to a biological emergency. This is a major national readiness gap.

Recommendation 6.1.2: Establish a National Strategic Bio Readiness Reserve modeled on the Civil Reserve Air Fleet.

- Enact legislation authorizing the National Strategic Bio Readiness Reserve (NSBRR) as a standing, public-private reserve of design and discovery, compute power, advanced development, and biomanufacturing capacity, modeled on the Civil Reserve Air Fleet.
- Participating firms should receive guaranteed procurement preferences, standby payments, or other contractual incentives in exchange for committing a defined share of compute and production capacity during declared emergencies.
- Specify activation triggers, response timelines, incentive structures—including minimum government contract guarantees, liability and indemnification provisions—and annual audit requirements. Activation authority should be clearly delineated—potentially requiring a joint determination or consultation by the secretary of defense and the secretary of health and human services. Staged activation levels would distinguish between regional outbreaks and major public health emergencies.
- Conduct a DPA Section 705 Industrial Base Survey of each NSBRR segment and a parallel voluntary request for information on laboratory and manufacturing capability to properly scope the reserve.
- The federal government should conduct regular readiness exercises (e.g., annually or biennially) with participating firms to test surge production timelines, supply chain resilience, and distribution capabilities before future emergencies occur.
- Establish a specific Genesis Mission challenge to build bio emergency response compute capacity, including pre-negotiated access agreements for crisis activation.

Finding 6.1.3: The US government lacks a comprehensive, current assessment of frontier AI CBRN capabilities and safeguards. Although DPA authorities to compel or incentivize the deployment of safeguards exist, they are underutilized.

Recommendation 6.1.3.1: Strengthen CBRN safeguards on AI models.

- Perform a DPA Section 705 Industrial Base Survey of frontier AI developers' CBRN capabilities and safeguards including mandatory pre-deployment testing results, with specific focus on biosecurity and CBRN risks.
- Provide incentives for the development and procurement of frontier AI models with built-in CBRN safeguards using DPA Title III authorities. Use DPA Title I authorities to compel deployment of known, technically feasible safeguards when firms have chosen not to implement them.
- Design and fund a challenge under BARDA DRIVE to develop novel CBRN safeguards that can prevent, deter, or detect AI-enabled bioweapon development.

Finding 6.2.1: The US government's protocols for responding to a biological incident of unknown or contested origin, including the Biological Incident Notification and Assessment Protocol and the NSC-convened interagency response mechanism, depend on robust oversight and a well-trained civil service. Without a restored and regularly exercised decision architecture, the speed advantage that early detection and attribution can provide cannot be translated into effective response.

Recommendation 6.2.1.1: Restore and exercise the interagency biological incident decision architecture.

- The NSC should reestablish senior-level oversight of the Biological Incident Notification and Assessment Protocol and associated interagency response mechanisms.
- Congress should mandate and fund regular senior-level exercises of the biological incident response system across HHS, DOD, DHS, and the FBI, with performance metrics publicly reported.
- Data-sharing agreements between public health, law enforcement, and intelligence agencies that enable rapid response to an incident of unknown origin should be reviewed, updated, and confirmed in place before the next crisis.

Recommendation 6.2.1.2: Establish a National Clinical Trials Emergency Authority.

- Congress should authorize a standing National Clinical Trials Emergency Authority (NCTEA) within HHS with operational command over federally supported research during declared public health emergencies.
- The NCTEA should have authority to designate priority platform trials, mandate hospital participation, and enforce radical transparency standards for all federally designated emergency trials.

Recommendation 6.2.1.3: Establish a formal pathway for mutual recognition of clinical trial results from allied-country regulatory systems.

- The FDA should develop a formal framework for mutual recognition of clinical trial data generated by equivalently rigorous regulatory bodies belonging to US allies.
- Trials conducted to FDA-compatible standards should be eligible for FDA approval application without duplicative domestic trials.

Finding 6.2.2: The US clinical trial system cannot generate evidence at the speed of a biological crisis. COVID-19 demonstrated that a system optimized for hypothesis-driven, grant-funded research fails as a crisis response tool. The administrative, financing, and data infrastructure required for rapid, decisive evidence generation does not exist in peacetime and cannot be improvised under crisis conditions.

Recommendation 6.2.2.1: Build a crisis-ready learning clinical trial system as permanent national security infrastructure.

- Congress should provide the NIH with specific mandates for emergency surge capacity, including authority to deploy large-scale funding within days to weeks and a protected

Emergency Research Reserve Fund with statutory triggers modeled on the G7 100-day mission frameworks.

Recommendation 6.2.2.2: Create a standing national clinical trial infrastructure.

- Congress should mandate that any hospital with over 100 beds participating in CMS maintain a standing clinical trial office, supported by pre-negotiated IDIQ reimbursement contracts; designate a permanent emergency clinical trial officer at each participating institution with authority to prioritize national protocols during a declared emergency.
- Congress should designate clinical research data as a public utility, adopt a privacy-protective national patient identifier, establish a member-governed nonprofit clinical trial data access broker with standardized and cyber-secure data, and authorize emergency data-sharing authority modeled on the Defense Production Act for use during declared public health emergencies.
- Congress and federal agencies should fund and maintain disease-agnostic platform trial infrastructures with preapproved master protocols for priority syndromes, and require the FDA to publish prespecified, pathogen-class guidance on evidentiary requirements and acceptable endpoints before the next biological emergency.

Finding 6.2.3: The private sector plays an indispensable role in developing and manufacturing pandemic-relevant therapeutics, but persistent market failures discourage sustained investment in infectious disease research and development during peacetime. Without predictable economic incentives and structured government engagement, promising medical countermeasures may never reach clinical testing, and critical manufacturing capacity may not be available when a biological emergency occurs.

Recommendation 6.2.3: Establish durable public-private mechanisms to sustain investment in pandemic-relevant therapeutics and manufacturing capacity.

- Congress should authorize and fund advance market commitments for priority pandemic countermeasures to provide predictable demand signals and reduce private-sector investment risk.
- HHS should establish a permanent industry-government coordination forum to align expectations regarding clinical trial requirements, regulatory pathways, manufacturing readiness, and emergency response planning before a crisis emerges.

Finding 6.2.4: Clinical trial readiness depends not only on infrastructure and authorities, but also on the routine exercise of the systems, relationships, and procedures that enable rapid research during a public health emergency. Without regular testing and evaluation, critical capabilities—rapid contracting, centralized review, data sharing, and trial activation—may degrade, creating delays when a biological crisis occurs.

Recommendation 6.2.4: HHS should conduct recurring national-scale exercises that stress-test the full clinical trial system. HHS should establish measurable readiness indicators—including time to protocol activation, time to first patient enrollment, time to contract execution, time to data availability, and participant representativeness—and publicly report performance results.

The NCTEA (described above) should use these exercises to identify operational bottlenecks and produce after-action assessments with corrective action plans.

Chapter 7

Finding 7.1: The United States lacks the functional biological data infrastructure, high-containment defense research and development capacity, and operational response teams needed to maintain deterrence by resilience against AI-enabled biological threats. Existing programs are fragmented across agencies without a unifying institutional structure, standing operational commitment, or the legal authorities necessary for durable public-private engagement.

Recommendation 7.1: Charter a National Institute for Bioresilience (NIBR) as federally authorized public-private national infrastructure for biological AI and data.

- Congress should enact legislation creating and funding the National Bioresilience Institute as a federally chartered, public-private consortium. The statute should define NIBR's mission; specify its governance structure; designate a lead agency and its relationship with DOE, DOD, HHS, NIST, and DHS; and direct the lead agency to specify data types, establish data-sharing processes, and clarify applicable privacy protections and legal safe harbors.
- Congress should authorize and appropriate dedicated compute and data infrastructure for NIBR sufficient to train, evaluate, and serve frontier biological AI models at national scale—on the order of thousands of frontier-class accelerators, petabyte-scale curated biological datasets, secure enclaves for sensitive functional data, and pre-negotiated access to current frontier general-purpose and biological foundation models (Evo, ESM, AlphaFold-class, and successors). Capacity should support the conduct of high-containment research on AI-enabled biothreats and operational reach-back by BEST teams during incident response. NIBR could also use its expertise to develop recommendations and guardrails mitigating the risks of autonomous labs (discussed in Chapter 3).
- Congress should direct NIBR to establish formal partnerships with self-driving laboratories and cloud-laboratory operators for industrialized functional data generation, and joint operating authority with DOE national laboratories as well as with DOD and DHS labs (USAMRIID and NBACC) for high-containment research and dual-use evaluation.
- DOE, in coordination with DOD, HHS, NIST, and DHS, should take initial steps to implement NIBR's base tier leveraging and aligning with the Genesis Mission where appropriate, including creating federated biological data and controlled workspaces on Genesis infrastructure, establishing initial partnerships with private firms and universities through CRADAs and MOUs, and curating AI-ready biological datasets and AI-enabled workflows tied to DOE's biotechnology challenge.
- OSTP should examine different models for incentivizing private entities to participate across NIBR's three tiers and in NSBRR's reserve capability model (see Chapter 6).

- Congress should enact legislation establishing Bio Emergency Services Teams (BEST) as a cross-cutting operational response initiative drawing on NIBR analytical reach-back and NSBRR surge capacity, defining mission triggers and BEST's relationship with states, FEMA, CDC, ASPR, DOD, and national laboratories, and providing for regular preparedness exercises.

Finding 7.2. National resilience requires a robust private biotechnology research and manufacturing sector. The US private biotech sector is a national security asset whose health directly determines US biodefense capacity. Supply chain dependencies in biomanufacturing, APIs, sequencing, and drug discovery create exploitable vulnerabilities, and the federal government has not systematically treated private-sector biotech strength as a national security objective. Major cuts in NIH, NSF, and university funding along with reductions in the S&T workforce in key agencies are having direct impacts on America's capacity to win the bio race. Sustained, multiyear investments in research and human capacity are urgently needed to grow American talent to power the US AI-enabled biotechnology ecosystem, both in government and the private sector. In addition, offshoring of bio manufacturing and reliance on China for the bio supply chain pose long-term threats to the sector's sustainability. Furthermore, the US must have a skilled workforce to ensure the vitality and innovative capacity of the sector.

Recommendation 7.2.1: Rebuild and protect US biotech innovation capacity for the long term.

- Fund relevant research at levels commensurate with the competitive challenge.
- Extend financial incentives to on-shore pharmaceutical discovery and production through NIH grant funding, ARPA-H project funding, enhanced R&D tax credits, favorable capital gains treatment for biotech investments, and BARDA where appropriate.
- Expand incentives for domestic biotechnology innovation through NIH, ARPA-H, BARDA, enhanced R&D tax credits, and long-term support for translational biotechnology infrastructure and clinical trial capacity.

Recommendation 7.2.2: Work to reduce current pharmaceutical supply chain dependencies on China.

- Congress should direct USTR, HHS, DOD, and the Department of Commerce to conduct a comprehensive assessment of US dependency on foreign sources for critical pharmaceutical ingredients, precursor chemicals, biologics manufacturing inputs (including antibiotics and antivirals) essential to national defense and pandemic response. This should include assessment of how the capabilities of US allies and partners could contribute to advancing the objective of reducing dependencies on China.
- Congress should establish a national security review process for pharmaceutical in-licensing and out-licensing agreements involving China, as well as data-sharing arrangements that could create systemic vulnerabilities in critical sectors of the US biotechnology ecosystem. This review could be modeled on the Committee on Foreign Investment in the United States (CFIUS).

Recommendation 7.2.3: Strengthen US industrial capacity and supply chain security.

- Build support for a Boost American Bio Manufacturing Act modeled on the CHIPS and Science Act to reduce growing dependencies on China via a comprehensive industrial policy supporting the US biotech and pharmaceutical industries. This could include targeted subsidies, public R&D investment, tax incentives, workforce development, place-based economic strategies, and collaborative production with aligned foreign partners.
- Establish federal incentives for strengthening US biotech and pharmaceutical domestic manufacturing capacity, reducing single-point-of-failure risks (such as dependency on Chinese contract bio manufacturers), and improving long-term supply chain resilience. This could include investment tax credits, guaranteed procurement floors, accelerated depreciation, and BARDA- and ARPA-H-supported manufacturing partnerships for strategically essential products.
- Designate domestic genomic sequencing platform capacity as a critical supply chain priority alongside API, KSM, and CDMO manufacturing. Direct the Department of Commerce and BARDA to assess current US market share in high-throughput sequencing platforms, identify domestic capacity gaps, and develop incentive programs for domestic sequencing platform development and manufacturing analogous to those proposed for pharmaceutical APIs.
- Extend the procurement restrictions established under the NDAA for named Chinese biotechnology entities to additional categories of Chinese-origin sequencing and biological AI infrastructure.

Recommendation 7.2.4: Ensure access to a skilled workforce necessary to support the biotechnology sector.

- Create targeted federally funded fellowships in AI and bioscience to create a pipeline of AI-bio scientists with government service experience, appropriate clearances, and institutional knowledge.
- Enable rotational exchanges between government laboratories and leading AI firms.
- Provide special compensation authority and competitive retention bonuses to reduce the federal-private wage differential for AI-biology researchers to make long-term government careers feasible for top-tier AI-biology experts.

Finding 7.3. US security and resilience are directly related to global biosecurity and resilience. Recent reductions in US support for global health institutions have reduced our insight into and ability to respond to global health threats. Strengthening international capacity and resilience through public-private partnerships is a sustainable way to address these challenges.

Recommendation 7.3.1: Expand US engagement and investment in global biosecurity. Redeploy US public health experts to countries globally.

- Expand foreign assistance for biosurveillance and attribution. Include biosecurity provisions in foreign assistance grants and contracts.
- Reengage with WHO, Gavi, Global Fund, CEPI, and the Pandemic Fund.

Recommendation 7.3.2: Establish new pathways for global public-private partnerships to support bioresilience.

- Create a global senior leaders network initiative with a philanthropically supported secretariat. The initiative should focus on near-term actions and incentives to create sustainable biosecurity and resilience globally. Build out a sustainable process that continually solicits national actions improving global prevention and resilience.
- Build MCM regional manufacturing hubs through development banks and G20 DFIs (building on the G7 DFI MCM Surge Financing Initiative, 100 Days Mission, and CEPI 3.0).

Recommendation 7.3.3: Establish a US-China Track 1.5 Dialogue on AI-Enabled Biological Risks.

- Convene a sustained Track 1.5 dialogue involving government officials, AI and biotechnology leaders, scientists, and nongovernmental experts from both countries focused on reducing catastrophic risks arising from the convergence of AI and biotechnology.
- Prioritize practical risk-reduction measures, including DNA synthesis screening, laboratory biosafety and biosecurity practices, and approaches to preventing the design or creation of potentially pandemic-capable biological agents without appropriate human oversight.
- Develop confidence-building measures and crisis communication mechanisms to reduce the risks of misperception, misattribution, or escalation following unusual biological incidents or advances in high-risk biological capabilities.

Finding 7.4: Following US withdrawal from WHO, the United States lacks adequate bilateral biosurveillance architecture to detect emerging biological threats as early as multilateral surveillance systems previously enabled it to do. This gap is not merely diplomatic; it creates potential for both intelligence loss and warning failure with direct national and economic security consequences.

Recommendation 7.4: Build independent biosurveillance capacity and bilateral biosecurity engagement infrastructure to replace multilateral coverage lost through WHO withdrawal.

- The State Department and HHS should negotiate and implement biomonitoring MOUs with priority partners mandating real-time genomic data sharing, standardized data formats, minimum sequencing coverage, and mutual notification timelines.
- NIBR should build and operate independent global pathogen monitoring capacity using federated sequencing nodes and baseline characterization to ensure the United States detects emerging threats without dependence on WHO reporting systems.
- The US government should convene an annual Biosecurity and Bioresilience Summit with senior officials, scientific leads, and private sector biosecurity actors from priority partner nations to institutionalize bilateral coordination and provide a forum for sensitive threat intelligence sharing.

APPENDIX 2

Legislative Proposals on Biosecurity and AI

Representative Biosecurity, Biodefense, and Pandemic Preparedness Legislation (118th–119th Congress)

Congress	Bill	Sponsor(s)	Key Provisions	Markup?	Floor Vote?
118th	H.R. 4702 – Securing Gene Synthesis Act	Eshoo	• Establishes regulatory framework for gene synthesis providers • Requires aequivalence screening • Requires customer verification and records-keeping • Institutes federal procurement restrictions • Jump-starts industry standards development	No	No
118th	H.R. 4704 – AI and Biosecurity Risk Assessment Act	Eshoo, Crenshaw	• Mandates a formal federal risk assessment of risks at the intersection of AI and biosecurity • Requires identification of capability thresholds and actors to misuse AI for bio • ASPR must assess existing biodefense and public health preparedness systems for AI risks • Requires federal assessment of safeguards, including gene synthesis screening, data access controls, and model release practices	No	No

118th	S. 2356 – Gene Synthesis Safety and Security Act	Hickenlooper, Budd	• Requires HHS to update and formalize gene synthesis screening guidance, potentially moving toward enforceable standards • Establishes baseline security practices for providers, including sequence screening and customer vetting • Encourages standardization across the industry, reducing uneven compliance • Promotes international coordination to mitigate cross-border evasion risks • Focuses on strengthening existing voluntary frameworks rather than creating a fully new regulatory regime	No	No
118th	H.R. 8333 – BIOSECURE Act	Wenstrup et al.	• Institutes ban on federal procurement by and contracts with a “biotechnology company of concern” • Targets foreign adversary-linked firms • Institutes compliance audits	Yes	Yes
118th	S. 3558 – Prohibiting Foreign Access to American Genetic Information Act	Peters, Hagerty	• Prohibits federal agencies from contracting with foreign adversary-linked biotech companies • Requires DNI-led intelligence assessment of risks related to foreign access to US genomic data • Establishes data security expectations for entities handling sensitive genetic data • Addresses risks from biotech services that process, store, or analyze genomic information • Reinforces supply chain scrutiny similar to BIOSECURE but with stronger data-security framing	Yes	No
118th	H.R. 9194 – Nucleic Acid Standards for Biosecurity Act	Caraveo, McCormick	• Directs NIST to develop technical standards for nucleic acid sequence screening systems • Supports creation of reference datasets and benchmarking tools for identifying sequences of concern • Establishes public-private collaboration framework involving industry, academia, and government • Promotes interoperability and consistency across providers • Focuses on technical infrastructure rather than regulatory mandates, enabling future regulation	Yes	No

118th	H.R. 2791 – NBACC Authorization Act	Trone	• Codifies DHS National Biodefense Analysis and Countermeasures Center (NBACC) in statute • Defines mission areas: biothreat characterization, attribution science, countermeasure support • Enables interagency access to high-containment research capabilities • Formalizes NBACC's role in supporting national biodefense and intelligence missions	No	No
118th	S. 1210 – NBACC Authorization Act	Cardin	• Senate companion to H.R. 2791 • Integrates NBACC authorities into the Homeland Security Act • Clarifies operational scope and governance structure of the center	No	No
118th	S. 2333 – Pandemic and All-Hazards Preparedness and Response Act	Casey et al.	• Reauthorizes core public health emergency preparedness authorities under PAHPA • Strengthens ASPR's role in coordinating federal emergency response • Enhances Strategic National Stockpile management and supply chain resilience • Expands biosurveillance and data-sharing systems • Supports BARDA's role in advanced development of medical countermeasures • Addresses lessons learned from COVID-19 response gaps	Yes	No
118th	H.R. 1425 – No WHO Pandemic Preparedness Treaty Without Senate Approval Act	Tiffany	• Requires any WHO pandemic agreement to be treated as a treaty requiring Senate ratification • Limits executive authority to enter binding international pandemic arrangements • Strengthens congressional oversight of global health governance commitments	Yes	Yes
119th	H.R. 3029 – Nucleic Acid Standards for Biosecurity Act	Salinas, McCormick	• Reintroduces NIST-led effort to develop DNA synthesis screening standards and benchmarks • Expands stakeholder engagement mechanisms • Supports development of technical tools for automated sequence risk detection • Positions NIST as a central coordinator for biosecurity standards infrastructure	Yes	No

119th	S. 3469 – BIOSECURE Act 2025	Peters, Hagerty	• Reintroduces BIOSECURE framework • Expands definition of biotechnology companies of concern • Maintains procurement and funding prohibitions across federal ecosystem • Continues emphasis on genomic data security and supply chain decoupling	No	No
119th	S. 3741 – Biosecurity Modernization and Innovation Act	Cotton, Klobuchar	• Establishes a Commerce Department–led regulatory regime for nucleic acid synthesis • Requires mandatory screening of sequences and customers across providers • Creates a federally maintained “sequences of concern” list • Introduces a regulatory sandbox to support innovation while managing risk • Seeks to streamline fragmented federal biosecurity/biosafety authorities • Represents a shift toward comprehensive statutory regulation, not just guidance	No	No
119th	H.R. 7653 – Biodefense Diplomacy Enhancement Act	Self et al.	• Directs State Department to expand international coordination on biotechnology and biodefense • Strengthens cooperation with allies and NATO on biosecurity standards and threat response • Promotes export controls and global norms for sensitive biotech • Positions biodefense as a foreign policy priority area	Yes	No
119th	S. 501 – Strategy for Public Health Preparedness and Response to Artificial Intelligence Threats	Budd, Markey	• Requires HHS to develop national strategy for AI-related public health threats • Addresses both risk (misuse) and opportunity (surveillance, response) • Requires coordination across public health and national security agencies • Builds on earlier AI-bio risk assessment efforts	No	No
119th	S. 891 – Bipartisan Health Care Act	Wyden, Sanders	• Incorporates PAHPA reauthorization elements into broader health legislation • Supports public health infrastructure modernization • Enhances emergency preparedness and response capacity • Integrates pandemic preparedness with broader health system policy	No	No

119th	H.R. 4207 – No WHO Pandemic Preparedness Treaty Without Senate Approval Act	Tiffany	• Reintroduces requirement for Senate approval of WHO pandemic agreements • Continues effort to limit executive authority in global health governance	No	No
119th	S. 1983 – No WHO Pandemic Preparedness Treaty Without Senate Approval Act	Johnson	• Senate companion to H.R. 4207 • Reinforces treaty-based interpretation of international pandemic commitments	No	No
119th	H.R. 5747 – Preventing Illegal Laboratories and Protecting Public Health Act of 2025	Griffith	• Requires federal agencies (including HHS, CDC, and DHS) to identify, track, and respond to unauthorized or illegal biological laboratories operating in the United States • Requires development of interagency coordination mechanisms for detecting and investigating unlicensed or unsafe lab operations • Requires strengthening of inspection, enforcement, and shutdown authorities for facilities handling biological agents without proper authorization • Requires improved information sharing among federal, state, and local public health and law enforcement entities • Requires development of guidance for state and local officials on identifying illegal lab activity and responding to biohazards • Requires assessment of public health risks posed by unregulated biological research and materials handling	No	No

119th	S.1387 / H.R.2756 – National Biotechnology Initiative Act of 2025	Young, Bice	• Requires establishment of a national biotechnology strategy coordinating federal efforts across HHS, DOE, DOD, NSF, and other agencies • Requires creation or expansion of a National Biotechnology Coordination Office to align policy, funding, and research priorities • Requires development of long-term R&D priorities in biotechnology, including biosecurity, biomanufacturing, and synthetic biology • Requires support for domestic biomanufacturing capacity and supply chain resilience • Requires programs to secure biotechnology infrastructure from foreign adversary influence or exploitation • Requires coordination with industry and academia to accelerate innovation while managing safety and security risks • Requires reporting to Congress on progress, risks, and strategic gaps in US biotechnology leadership	No	No
119 th	S. 4729 – Biosecurity Smuggling Deterrence Act	Cotton	• Would impose mandatory minimum sentences for conspiracy to smuggle biological agents into the United States and for making false statements to federal agents in connection with such smuggling, and for other purposes	No	No
119th	S. 4023 / H.R. 7832 – America’s Living Library Act	Young, Khanna	• Launches a national genomic mapping effort: Directs the Department of the Interior to collect, catalog, and sequence DNA from animals, plants, fungi, and microbes on US public lands • Establishes standardized, AI-ready genomic databases and physical specimen collections to support research and innovation in medicine, agriculture, and biotechnology • Places the project within the US Geological Survey and frames biological data as a strategic national asset	No	No
119th	S. 4227 – Preventing Illegal Laboratories and Protecting Public Health Act	Cortez Mastro	• Requires distributors of designated highly pathogenic agents to verify purchasers’ identities and intended use, and maintain detailed electronic records of all transfers to improve traceability and oversight • Directs the federal government to conduct recurring evaluations of US high-containment labs, develop national biosafety and biosecurity standards, and establish a federal biosafety team to coordinate with state and local officials	No	No

119th	S. 4363 – Engineering Biology Readiness Act	Kaine	• Directs federal agencies to analyze current and emerging national security and public safety risks from engineering biology (including misuse or accidental release) and identify gaps in US preparedness • Requires recommendations to update biodefense, biosafety, and biosecurity authorities, regulations, research priorities, and programs, including proposed legislative changes and funding needs, developed with input from industry, academia, and civil society	No	No
119 th	H.R. 8981 – SCALE Biology Act	McClain Delaney	• Creates a dedicated program at NIST to develop measurement science, reference materials, and technical standards for engineering biology, biomanufacturing, and biometrology • Authorizes NIST to partner with universities, national labs, industry, and other federal agencies to develop testing capabilities, support collaborative research, train the engineering biology workforce, and improve technology commercialization	No	No
119th	S. 2676 / H.R. 7801 – Cloud LAB Act of 2026	Young, Obernolte	• Requires federal support for development and use of cloud-connected or remote-access laboratory platforms • Requires establishment of secure access controls and identity verification mechanisms for remote laboratory operation • Requires development of standards for monitoring, logging, and auditing experiments conducted through cloud labs • Requires evaluation of biosecurity risks associated with remote experimentation, including misuse or unauthorized access • Requires coordination with NIST and other agencies to develop technical standards for secure cloud lab infrastructure	No	No

119th	S. 854 / H.R.1864 – Risky Research Review Act	Paul, Griffith	• Requires establishment or strengthening of federal review processes for high-risk (“dual-use” or “gain-of-function”) biological research • Requires identification of research categories that pose elevated biosecurity or biosafety risks, including work that could enhance pathogen capabilities • Requires development of standardized criteria for evaluating risks and benefits of such research • Requires agencies to implement or expand pre-funding and pre-execution review mechanisms for risky research proposals • Requires coordination across federal agencies to ensure consistent oversight of dual-use research of concern (DURC) • Requires development of risk mitigation requirements and safeguards for approved research	No	No
119th	S.4770 / H.R. 9307 – Web of Biological Data Act of 2026	Young, Padilla, Rounds, Kim; Auchincloss, Van Epps	• Directs the Department of Energy to establish the Web of Biological Data (WOBD)—a secure, centralized portal where US researchers can access high-quality, AI-ready biological datasets from federal and other participating sources • A DOE national laboratory would launch an initial pilot within one year to develop cybersecurity protections, data curation methods, and biological data management tools • Treats biological data as a strategic national asset	No	No
119th	S.4069 / H.R.7907 – AI-Ready Bio-Data Standards Act	Young; Khanna	• Directs NIST to develop definitions, standards, and frameworks so biological datasets are consistently formatted and usable for AI models • Requires qualifying federally funded research projects to produce AI-ready biological data, while minimizing compliance burdens, testing the standards • Establishes a technical advisory group, engages academia and the private sector, aligns federal agency data policies, and requires ongoing reporting to Congress to improve interoperability and support US biotechnology competitiveness	No	No

Representative AI Governance Legislation³⁶³ (118th–119th Congress)

Congress	Bill	Sponsor(s)	Key Provisions	Markup?	Floor Vote?
118th	H.R. 5628 – Algorithmic Accountability Act of 2023	Clarke	• Requires FTC rulemaking to mandate algorithmic impact assessments • Evaluates bias, discrimination, privacy, and security risks • Requires documentation of training data, model design, and intended uses • Imposes ongoing monitoring and mitigation obligations for high-risk systems • Establishes ex ante governance rather than reactive enforcement	No	No
118th	H.R. 3369 – Artificial Intelligence Accountability Act	Harder	• Directs NTIA to study AI accountability mechanisms across sectors • Evaluates effectiveness of accountability tools in reducing risk • Requires stakeholder consultations (industry, academia, public) • Requires analysis of definitions of trustworthy/responsible AI • Produces policy recommendations to Congress for future regulation	Yes	No
118th	S. 3205 – Federal Artificial Intelligence Risk Management Act of 2023	Warner	• Requires federal agencies to adopt NIST AI Risk Management Framework • Standardizes processes for identifying, assessing, and mitigating AI risks • Applies to procurement and deployment decisions • Establishes baseline governance expectations for federal contractors	No	No

³⁶³ Does not include broad innovation or sector-adoption bills unless they imposed safety, disclosure, liability, or governance obligations.

118th	S. 3312 – Artificial Intelligence Research, Innovation, and Accountability Act of 2023	Thune et al.	• Establishes regulatory framework for high-impact and critical-impact AI systems • Requires detailed transparency reports on system purpose, data, performance, and safeguards • Mandates pre-deployment and recurring risk assessments • Creates TEVV-style testing, evaluation, and validation requirements • Authorizes enforcement including civil penalties and deployment restrictions	Yes	No
118th	S. 4178 – Future of Artificial Intelligence Innovation Act of 2024	Cantwell et al.	• Codifies AI Safety Institute at NIST • Develops testing infrastructure including red-teaming and benchmarking • Supports standards, metrics, and evaluation tools • Facilitates international coordination on AI safety standards	Yes	No
118th	S. 5616 – Preserving American Dominance in Artificial Intelligence Act of 2024	Romney et al.	• Creates AI Safety Review Office within Commerce • Requires pre-deployment review of frontier models • Mandates submission of evaluation results by developers • Allows review, rereview, and potential deployment prohibition • Focus is on CBRN, cyber, and national security risks	No	No
118th	H.R. 1736 – Generative AI Terrorism Risk Assessment Act	Pfluger	• Requires DHS, in coordination with the DNI, to conduct annual assessments of terrorism threats involving generative AI • Requires analysis of how terrorists may use generative AI for radicalization, recruitment, propaganda, and operational planning • Requires evaluation of risks that AI could enable or enhance chemical, biological, radiological, or nuclear (CBRN) threats • Requires development of recommendations for countermeasures to mitigate AI-enabled terrorism risks	No	No

118th	S. 3162 – TEST AI Act	Luján et al.	• Requires NIST/DOE AI testbeds • Develops TEVV methodologies • Simulates adversarial conditions • Identifies vulnerabilities and failure modes • Develops safeguards and standards • Coordinates across agencies and labs • Establishes pilot programs	No	No
118th	S. 4230 / H.R. 9737 – Secure Artificial Intelligence Act of 2024	Warner, Tillis; Ross, Beyer	• Requires NIST to evaluate whether existing vulnerability-management frameworks (including the National Vulnerability Database (NVD)) are adequate for artificial intelligence systems • Requires NIST to update, adapt, or develop new processes to ensure that AI-related security vulnerabilities can be identified, categorized, and tracked • Requires development of a taxonomy or classification system for AI vulnerabilities • Requires the CISA to adapt the Common Vulnerabilities and Exposures (CVE) program or establish a parallel mechanism capable of handling AI-specific vulnerabilities • Requires establishment of a public, voluntary database or mechanism to track AI safety and security incidents	No	No
118th	H.R. 9497 – AI Advancement and Reliability Act of 2024	Obernolte et al.	• Establishes NIST Center for AI reliability • Develops evaluation methods and benchmarks • Advances measurement science • Creates multistakeholder consortium • Expands hiring authority • Authorizes funding	No	No
118th	S. 4769 – VET AI Act	Hickenlooper et al.	• Requires NIST, DOE, and NSF to develop TEVV guidelines • Requires specification of minimum documentation and disclosure expectations to support evaluation • Addresses privacy, transparency, risk mitigation • Requires development of recommendations regarding the use of third-party auditors and accredited testing entities	No	No

118th	H.R. 9466 – AI Development Practices Act	Obernolte et al.	• Requires NIST to identify and catalog current industry practices related to the development of AI systems • Requires evaluation of how these practices address key risk and performance domains • Requires NIST to identify best practices for communicating AI system information	No	No
118th	S. 3050 – AI Advancement Act	Peters et al.	• Creates AI bug bounty programs • Requires the federal government to conduct a comprehensive analysis of vulnerabilities in AI systems used for national security and defense • Requires evaluation of how AI systems could be exploited by adversaries • Requires assessment of the security risks associated with AI integration into critical infrastructure and defense systems	No	No
119th	TRUMP AMERICA AI Act (not formally introduced yet)	Blackburn	• Establishes a single federal “rulebook” for AI, replacing fragmented state-level regulation with a unified national framework • Imposes a statutory “duty of care” on AI developers and deployers to prevent and mitigate foreseeable harms arising from AI systems • Creates a federal liability regime for AI systems, moving toward a product-liability-style framework • Sunsets (repeals) Section 230 protections after a transition period, significantly expanding legal exposure for AI platforms and online services • Requires AI systems to incorporate child safety protections • Establishes federal policy to protect intellectual property rights in AI training and outputs	No	No
119th	S. 2938 – Artificial Intelligence Risk Evaluation Act of 2025	Hawley, Blumenthal, Blackburn	• Creates a DOE program to evaluate frontier AI models before deployment • Requires developers to provide information needed for government risk evaluations • Assesses risks including loss of control, cyber misuse, CBRN enablement, model theft, and jailbreak vulnerabilities	No	No

119th	S. 1633 – TEST AI Act of 2025	Luján, Durbin, Blackburn, Risch	• Directs NIST and DOE to establish AI testing and evaluation testbeds • Develops standards for security, reliability, robustness, privacy, and bias • Requires a national testing strategy and regular reports to Congress	No	No
119th	H.R. 8094 – AI Foundation Model Transparency Act of 2026	Beyer, Lawler, Jacobs, Fitzpatrick	• Directs the FTC to establish transparency requirements for covered foundation models • Requires disclosures regarding training data, capabilities, evaluations, and known risks • Improves information available to regulators, researchers, and downstream users	No	No
119th	H.R. 1694 – Artificial Intelligence Accountability Act	Harder, Kelly	• Requires a federal study of AI accountability mechanisms • Evaluates options for audits, certifications, liability rules, and independent oversight • Recommends legislative and regulatory approaches for trustworthy AI	No	No
119th	S. 2164 / H.R. 5511 – Algorithmic Accountability Act of 2025	Wyden; Clarke	• Requires impact assessments for high-impact automated decision systems • Evaluates risks of discrimination, privacy harms, and other adverse impacts • Authorizes FTC oversight and periodic reassessments after deployment	No	No
119th	H.R. 7058 – Foreign Adversary AI Risk Assessment and Diplomacy Act	Himes	• Requires assessments of AI threats posed by foreign adversaries • Directs development of diplomatic and national security strategies to address those threats • Requires recurring reporting on emerging international AI risks	No	No
119th	H.R. 8283 – Deterring American AI Model Theft Act of 2026	Huizenga	• Addresses theft of advanced US AI models by foreign actors • Authorizes sanctions and visa restrictions for certain offenders • Improves coordination on protecting AI intellectual property and model weights	No	No
119th	H.R. 5287 – China Advanced Technology Monitoring Act	Vindman	• Requires annual assessments of China's AI and semiconductor capabilities • Supports national security planning through technology monitoring • Identifies strategic advances relevant to US competitiveness	No	No

119th	H.R. 6996 – Full AI Stack Export Promotion Act	Sponsor not verified	• Promotes export of US AI hardware, software, and models to trusted partners • Supports strategic competition with China through allied technology deployment	No	No
119th	S. 4113 – AI Guardrails Act of 2026	Slotkin	• Establishes guardrails for Defense Department AI systems • Focuses on safe, reliable, and accountable military AI deployment	No	No
119th	S. 1792 – AI Whistleblower Protection Act	Grassley, Coons, Blackburn, Klobuchar, Hawley, Schatz	• Protects employees who disclose AI safety, security, or legal concerns • Prohibits retaliation and establishes enforcement mechanisms	No	No
119th	S. 1837 – DEFIANCE Act of 2025	Sponsor not verified	• Creates a federal civil cause of action for victims of AI-generated intimate-image deepfakes • Allows recovery of damages and other legal relief	Yes	Passed Senate
119th	S. 146 – TAKE IT DOWN Act	Cruz, Klobuchar	• Criminalizes distribution of nonconsensual intimate images, including AI-generated deepfakes • Requires covered online platforms to remove reported content within statutory deadlines	Yes	Yes
119th	S. 1367 / H.R. 2794 – NO FAKES Act	Coons, Blackburn, Klobuchar, Tillis	• Creates a federal property right in an individual's voice and likeness • Allows lawsuits over unauthorized AI-generated digital replicas while preserving exceptions such as news and parody	No	No
119th	H.R. 5272 – Protect Elections from Deceptive AI Act	J. Johnson	• Restricts materially deceptive AI-generated election communications • Creates civil remedies for candidates harmed by prohibited deepfakes	No	No
119th	H.R. 5388 – American AI Leadership and Uniformity Act	Baumgartner	• Establishes a federal AI governance framework • Temporarily preempts certain state AI laws while the federal framework is developed	No	No
119th	S. 2750 – SANDBOX Act	Cruz	• Creates a federal AI regulatory sandbox • Allows qualifying developers to test AI systems under temporary regulatory flexibility	No	No
119th	S. 4216 – GUARDRAILS Act	Schatz, Coons, Blunt Rochester, Murphy, Kim, Duckworth	• Prevents implementation of federal AI policies that broadly preempt state law • Preserves states' authority to regulate AI	No	No

119th	S. 3557 – States’ Right to Regulate AI Act	Markey, Wyden, Sanders, Schiff, Booker, Luján, Klobuchar, Padilla, Durbin	• Blocks use of federal funds to implement AI preemption policies • Affirms states’ ability to enact and enforce AI regulations	No	No
-------	--	---	--	----	----

APPENDIX 3

Perspective on Liability for Bio-risks of AI Models

Note: This information and analysis are provided to inform further exploration of options for liability regimes relevant to the bioconvergence challenge.

Today, the default tort liability doctrine for AI models is not a single AI-specific statute. It is ordinary state tort law applied to AI-related fact patterns as they reach courts. Tort liability regimes are generally structured in four parts: (1) establishment of a duty of care, (2) breach of that duty, (3) a causal link flowing from the breach, and (4) the harm caused or damages incurred by the breach. Thus, in order to bring a tort claim against an AI model or biological design tool (BDT) developer, plaintiffs would have to demonstrate causation between the use of AI-generated outputs and the “concrete[,] particularized, [and] actual or imminent” injury in fact.³⁶⁴

Given the range and severity of risks associated with BDTs, *ex post* tort liability should not be the primary mechanism to prevent harm. *Ex ante* regulations that ensure the models meet certain safety benchmarks prerelease and structured access regimes that ensure models are only used by those without ill intent post-release are key elements of prevention here. Absent real regulation, tort liability can shift private sector incentives in the direction of safety—and even if some regulations were adopted, tort liability would layer onto them as an additional legal safeguard.

Since tort law is largely state common law, the applicable standard of care and evidentiary expectations can vary across jurisdictions. Absent new legislation—either state or federal—judges and juries may struggle with adjudicating who owes what duty, to whom, and who “caused” the harm in a legally cognizable way.³⁶⁵ A new law could clarify these issues by imposing a specific duty of care on AI model developers, in this case on the application-layer developers, i.e. the BDT model developers, since they are the least cost avoiders. Such a law could take one of two viable paths.

Negligence Opinions

The first path is one that adopts a negligence standard similar to California Senate Bill (SB) 1047, which was passed by the legislature and vetoed by Gov. Gavin Newsom.³⁶⁶ SB 1047 represented one of the most detailed state-level attempts to construct a liability regime for frontier AI models. Though it failed, its liability architecture is among the most fully developed templates currently

364 *Lujan v. Defs. of Wildlife*, 504 US 555, 557 (1992); Ketan Ramakrishnan et al., *US Tort Liability for Large-Scale Artificial Intelligence Damages* (RAND Corporation, Aug. 21, 2024), https://www.rand.org/content/dam/rand/pubs/research_reports/RRA3000/RRA3084-1/RAND_RRA3084-1.pdf.

365 Ramakrishnan et al., *US Tort Liability for Large-Scale Artificial Intelligence Damages*.

366 Alex Friedland, “Governor Newsom Vetoes Sweeping AI Regulation, SB 1047,” Center for Security and Emerging Technology, Oct. 17, 2024, <https://cset.georgetown.edu/article/governor-newsom-vetoes-sweeping-ai-regulation-sb-1047/>.

available in US law and can be used to model any state or federal liability regime covering AI models and BDTs. While the bill applied only to frontier-scale models, defined as those that cost more than \$100 million to train and were trained using more than 10^{26} floating-point operations (FLOPs), the standard of care it used could be adopted for a larger share of models.³⁶⁷ SB 1047 proposed a “reasonable care” standard for liability that drew from OpenAI’s Preparedness Framework, Google’s Frontier Safety Framework, and Anthropic’s Responsible Scaling Policy.^{368,369}

Before beginning to train a covered model, developers were required to draft a written Safety and Security Protocol (SSP) addressing foreseeable risks of catastrophic misuse, implement the capability to perform a full model shutdown, and designate a senior individual responsible for compliance.³⁷⁰ The SSP was to be reviewed annually and retained for the life of the model plus five years. The attorney general could demand access to an unredacted SSP at any time. Before public deployment, developers were required to conduct pre-deployment testing, including adversarial red-teaming, to assess whether the model could contribute to “critical harm,” and to retain results in a form allowing third-party replication. The bill’s liability trigger was specifically tied to SSP compliance: A developer that suffered a critical harm incident but had maintained a sufficient SSP in good faith had a meaningful basis to resist liability. Conversely, a developer that failed to maintain the SSP, skipped required testing, or could not demonstrate adherence to industry best practices was presumptively in breach of the reasonable care standard—a rebuttable presumption approach rather than true strict liability. The critical harm standard would cover BDTs whose output materially contributed to the design or synthesis of a biological weapon capable of mass casualties. While SB 1047 linked enforcement exclusively to the California attorney general, an alternative framework could additionally create an explicit private right of action.

Ferae Naturae

An alternative, stricter path than that of SB 1047 would be to adopt the doctrine of *ferae naturae*, or the design defect standard from product liability law to impose strict liability on AI model developers.³⁷¹ The doctrine of *ferae naturae* holds that if an individual introduces a nonindigenous wild animal to a locale, and damages occur to another’s person or property due to “a dangerous propensity that is characteristic of wild animals of the particular class, or of which the possessor knows or has reason to know,” the possessor is held strictly liable.³⁷² Strict liability in this context means that “the possessor has exercised the utmost care to confine the animal, or otherwise prevent it from doing harm,” and even if the possessor believes they have “tamed” the animal’s “dangerous propensities.”³⁷³

Those who support *ferae naturae* strict liability argue that similar to wild animals, AI models are inherently unpredictable and nondeterministic (due to the black box nature of their neural networks). They can be trained (fine-tuned with reinforcement learning) without completely

367 California State Legislature, *SB-1047 Safe and Secure Innovation for Frontier Artificial Intelligence Models Act*, 2023–2024 Regular Session, https://leginfo.ca.gov/faces/billTextClient.xhtml?bill_id=202320240SB1047.

368 Ibid.

369 Ketan Ramakrishnan, “Tort Law Is the Best Way to Regulate AI,” *Wall Street Journal*, Sept. 24, 2024, <https://www.wsj.com/opinion/tort-law-is-the-best-way-to-regulate-ai-california-legal-liability-065e1220>.

370 California State Legislature, *SB-1047*.

371 Trent S. Kannegieter, “Nondeterministic Torts: A Technical Approach to AI Liability,” 135 *Yale Law Journal* 1707 (2026).

372 *Restatement (Second) of Torts* § 507 (Am. L. Inst. 1977).

373 Ibid. § 507 cmt. c.

eliminating the underlying dangers, and safeguards are applied externally to reduce risk (i.e., cages for wild animals, classifiers and filters for AI models) without eliminating it.³⁷⁴ Similarly, deploying an unbounded, nondeterministic system into a safety-critical environment constitutes a defective product design. Because standard quality assurance testing is inherently incomplete for nondeterministic systems like AI models, the design itself is fundamentally risky when error cost is high.³⁷⁵ Under modern product liability law, this triggers strict liability, meaning deployers can be held legally responsible even if they followed best practices and did not act negligently or with malintent.

Section 230

Either of the approaches described could benefit from another legal shift. Section 230 of the Communications Decency Act (47 U.S.C. § 230) immunizes interactive computer service providers from liability for content created by third-party users. AI developers have attempted to invoke Section 230 as a shield against liability for harms caused by AI-generated content.³⁷⁶ However, Section 230 was designed to protect passive conduits for third-party speech, not active AI systems that generate their own outputs. In the biological AI context, Section 230 immunity is particularly weak. An AI biological design tool does not publish third-party information; it generates novel biological sequences through its own computational processes. The generated sequence is the AI developer's own "content" in every meaningful sense. Accordingly, a law clarifying that generative AI outputs are not third-party content under Section 230's text would strengthen the liability imposed on AI model developers.

Relatedly, there is scholarly disagreement over whether imposing tort liability on AI developers for harmful outputs constitutes government regulation of speech, in which case the law would need to survive some form of judicial scrutiny depending on what type of speech the AI output is categorized as.³⁷⁷ Alternatively, it is also plausible that such a law would be interpreted as simply regulating conduct (the design and deployment of a commercial AI system) with only incidental effects on expression, which would receive more deferential review.³⁷⁸

Finally, liability could also be layered at the access level to incentivize robust compliance. In a fashion similar to how US Treasury Department sanctions issued by the Office of Foreign Assets Control utilize a strict liability standard for civil penalties under 50 U.S.C. § 1705, a legally mandated or supported structured access regime could institute know-your-customer (KYC) requirements with a similarly strict liability standard.³⁷⁹

374 Kannegieter, "Nondeterministic Torts."

375 Ibid.

376 Peter J. Benson and Valerie C. Brannon, *Section 230 Immunity and Generative Artificial Intelligence*, Legal Sidebar LSB11097 (Dec. 28, 2023), <https://www.congress.gov/crs-product/LSB11097>.

377 Ilan Kogan, "Artificial Intelligence, Existential Risk, and the First Amendment," 27 *University of Pennsylvania Journal of Constitutional Law* 156 (2025).

378 Peter Salib, "AI Outputs Are Not Protected Speech," 102 *Washington University Law Review* 83 (2024).

379 Liana W. Rosen, *U.S. Sanctions: Overview for the 119th Congress*, In Focus IF12390 (Congressional Research Service, Sept. 4, 2025), <https://www.congress.gov/crs-product/IF12390>.

APPENDIX 4

Further Discussion of the Defense Production Act

The Defense Production Act (DPA) of 1950 is the primary source of presidential authorities to expedite and expand the supply of materials and services from the US industrial base needed to “promote the national defense.”³⁸⁰ Relevant DPA authorities include:

- Title I (Priorities and Allocations) allows the president/delegates to require private sector actors to *accept and prioritize contracts* for materials and services necessary to promote the national defense and to *allocate* materials, services, and facilities necessary to promote the national defense.
- Title III (Expansion of Productive Capacity and Supply) allows the president/delegates to *financially incentivize* the domestic industrial base expansion for the production and supply of critical materials and goods for national defense purposes, with specific limitations and requirements.
- Title VII (General Provisions) includes many key definitions and additional DPA authorities. Two particularly relevant parts of Title VII are Sections 705 and 708. Section 705 allows the president/delegates to *survey* the private sector to obtain information relevant to national defense goals. Section 708 allows the president/delegates to *establish voluntary agreements* with private industry *immunized from antitrust liability*.

Since its inception, the DPA has been reauthorized more than fifty times for varying periods.³⁸¹ Historically, the DPA has typically been reauthorized through standalone legislation, but in both fiscal years 2019 and 2026, the National Defense Authorization has been used as the reauthorization vehicle. In December 2025, the FY2026 NDAA extended the DPA’s expiration date from Jan. 30, 2026 to Sept. 30, 2026.³⁸² At the time of writing, this remains the current expiration date of the DPA.

Organization of DPA Authority

The DPA has also been amended several times over the decades, many times through reauthorization measures. The DPA Reauthorization of 2009 established the central organizing

380 ASPR, “Defense Production Act Overview,” US Department of Health and Human Services, accessed June 29, 2026, <https://aspr.hhs.gov/legal/DPA/Pages/Overview.aspx>.

381 Recent reauthorization extension periods have lasted for five years (2014–2019), six years (2019–2025), and less than one year (December 2025–September 2026). In the past forty years, there have been at least two lapses in authorization (October 2025–November 2025 and October 1990–August 1991). Christopher T. Mann, *Reauthorizing the Defense Production Act*, CRS Insight IN12484 (Congressional Research Service, Jan. 14, 2025), <https://www.congress.gov/crs-product/IN12484>.

382 Ibid.

body for DPA authority, the Defense Production Act Committee (DPAC), which is authorized under Title VII of the statute. The DPAC was intended to serve as the primary interagency mechanism for advising the president about the use of the DPA authority and coordinating DPA activities across government.

Following the DPAC's establishment, President Obama issued Executive Order 13603 on "National Defense Resources Preparedness" to provide additional guidance. EO 13603 designates the secretary of homeland security as chair of the DPAC, which was then sub-delegated to the FEMA administrator (given FEMA's role as the interagency DPA convener before DHS's existence³⁸³). Importantly for purposes of this report, EO 13603 also delegates the president's authority under the DPA for "health resources"³⁸⁴ to the secretary of health and human services. The EO similarly delegates DPA authority for other sectors to corresponding secretaries (e.g., Department of Agriculture for food resources, Department of Energy for energy resources). The sector-specific delegations under EO 13603 are nonexclusive, identifying the "lead agency" for defining, prioritizing, and allocating resources in the relevant sector.

The DPAC was again restructured during the DPA's 2014 reauthorization, and recently members of Congress and other stakeholders have proposed several other substantive amendments to the DPA that include restructuring the DPAC, increasing the carryover limit for the DPA Fund, expanding Title III authorities, and more.³⁸⁵

DPA and "Health Resources"

The Administration of Strategic Preparedness and Response (ASPR) is an HHS division that leads the nation's medical and public health preparedness and response programs for disasters and public health emergencies. This includes pandemics and bioterrorism/biological threats. The ASPR thus plays a major role in managing HHS' DPA authority. ASPR leads and implements the use of HHS' delegated DPA Title I authorities to ensure the US has a robust supply of necessary "health resources." For this purpose, ASPR has developed the Health Resources Priorities and Allocations System (HRPAS).³⁸⁶

ASPR also directs the HHS DPA Title III Program, which is responsible for strengthening national defense capabilities in "health resource" domains through the use of Title III financial incentive authorities. Specifically, within ASPR, the Office of Industrial Base Management and Supply Chain (IBMSC) is responsible for managing and administering Title III authorities to expand domestic production capabilities for, *inter alia*, medical countermeasures, essential medicines, and critical inputs.

383 Exec. Order No. 12919, 59 Fed. Reg. 29,525 (June 7, 1994).

384 Exec. Order No. 13603, § 801(i), 77 Fed. Reg. 16,651, 16,672 (Mar. 22, 2012) (defining "health resources" drugs, biological products, medical devices, materials, facilities, health supplies, services and equipment required to diagnose, mitigate or prevent the impairment of, improve, treat, cure, or restore the physical or mental health conditions of the population).

385 Christopher T. Mann, *Reauthorizing the Defense Production Act*, CRS Insight IN12484 (Congressional Research Service, Jan. 14, 2025), <https://www.congress.gov/crs-product/IN12484>; Michael Brown, "Reauthorize the Defense Production Act—But Make Changes for the 21st Century," *Breaking Defense*, May 29, 2024, <https://breakingdefense.com/2024/05/reauthorize-the-defense-production-act-but-make-changes-for-the-21st-century/>.

386 ASPR, "Title I Priorities and Allocations," US Department of Health and Human Services, <https://aspr.hhs.gov/legal/DPA/Pages/Title-I.aspx>; "Health Resources Priorities and Allocations System (HRPAS)," 89 Fed. Reg. 9,143 (Feb. 9, 2024), <https://www.federalregister.gov/documents/2024/02/09/2024-01947/health-resources-priorities-and-allocations-system-hrpas>.

Information Gathering

Before major federal actions concerning the private sector capabilities, the government typically undertakes information gathering to identify the relevant private industry participants, their current and emergency-response capabilities, and any industry shortcomings. As of now, there are two key methods available: voluntary requests for information; and issuance of compulsory surveys using DPA authority.

Requests for Information

ASPR can issue a request for information (RFI) for use by the HHS DPA Title III Program to solicit similar information from necessary industry segments. The purpose of such an RFI would be to improve understanding of current and potential future gaps in the domestic supply chain and how such shortfalls would benefit from HHS' use of DPA Title III financial incentive authorities.

For example, in February 2024, ASPR issued an RFI to gather market research on the domestic production capability and capacity of the Public Health Industrial Base (PHIB). This RFI focused specifically on essential medicines, medical countermeasures, and critical inputs. It asked PHIB members for specifics regarding production capabilities, including questions about manufacturing equipment, current constraints, and BARDA TRL and MRL stages.³⁸⁷ It also asked for information about both the “monthly warm-base manufacturing requirement versus the maximum surge capacity” of productive capabilities, seeking to understand the gap between current operational and emergency-response capabilities.

However, the RFI process is both voluntary and somewhat imprecise, as recipients cannot be chosen strategically.

Industrial Base Surveys—DPA Section 705

The Bureau of Industry and Security (BIS) within the Department of Commerce conducts studies to assess the capabilities of the US industrial base to support the national defense for DPA purposes.³⁸⁸ Under the DPA, Industrial Base Surveys need to be related to assessment of capabilities of the US industrial base to support national defense, and must be well-defined and nonduplicative.³⁸⁹ In partnership with the requesting agency (e.g., HHS), BIS develops a list of persons to be surveyed based on their likelihood to have relevant information.³⁹⁰ Neither the proximity of association within the supply chain nor the nature of the “person” selected is relevant to assessing the sufficiency of that person’s inclusion in the survey. Therefore, there is substantial latitude for who may be surveyed.

As for survey content, there is also a wide array of knowledge that can be gathered. Industrial Base Surveys may require information about employment, R&D, sources of supply, manufacturing processes, customers, business strategy, finances, and other factors affecting the industry’s ability to meet both the needs of national defense programs (such as the National Strategic Bio

387 BARDA uses Technology Readiness Levels (TRLs) to measure the maturity of a medical technology and Manufacturing Readiness Levels (MRLs) to measure whether that technology can be manufactured reliably at scale.

388 50 U.S.C. § 4555; Exec. Order No. 13603, § 104, 77 Fed. Reg. 16,651 (Mar. 22, 2012).

389 Bureau of Industry and Security, “US Industrial Base Surveys Pursuant to the Defense Production Act of 1950,” 80 Fed. Reg. 41,430 (July 15, 2015).

390 Ibid.

Readiness Reserve, or NSBRR) and international competitiveness. Survey recipients must also respond appropriately. The DPA provides civil remedies and criminal penalties for failure of survey recipients to supply the information sought.

HHS and BIS are well-suited to develop more extensive lists of survey recipients and the substantive scope required for the formation of entities called for in this report, namely the NSBRR and the National Institute for Bioresilience, or NIBR. The survey should inquire into both baseline and emergency response capabilities and could even seek to understand how survey recipients would respond to each activation tier contemplated.

The advantages of conducting an ASPR/HHS RFI over an Industrial Base Survey are (1) greater local expertise within the HHS, allowing for a broader and more thorough information gathering process, and (2) closer proximity to the HHS DPA Title III Program, which will likely be heavily involved in the NSBRR's development. Conversely, such an RFI is voluntary as opposed to required by the compulsory approach of an Industrial Base Survey under DPA Section 705. Also, an Industrial Base Survey can more easily be used in conjunction with Title I compulsory procurement and allocation authorities, by seeking answers to specific questions that will allow for the necessary specification of immediate courses of action under Title I.

Antitrust Exemption—DPA Section 708

The COVID-19 pandemic demonstrated that scientific breakthroughs can occur rapidly while manufacturing and supply chains lag. Operation Warp Speed, leveraging DPA authorities and BARDA funding, dramatically accelerated vaccine development. Yet scaling production required *ad hoc* negotiations, improvisational contracting, and extensive federal coordination with manufacturers and suppliers. The federal government relied heavily on DPA prioritization authority³⁹¹ and voluntary agreements under DPA Section 708³⁹² to align industry production.³⁹³

DPA Section 708 authorizes the president to bring industry together to voluntarily plan and coordinate actions to support national defense preparedness, including expanded production capacity and supply. The statute allows participants to invoke a defense to certain antitrust laws when they act in accordance with an approved voluntary agreement.³⁹⁴ For example, in August 2020, the Federal Emergency Management Agency (FEMA) published its final version of a voluntary agreement under the DPA Section 708, creating a Committee for the Distribution of Healthcare Resources Necessary to Respond to a Pandemic.³⁹⁵ The committee, given a five-year mandate, was created with the articulated purpose to “maximize the effectiveness of the manufacture and distribution of Critical Healthcare Resources nationwide.” Members of the committee included the chairperson; representatives from FEMA, HHS, DOJ, and other federal agencies; and private sector participants. Then, the committee created subcommittees and plans of action for pandemic response workstreams. Each plan of action related to one or more

391 50 U.S.C. § 4511.

392 50 U.S.C. § 4558; 44 C.F.R. 332.

393 Michael H. Cecire and Heidi M. Peters, *Defense Production Act (DPA): Recent Developments in Response to COVID-19*, CRS Insight IN11470 (Congressional Research Service, July 28, 2020), https://www.congress.gov/crs_external_products/IN/PDF/IN11470/IN11470.5.pdf.

394 Valerie C. Brannon and Michael H. Cecire, *The Role of Section 708 of the Defense Production Act in the Federal Government's Response to COVID-19: Antitrust Considerations*, Legal Sidebar LSB10534 (Congressional Research Service, Aug. 20, 2020), <https://www.congress.gov/crs-product/LSB10534>.

395 FEMA, “Voluntary Agreement Under Section 708 of the Defense Production Act; Manufacture and Distribution of Critical Healthcare Resources Necessary to Respond to a Pandemic,” 85 Fed. Reg. 50,035 (Aug. 17, 2020).

of the committee's twelve objectives and entailed a "plan of action agreement."

Within the duration of the plan of action agreement, private sector participants could use the DPA Section 708 provisions as a defense to any civil or criminal action brought for violation of antitrust laws (or similar state laws) with respect to any action taken under the plan of action. Of course, to invoke this defense, participants had to demonstrate their actions were within scope, or were: (1) taken in the course of developing or carrying out the plan of action/agreement, (2) in accordance with the terms of the plan of action/agreement, and (3) otherwise in compliance with the provisions of DPA Section 708.³⁹⁶

A voluntary agreement under DPA Section 708 may meaningfully improve the efficiency and efficacy of the NSBRR proposed in this report and its formation by enabling coordination between and among public and private resources. A committee similar to FEMA's COVID-19 Pandemic-Response Committee could complement the formation of the NSBRR effectively. Though the August 2020 Pandemic-Response Committee focused mainly on the manufacture and distribution of "Critical Healthcare Resources," a NSBRR-voluntary agreement under Section 708 may add some additional aims. In fact, it may be most prudent to establish a committee tasked with assisting the formation of the NSBRR itself, such that supply chains and emergency response capacities among the three industry segments can be properly coordinated. Such a committee would allow for data sharing among NSBRR participants for coordination purposes in a way that could responsibly account for sensitive information, trade secrets, and antitrust concerns while allowing setup of a robust NSBRR. Like the previous Pandemic-Response Committee, a future committee would entail oversight by designated representatives from FEMA, HHS, and other relevant federal agencies. The attorney general and chair of the FTC, or their delegates, may also join the committee and attend meetings.

Furthermore, with the DPA up for reauthorization in September 2026, Congress may consider modestly amending Section 708 for these purposes. A Congressional Research Service report on FEMA's proposed voluntary agreement recommended that Congress consider amending Section 708 to remove its exclusion of "unfair or deceptive acts or practices" from the scope of its antitrust defense, so as to allow use of the Section 708 defense in state law consumer protection cases.³⁹⁷ This would be a significant step toward defining the already quite limited Section 708 defense in a clear way that leverages the statutory scheme's advantages in allowing industry collaboration for specific, government-monitored national defense purposes.

To be clear, establishing a voluntary agreement setting up a NSBRR committee under DPA Section 708 would allow private sector participants to coordinate and communicate with each other and government agencies in ways they may not otherwise be able or likely to. Subcommittees established for (or even within) each market segment identified in the NSBRR proposal can work together in competitively sensitive areas without the fear of an antitrust lawsuit (as long as they work within the boundaries and under supervision of the committee). For example, firms with certain levels of laboratory capacity can come together and share technical information about the resources they would each need, dividing up roles or duties in relation to specific response scenarios. It also allows for formal coordination between government agencies and private industry, keeping actors on the same page and an open channel of communication by which to arrive at quick, collaborative solutions.

³⁹⁶ Ibid.

³⁹⁷ Brannon and Cecire, *The Role of Section 708*.

Gaps for NSBRR and NIBR Implementation

Chapters 6 and 7 propose two new capabilities to strengthen America’s resilience to biological threats and reinforce deterrence by denial. The National Strategic Bio Readiness Reserve (NSBRR) would establish a biological strategic reserve that harnesses private-sector capabilities and expertise to respond quickly and effectively to localized public health threats and catastrophic disasters. The National Institute for Bioresilience (NIBR) would create a public-private data, compute, and model infrastructure that enables low-risk but cutting-edge biological research; supports R&D collaboration between government labs and private partners on high-risk pathogens; and deploys Bio Emergency Services Teams (BEST) to respond to biological incidents. While current law and policy provide a helpful framework for the two initiatives, tailored authorities and funding would ensure their durability and flexible operation, allowing the United States to proactively shape its response to a dynamic biological threat landscape.

A. NSBRR

Current Law and Policy

Public Health Service Act authorities support a reserve-style biological response architecture. The Public Health Service Act directs HHS to establish interagency arrangements under which it can assume operational control of emergency public health and medical response assets in the event of a public health emergency, while coordinating with other federal agencies.³⁹⁸ The act authorizes several reserve-like mechanisms, including the National Disaster Medical System,³⁹⁹ the Medical Reserve Corps,⁴⁰⁰ and the USPHS Ready Reserve Corps.⁴⁰¹ Additional HHS authorities support grants and cooperative agreements⁴⁰² with public and private entities around public health emergency preparedness and acquisition of countermeasures.⁴⁰³

398 42 U.S.C. § 300hh.

399 The NDMS expressly contemplates private-entity participation agreements, reimbursement arrangements, ongoing readiness activities, and exercises; and it authorizes appointment of intermittent disaster-response personnel with federal liability protection while acting within scope. 42 U.S.C. § 300hh-11.

400 42 U.S.C. § 300hh-15.

401 42 U.S.C. § 204.

402 42 U.S.C. § 247d (In a public health emergency, “the Secretary may take such action as may be appropriate to respond to the public health emergency, including making grants, providing awards for expenses, and entering into contracts and conducting and supporting investigations into the cause, treatment, or prevention ...”); *Centers for Disease Control and Prevention*, “Public Health Emergency Preparedness (PHEP) Program,” last reviewed Jan. 9, 2025, <https://www.cdc.gov/readiness/php/phep/index.html>; ASPR, “Hospital Preparedness Program (HPP),” US Department of Health and Human Services, accessed June 29, 2026, <https://aspr.hhs.gov/HealthCareReadiness/HPP/Pages/default.aspx>.

403 ASPR, *Justification of Estimates for Appropriations Committees, Fiscal Year 2026*, US Department of Health and Human Services (2025), <https://www.hhs.gov/sites/default/files/fy-2026-gdm-cj.pdf>; 42 U.S.C. § 247d-6b. The PREP Act provides immunity from

Defense Production Act (DPA) authorities are directly applicable to the private capacity reserve model. The statute codifying the Civil Reserve Air Fleet (CRAF)—described further in Chapter 6—defines reserve aircraft as allocated under Section 101 of the DPA or made available under DOD contracts.⁴⁰⁴ Federal agencies have described Title III’s industrial base investment authority as a tool for expanding domestic public health manufacturing capacity and ensuring supply chain resilience,⁴⁰⁵ while Title VII provides for voluntary agreements with industry to address threats to the United States’ preparedness programs.⁴⁰⁶

Existing emergency response policies create a policy and legal framework for operational deployment in disasters. The Stafford Act⁴⁰⁷ and FEMA National Response Framework⁴⁰⁸ create a policy architecture for emergency and disaster response, while Emergency Service Function 8 coordinates federal assistance specifically around public health and medical services.⁴⁰⁹ Activation of NSBRR’s reserve capacity and BEST teams (described below) would accordingly align and require coordination with existing emergency response functions.

HHS authorities allow for interagency data sharing at incident sites, although more comprehensive authorities may be necessary for NSBRR’s response capabilities. While Congress has directed HHS to facilitate interagency and external data use agreements to prepare for, monitor and respond to declared or potential public health emergencies,⁴¹⁰ additional authorities could create more flexibility for NSBRR’s emergency response teams to aggregate and conduct analyses across critical public and private datasets, consistent with their role as a biological “fire department.”

Legal Gaps

1. Express authority to reserve private biological capacity in peacetime

While CRAF is a voluntary program in which carriers commit assets through contractual arrangements, Congress later reinforced the model with a statutory policy of maintaining CRAF readiness through peacetime augmentation.⁴¹¹ There is no comparably specific statute for reserving diagnostics, sequencing, or other pathogen response capacities across private firms. Legislation strengthening interstate reciprocity or emergency practice authority may also be valuable to facilitate the rapid deployment of clinical or private lab personnel across state lines.

2. A bespoke certification and modular capability framework

liability for covered persons and countermeasures when the HHS secretary issues an emergency declaration. 42 U.S.C. § 247d-6d.

404 10 U.S.C. § 961.

405 ASPR, “Title III: Expansion of Productive Capacity and Supply,” US Department of Health and Human Services, accessed June 29, 2026, <https://aspr.hhs.gov/legal/DPA/Pages/Title-III.aspx>.

406 50 U.S.C. § 4558. Title VII also calls for an “executive reserve,” creating a direct precedent for holding private-sector expertise in reserve for national emergencies. GAO, *Defense Production Act: Information Sharing Needed to Improve Use of Authorities*, GAO-25-107688 (June 2025), <https://www.gao.gov/assets/gao-25-107688.pdf>.

407 42 U.S.C. § 5121 et seq.

408 FEMA, *National Response Framework*, 5th ed. (2023), <https://www.fema.gov/emergency-managers/national-preparedness/frameworks/response>.

409 ESF 8 includes procurement and distribution of vaccines and other medical supplies. ASPR, “Emergency Support Functions,” US Department of Health and Human Services, accessed June 29, 2026, <https://aspr.hhs.gov/legal/Pages/Emergency-Support-Functions.aspx>.

410 Pandemic and All-Hazards Preparedness and Advancing Innovation Act of 2019, Pub. L. No. 116-22, § 306, 133 Stat. 905, 941 (2019).

411 US Department of Transportation, “Civil Reserve Air Fleet,” last updated Feb. 23, 2024, <https://www.transportation.gov/mission/administrations/intelligence-security-emergency-response/civil-reserve-airfleet-allocations>; 10 U.S.C. § 9517.

NSBRR's model depends on modular certification of facilities and biological services, such as special pathogen diagnostics, high-containment R&D support, and field sequencing. A national legal taxonomy and certification framework would accordingly facilitate NSBRR's reserve capability model.

3. Liability, indemnification, and compensation rules

Existing liability protections for participation in public health emergency response are tied to covered countermeasures and an explicit emergency declaration by the HHS secretary.⁴¹² The breadth of activities envisioned under NSBRR, including with respect to peacetime preparation, necessitates a broader, NSBRR-specific liability and indemnification regime.

4. Dedicated appropriations or no-year readiness funding

The National Strategic Computing Reserve blueprint, described below, demonstrates the importance of sustained funding to implement readiness programs and exercises.⁴¹³ Designated appropriations, including possibly no-year funds, would enable the development and continued readiness of NSBRR's reserve capability.

B. NIBR

Current Law and Policy

Congress has authorized analogous public-private entities advancing public health and innovation in critical technologies.⁴¹⁴ In 1987, Congress authorized the government to participate in SEMATECH, a government-industry research consortium that contributed to revitalizing the United States' semiconductor industry in the 1990s by accelerating innovation, establishing shared facilities for experimentation, and developing industry-wide standards. Over the course of five years, the DOD provided an annual \$100 million grant to defray R&D costs.⁴¹⁵

The National Strategic Computing Reserve (NSCR) provides a conceptual model for NIBR's base of coordinated research compute. The NSCR blueprint proposed a reserve of compute, software, data, and expertise spanning government, academia, nonprofits, and industry, explicitly modeled as a standing strategic reserve to be mobilized in crises.⁴¹⁶ NIBR goes a step further than the

412 42 U.S.C. § 247d-6d.

413 Dylan Cohen and Kush Patel, *National Strategic Computing Reserve Tabletop Exercise After Action Report*, IDA Paper P-3003599 (Institute for Defense Analyses, December 2024), <https://www.ida.org/-/media/feature/publications/n/na/national-strategic-computing-reserve-tabletop-exercise-after-action-report/3003599.ashx>.

414 GAO, *Federally Created Entities: An Overview of Key Attributes*, GAO-10-97 (October 2009), <https://www.gao.gov/assets/gao-10-97.pdf>. Public-health analogues include the CDC Foundation, which Congress authorized in 1992 as an independent nonprofit to support CDC, and the Foundation for the NIH, which Congress established to support NIH through public-private partnerships.

415 GAO, *Lessons Learned from SEMATECH*, GAO/RCED-92-283 (September 1992), <https://www.gao.gov/assets/rced-92-283.pdf>.

416 National Science and Technology Council, *National Strategic Computing Reserve: A Blueprint* (Executive Office of the President, October 2021), <https://www.nitrd.gov/pubs/National-Strategic-Computing-Reserve-Blueprint-Oct2021.pdf>. The NSCR has not been fully implemented. Following a 2024 tabletop exercise, the National Science and Technology Council observed that noncrisis operation of the NSCR could include serving as a repository for relevant information: a coordinating entity with detailed knowledge of existing and potential interagency cooperation mechanisms. The council also highlighted the need for formal authorities enabling government agencies to participate in the Reserve and share resources with other agencies. Cohen and Patel, *National Strategic Computing Reserve Tabletop Exercise After Action Report*.

NSCR blueprint as it envisions a standing source for computational power and biological data, focused on ongoing research coordination rather than crisis deployment.⁴¹⁷

Current biosafety regulations, data security guidance, and collaboration models supply a regulatory framework for NIBR's public-private research collaboration on high-risk pathogens. Commerce and HHS regulations require registration, designated responsible officials, security barriers, and security-risk assessments for access to select agents and toxins.⁴¹⁸ NIST has provided guidance to manage the cybersecurity and privacy risks that arise from processing genomic data.⁴¹⁹ Moreover, the National Nuclear Security Administration's tech transfer program facilitates robust collaboration between national labs and industry partners in commercializing high-risk technologies.⁴²⁰ These authorities provide a concrete basis for tiered access controls, misuse prevention measures, and high-risk tech collaborations between government labs and the private sector, core components of NIBR's mid-tier defense R&D.

Existing biodefense and public health law supports countermeasure and response functions during public health emergencies. The Public Health Service Act authorizes ASPR within HHS to maintain a stockpile of drugs, vaccines, and medical devices, and to deploy the stockpile in response to public health emergencies.⁴²¹ Congress similarly established the Public Health Emergency Medical Countermeasures Enterprise, an interagency partnership of federal medical countermeasure experts, to develop strategies for enhancing medical countermeasures for CBRN threats in the United States and advising the HHS secretary on improving access to medical products during public health emergencies.⁴²² These programs provide compelling precedents for preparing and deploying capabilities in response to public health crises,⁴²³ although they do not extend explicitly to forward deployment of operational teams with biological diagnostic capabilities and related technical expertise.⁴²⁴

Existing antitrust safe-harbor law provides a legal backstop for entities participating in NIBR's consortium model. As private companies engage with NIBR's innovation ecosystem and defense R&D tiers, legislation such as the National Cooperative Research and Production Act can

417 The NIH has pursued an analogous albeit more narrow effort to make commercial cloud environments available for biomedical research through negotiated partnerships with industry cloud providers, under its Strategic Plan for Data Science. National Institutes of Health, "About STRIDES," NIH STRIDES Initiative, accessed June 29, 2026, <https://cloud.nih.gov/about-strides/>.

418 42 C.F.R. § 73.10 (Restricting access to select agents and toxins; security risk assessments); 15 C.F.R. pt. 774 (Technology Export Controls on Instruments for the Automated Chemical Synthesis of Peptides). The Bureau of Industry and Security at the Department of Commerce has also relied on the Export Control Reform Act of 2018 to establish export controls on emerging biotechnologies.

419 National Institute of Standards and Technology, *Cybersecurity of Genomic Data*, NIST Interagency/Internal Report (IR) 8432 (December 2023), <https://doi.org/10.6028/NIST.IR.8432>; National Institute of Standards and Technology, *Genomic Data Cybersecurity and Privacy Frameworks Community Profile*, NIST IR 8467 2nd Public Draft (2PD) (December 2024), <https://doi.org/10.6028/NIST.IR.8467.2pd>.

420 DOE, National Nuclear Security Administration, "NNSA and Technology Transfer," accessed June 29, 2026, <https://www.energy.gov/nnsa/technology-transfer-1>; Advanced Research Projects Agency for Health, "ARPA-H BDF Toolbox," accessed June 29, 2026, <https://arpa-h.gov/explore-funding/programs/arpa-h-bdf-toolbox>.

421 42 U.S.C. § 247d–6b.

422 ASPR, "Public Health Emergency Medical Countermeasures Enterprise (PHEMCE)," US Department of Health and Human Services, accessed June 29, 2026, <https://aspr.hhs.gov/PHEMCE/Pages/default.aspx>.

423 Defense Production Act authorities may also be relevant for scaling secure compute, trusted infrastructure, and critical biosecurity-enabling production capabilities during crises.

424 ASPR deploys analogous National Disaster Medical System response teams that provide surge medical services during or after a disaster or public health emergency. ASPR, Administration for Strategic Preparedness and Response, *Justification of Estimates for Appropriations Committees, Fiscal Year 2024*.

ensure more predictable antitrust treatment and a limitation on damages, incentivizing broader participation.⁴²⁵

Legal Gaps

1. A bespoke NIBR charter and stable appropriations

While current law and authorities provide a legal and political foundation for much of NIBR's implementation, they do not expressly authorize the government to stand up the full concept as a durable organization with congressional backing. As NIBR's success depends on utilizing public-private partnerships creatively to build America's innovation and crisis response capacities, the lack of an explicit authorizing statute, and dedicated, long-term funding may limit the initiative's ability to adapt to a dynamic biological threat landscape.

2. Clear legal rules for sensitive biodata access and use

Existing privacy law, confidentiality protections, and agency-specific authorities were not designed for a model in which government and private actors gain broad, secure access to biological data via a centralized database; moreover, the application of these rules in training AI models is fragmented⁴²⁶ and the subject of several legislative proposals. The lack of a flexible IP framework may also limit the scope of research undertaken through NIBR, given restrictions on use of proprietary datasets.⁴²⁷ Greater clarity on the types of data NIBR may receive, what privacy protections and data security controls apply, and the legal safe harbors available for participants would create a more predictable framework for NIBR's implementation while ensuring appropriate safeguards remain in place.

3. Procurement and partnership authorities tailored to AI-bio infrastructure

While federal agencies have applied existing contracting and partnership authorities in analogous contexts, NIBR may need explicit authority for specific functions, such as rapidly scaling Genesis challenge programs or establishing specialized IP regimes for jointly produced models and biodata products. Congress could consider adapting models such as the Biomedical Advanced Research and Development Authority's flexible use of other transaction agreements for medical countermeasures to the AI and biotech context.⁴²⁸

4. Express authority for operational biological emergency teams

While NIBR's model of biological emergency teams is analogous to the DOE's Nuclear Emergency Support Teams (NEST), NEST relies on a mature body of statutes and specific presidential policies assigning DOE and NNSA radiological and nuclear response missions.⁴²⁹ There is currently no biological equivalent under federal law with the same clarity of mission, chain of command,

425 15 U.S.C. § 4301 et seq.; US Department of Justice, Antitrust Division, "Filing a Notification Under the NCRPA," last updated Apr. 1, 2025, <https://www.justice.gov/atr/filing-notification-under-ncrpa>.

426 Craig Konnoth, "AI and Data Protection Law in Health," in *Research Handbook on Health, AI and the Law*, ed. Timo Minssen and others (Edward Elgar Publishing, 2024), <https://www.ncbi.nlm.nih.gov/books/NBK613196/>.

427 National Science and Technology Council, *National Strategic Computing Reserve: A Blueprint*.

428 ASPR, "Other Transaction Agreements," US Department of Health and Human Services, accessed June 29, 2026, <https://aspr.hhs.gov/AboutASPR/ProgramOffices/BARDA/Pages/Other-Transaction-Agreements.aspx>.

429 FEMA, *NIRT Fact Sheet: Nuclear Incident Response Team (NIRT)* (September 2023), https://www.fema.gov/sites/default/files/documents/fema_oet-nirt-factsheet.pdf; DOE, *Nuclear Emergency Support Team (NEST)*, accessed June 29, 2026, <https://nsarchive2.gwu.edu/nukevault/ebb270/04.pdf>.

and standing readiness obligation.⁴³⁰ Moreover, the complexity of necessary support functions, including onsite diagnostics, reach-back modeling, regular training exercises, and potential federal-state coordination, would benefit from explicit legal codification.

430 For instance, DOE capabilities to prepare for and respond to a biological incident, many developed after the COVID-19 pandemic, focus on marshalling R&D capabilities or organizing internal capabilities and processes to coordinate the department's response. DOE, Office of Science, "National Virtual Biotechnology Laboratory (NVBL)," accessed June 29, 2026, <https://science.osti.gov/nvbl>; DOE, "US Government Investments Advance Biodefense and Pandemic Preparedness at the Department of Energy," accessed June 29, 2026, <https://www.energy.gov/science/articles/us-government-investments-advance-biodefense-and-pandemic-preparedness-department>; National Nuclear Security Administration, "Bioassurance Program," in US Government Investments Advance Biodefense and Pandemic Preparedness at the Department of Energy, <https://www.energy.gov/science/articles/us-government-investments-advance-biodefense-and-pandemic-preparedness-department>; DOE, Office of Environment, Health, Safety and Security, "DOE Biological Events Monitoring Team," accessed June 29, 2026, <https://www.energy.gov/ehss/emergency-response-resources>.