

Issue Brief

Anticipating Biological Risk



A Toolkit for Strategic
Biosecurity Policy

Author

Steph Batalis

Executive Summary

Biological threats from pathogens and toxins have the capacity to cause significant and widespread harm, regardless of whether they are natural or engineered, intentional or unintentional. Although recent news and policy discussions have focused on biological threats that are enhanced or enabled by artificial intelligence (AI), pathogens and toxins can already cause harm without this emerging technology. To counter biological threats regardless of source, policymakers need a range of governance tools and mitigation measures upon which to draw. The first step in building such a toolkit is to understand what the process resulting in biological harm looks like for various scenarios. Then a suite of policy options can be assembled to intervene at points throughout the process.

This report addresses two generalized scenarios that can result in biological harm. The first outlines the steps that a malicious actor may take to intentionally generate and deploy a pathogen or toxin (marked with a target icon throughout the report ). The second describes actions that could result in an unintentional laboratory accident during legitimate scientific research (marked with a test tube icon ). Each scenario involves completing a series of planning and physical stages, which offer multiple opportunities to build policy toolkits with a variety of mechanisms.

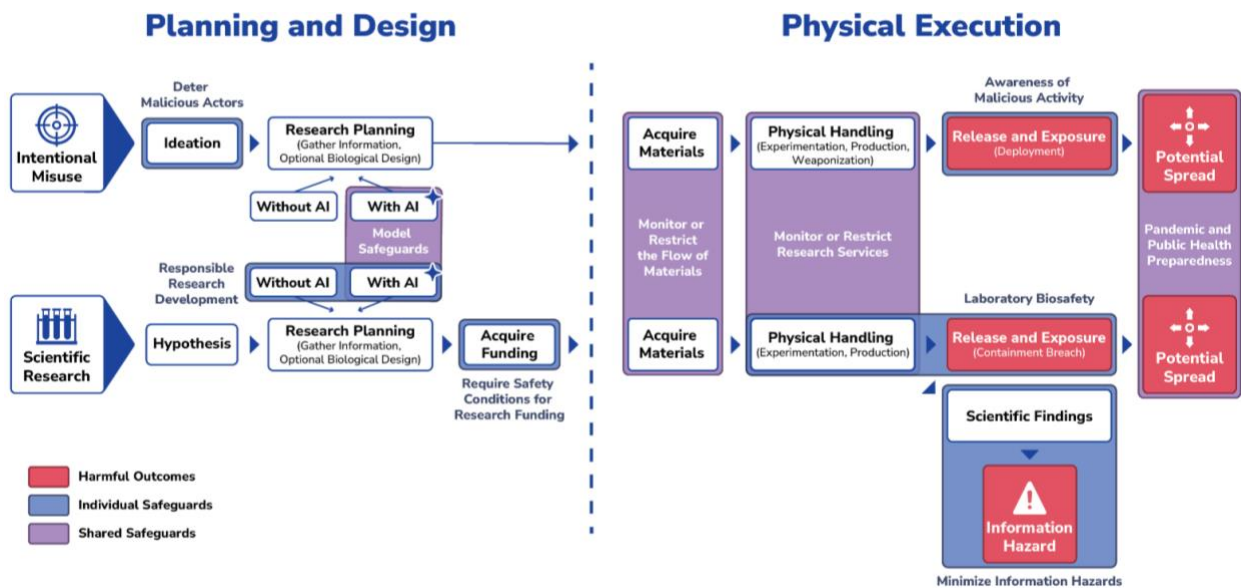
While describing each pathway, this report also takes stock of the many recommendations that have been made to strengthen U.S. biosecurity and biodefense, and maps these to the steps and scenarios where they would apply. In doing so, we identify governance gaps, along with opportunities to address them by implementing new safeguards or improving existing ones. Key findings include:

- Some safeguards apply to both intentional misuse and legitimate scientific research, while others are specific to one scenario.
- Research oversight mechanisms are primarily leveraged against federally funded research, leaving both regulatory and visibility gaps for non–federally funded research.
- Biosafety, biosecurity, and biodefense responsibilities span government missions, departments, and agencies and require increased coordination.
- Effective oversight would benefit from a biological risk framework that clearly and specifically defines concerning outcomes.

The options presented throughout this report are part of a comprehensive toolkit that policymakers can apply across the entire biological risk pathway. These solutions are

designed to mitigate biological harm from a variety of sources, including the AI-relevant concerns that are the focus of considerable current attention. By layering safeguards across multiple steps, policymakers can apply the tools described in this report to more effectively address both AI-enhanced and AI-agnostic threats without unduly hindering scientific innovation.

Figure A: Simplified View of Pathways that Result in Biological Harm from a Pathogen or Toxin*

















Source: CSET.

Note: Dashed blue line denotes the Planning-to-Physical transition, and red boxes denote harmful outcomes.

* Figure A in the executive summary is the same as Figure 4 below.

Table of Contents

Executive Summary.....	1
Introduction.....	4
A Policy Toolkit for Two Biorisk Pathways	6
Planning and Design and Physical Execution Phases	6
Phase 1: Planning and Design.....	8
Step: Ideation  or Hypothesis Generation 	8
Step: Research Planning (Information-Gathering, Optional Biological Design)  	11
Step: Acquire Research Funding 	17
Phase 2: Physical Execution.....	18
Step: Acquire Materials  	19
Step: Physical Handling (Experimentation and Production  , Weaponization )	22
Outcome: Release and Exposure (Intentional Deployment  or Accidental Containment Breach ).....	24
Outcome: Potential Spread  	26
Outcome: Scientific Findings and Potential Information Hazard 	27
Concluding Thoughts	29
Author.....	31
Acknowledgments.....	31
Endnotes.....	32

Introduction

Pathogens and toxins are at the center of some of the most pressing biodefense concerns. The threat landscape for biological harms from pathogens and toxins is large and diverse, ranging from naturally occurring diseases that pass from animals to humans, to the increasing incidence of antimicrobial-resistant infections, to accidental biological incidents such as laboratory accidents due to human or mechanical errors, to deliberate biological threats from malicious actors.¹

Recently, the concern has expanded to include biological threats that are intentionally or unintentionally enabled, accelerated, or otherwise augmented by artificial intelligence.² References to AI-enhanced biological threats have made their way into both public and policy discussions, including concerns that AI tools could help a potential bioterrorist to make their plans of misuse or to design more specific, targeted, or dangerous pathogens or toxins.³ Deciding how to govern these AI tools is particularly challenging, given that the same tools have already proven remarkably useful for legitimate scientific research, and are likely to contribute to future scientific, biomedical, public health, and environmental advances.

Policymakers have a range of options in their governance toolkits to combat foreseeable biological harms. These are not just “AI solutions,” but rather span the multistep process that results in biological harm. By layering safeguards across multiple steps, policymakers can more effectively mitigate both AI-enhanced and AI-agnostic threats. For example, AI regulation is receiving considerable current attention but would only prevent certain types of biological risks. While model safeguards should be pursued, they should be viewed as one tool in a larger safety toolkit. A multilayered approach also benefits from being able to draw upon a robust community of biosafety and biosecurity experts, who have already conducted decades of research on gaps and suggestions for U.S. biodefense.⁴



To implement layered safeguards, policymakers should first understand the multistep process that results in biological harm and which interventions could mitigate risk at various points. This report outlines these steps for two scenarios that involve using or generating pathogens or toxins: (1) malicious actors intentionally misusing biology; (2) laboratory accidents during legitimate scientific research that result in the unintentional release of a biological agent. These two scenarios are the focus of this report because they have been at the center of recent attention, especially in relation to concerns about the possibility for AI to enhance harmful outcomes.

This report also presents a suite of potential options to safeguard against accidental or intentional harm from pathogens and toxins, mapped to the most applicable step or steps in the pathway. These options include mitigation measures that have been previously suggested by various experts, along with considerations and challenges that policymakers should keep in mind while deciding which, if any, tools to leverage.

Mapping each scenario and its corresponding policy levers reveals the following key takeaways:

- **Some safeguards apply to both intentional misuse and legitimate scientific research, while others are specific to one scenario.** It will be important to ensure that future efforts identify the scenario they are trying to target in order to apply the correct policy tool. Future efforts could further assess how intentional misuse differs between state- and non-state actors, or which safeguards apply to domestic versus foreign actors.
- **Research oversight mechanisms are primarily leveraged against federally funded research, leaving both regulatory and visibility gaps for non-federally funded research.** Most of the research oversight discussed here uses federal research funding as the enforcement lever for safety or federal reporting measures. This leaves out research conducted without federal funding that could still result in harmful or unintended outcomes. New oversight would be needed to monitor or regulate this type of research.
- **Biosafety, biosecurity, and biodefense responsibilities span government missions, departments, and agencies and require increased coordination.** As the steps and policy interventions throughout the biological risk pipeline demonstrate, a whole-of-government approach is needed to effectively safeguard against potential harms while promoting scientific research and innovation. The currently fragmented U.S. biodefense ecosystem raises regulatory and information-sharing challenges that will require increased coordination, strategy, and integration among agencies to overcome.⁵
- **Effective oversight would benefit from a biological risk framework that clearly and specifically defines concerning outcomes.** While biological risk is frequently cited as a concern for new technologies such as AI, there is no comprehensive understanding of what exactly is considered a risk and which specific outcomes are considered to be concerning. The lack of a clearly interpretable biological risk framework hinders the ability to prioritize threats and to design and test their corresponding safeguards.

A Policy Toolkit for Two Biorisk Pathways

This report will discuss the general, simplified steps for two pathways: a malicious actor who intentionally generates pathogens and toxins for misuse (marked with a target icon throughout this report ), and laboratory accidents that occur during legitimate scientific research (marked with a test tube icon ). Policy options are presented alongside the step or steps that they would be most likely to affect, all of which serve different purposes and act through various mechanisms.

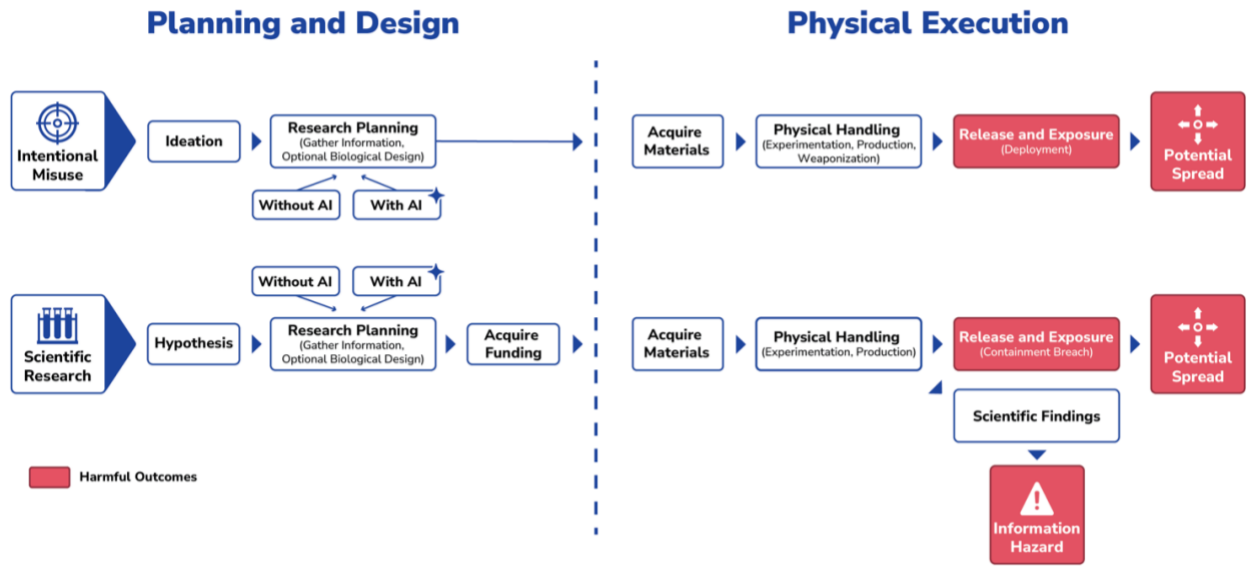
Note that while this report depicts these steps within a linear pathway, they can backtrack, repeat, or occur in a different order in real life. For example, research often involves troubleshooting experiments and iterating upon protocols, methods, and materials. These variations are not shown in the figures of this report in the interest of simplicity, but the associated risk mitigation strategies are applicable regardless of the steps' order.

Planning and Design and Physical Execution Phases

The pathway to generate a pathogen or toxin is divided into two broad phases, Planning and Design as well as Physical Execution (Figure 1). In the Planning and Design phase, users generate a design for their desired product, and a detailed step-by-step plan for creating it. In the Physical Execution phase, users bring the plan out of the computer screen or lab notebook and into the real world by physically producing the pathogen or toxin.*

* In discussing the impact of AI on biosecurity, this is often called the Digital-to-Physical Transition because it describes the phase at which a computer prediction becomes a physical object.

Figure 1: Simplified View of Pathways that Result in Biological Harm from a Pathogen or Toxin



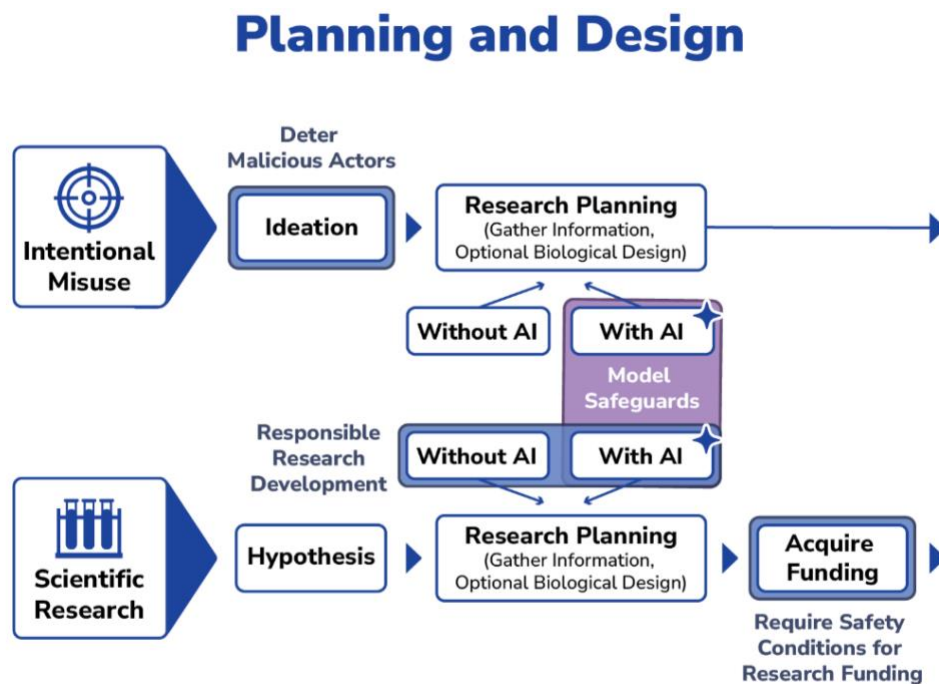
Source: CSET.

Note: Dashed blue line denotes the Planning-to-Physical transition, and red boxes denote harmful outcomes.

Phase 1: Planning and Design

The Planning and Design phase drafts the overall experimental and operational protocol before it is physically set into motion. It generally includes conceptualization, gathering information, and designing the experimental research plan. These steps do not yet incorporate the materials or physical actions that result in real-world infections. The general steps in the Planning and Design phase, and the policy options that could impact them, are shown in Figure 2.

Figure 2. Steps and Policy Options in the Planning and Design Phase




Source: CSET.

Note: Blue boxes denote safeguards that are specific to intentional misuse or scientific research, while purple boxes denote safeguards that apply to both scenarios.

Step: Ideation or Hypothesis Generation

Both malicious actors and well-intentioned scientists begin with an idea that guides subsequent steps in the pathway. For a malicious actor, this “ideation” step could be as general as deciding to pursue a biological weapon in the first place, or as specific as selecting a bioweapon type and target. For a researcher, “hypothesis generation”

describes making a testable, educated guess informed by previous research experience, specific scientific goals, or open questions in the scientist's field of interest.

- **Policy Option: Enhance biosurveillance and bioattribution to deter malicious actors.**  Malicious actors may be less likely to pursue biological weapons if they think they are unlikely to be successful or likely to be caught.

Demonstrating effective biosurveillance and bioattribution capabilities may increase the perception of these barriers: biosurveillance by rapidly detecting potentially disease-causing agents for mitigation before they can cause widespread harm, and bioattribution by determining the source of a biological agent, whether natural or engineered. Enhancing deterrence through biosurveillance and bioattribution could include the following actions:

- **Improve and expand U.S. biosurveillance.** At present, experts note that U.S. biosurveillance is hampered by fragmented federal responsibilities and the need for more effective programs and detection technologies.⁶ Experts point to a need for the United States to increase federal coordination, develop strategies and infrastructure for data collection and sharing among jurisdictions and agencies, and promote the development of pathogen-agnostic detection measures.⁷
- **Strengthen biological attribution to deter malicious actors.** Resources that could improve the technical capacity to attribute biological agents include more robust scientific tools, methods, and reference databases, along with an expanded bioattribution workforce with the skills to use them.⁸ Other opportunities include establishing multilateral agreements for access to samples and associated data, implementing a U.S. national plan for bioattribution that clarifies departmental and agency roles and responsibilities, and increasing support for international efforts such as the UN Secretary-General's Mechanism for Investigation of Alleged Use of Chemical and Biological Weapons.⁹

Box 1. Federal Research Oversight Mechanisms

In the United States, a range of federal oversight mechanisms and regulations for legitimate biological research apply throughout the research cycle. With the exception of the Federal Select Agents Program, each of these mechanisms applies only to federally funded research.

- The [NIH Guidelines for Research Involving Recombinant or Synthetic Nucleic Acid Molecules \(NIH Guidelines\)](#), issued by the **National Institutes of Health**, direct Institutional Biosafety Committees (IBCs) to review research proposals, conduct risk assessments, and determine containment procedures before researchers can initiate work involving recombinant or synthetic nucleic acid molecules, or organisms and viruses containing these molecules.
- The [United States Government Policy for Oversight of Dual Use Research of Concern and Pathogens with Enhanced Pandemic Potential \(DURC/PEPP Policy\)](#), issued by the **Office of Science and Technology Policy (OSTP)**, addresses risks related to these areas. The DURC/PEPP Policy requires investigators and their institution's Institutional Review Entities (IREs) to develop risk-benefit assessments and risk mitigation plans, and for the federal funding agency or department to review the potential risks and benefits of the proposed research to guide funding decisions.
- The [Federal Select Agent Program \(FSAP\)](#), jointly managed by the **Centers for Disease Control and Prevention (CDC)** within the U.S. Department of Health and Human Services (HHS) and the **Animal and Plant Health Inspection Service (APHIS)** within the U.S. Department of Agriculture (USDA), regulates the possession, use, and transfer of certain high-risk pathogens and toxins, called Biological Select Agents and Toxins (BSAT). FSAP registration is required for all entities in the United States that possess, use, or transfer any BSAT, and is not tied to federal research funding.
- The [Executive Order 14110 on Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence \(2023 EO on AI\)](#), issued by the **Biden Administration White House**, and the accompanying [Framework for Nucleic Acid Synthesis Screening \(OSTP Screening Framework\)](#), issued by **OSTP**, requires commercial synthetic nucleic acid providers to attest to screening orders to determine whether the requested sequence poses potential risk, and whether the customer has a legitimate use for it, for federally funded research projects.
- The [Biosafety in Microbiological and Biomedical Laboratories \(BMBL\)](#), issued by the **CDC** and **NIH**, describes biological risk groups and biosafety levels and establishes general biosafety practices, procedures, and equipment to guide the proper handling of infectious microorganisms and hazardous biological materials. The BMBL contains best practices to use as guidance but is not a regulatory document.

Step: Research Planning (Information-Gathering, Optional Biological Design)


Once an actor or group has decided to generate a pathogen or toxin, the next step is to find or create a detailed, step-by-step protocol to guide future physical action. To make such a plan, an individual needs to understand scientific concepts, find relevant laboratory techniques and experimental protocols, and identify life sciences material providers. Much of this information could inform both malicious and legitimate outcomes, because the underlying scientific processes are the same regardless of the intended use.

Scientific information can be accessed from a range of sources, including the researcher's preexisting scientific expertise, non-AI resources on the internet or at a library, and AI chatbots such as ChatGPT.¹⁰ A chatbot may be especially appealing to non-experts due to its perceived ease of use, conversational tone, and simple user interface. However, much of the information it offers is already widely available without that technology. Online resources for a range of expertise levels include detailed laboratory guides and protocols, DIYbio educational materials, YouTube videos, and many others.¹¹ Scientific publications are particularly rich sources of information because they are meant to be replicable and testable, and so include details such as scientific rationales, step-by-step protocols, and comprehensive lists of specific materials and where to obtain them. In addition to scientific concepts, malicious actors may also research logistical and operational information at this stage to inform their future plans.

Some plans involve additional biological design steps to design or modify a pathogen or toxin to have new characteristics. Researchers have studied pathogens and toxins this way for decades, for example to learn more about how a pathogen functions or to design therapies such as vaccines.* Malicious actors may use the same techniques to design or redesign elements of a pathogen or toxin to be more harmful, easier to

* Researchers can design protocols using longstanding, decades-old techniques to modify pathogens and toxins, including genetic engineering, serial passaging, reverse genetics, and pathogen recombination. More recently, non-AI computational modeling software has accelerated researchers' ability to design and predict specific modifications and outcomes. For more information on how and why researchers conduct gain and loss-of-function research on pathogens, see: Caroline Schuerger, Steph Batalis, Katherine Quinn, et al., "Understanding the Global Gain-of-Function Research Landscape," Center for Security and Emerging Technology, August 2023, <https://cset.georgetown.edu/publication/understanding-the-global-gain-of-function-research-landscape/>.

produce, or evade detection. Regardless of intent, one way to design pathogens and toxins uses the traditional scientific method and involves forming a hypothesis about which biological elements to modify, and how, to achieve a desired goal. Alternatively, AI-enabled biological tools, like Biological Design Tools (BDTs), predict biological and molecular features guided by user-defined criteria.* Users may find it beneficial to combine these methods, for example by drawing on their preexisting scientific expertise to select a biological element to modify, and then taking advantage of AI-enabled biological tools to predict specific alterations that will result in the desired outcome.

- **Policy Option: Employ model safeguards to govern how AI tools are developed, deployed, used, and monitored.**  While AI tools are not required to generate a pathogen or toxin, they can be incorporated into the Planning and Design phase and are thus one area toward which policy action could be directed. Potential safeguards span the entire AI lifecycle, and generally involve governing how certain models should be built and who should be able to develop them, what capabilities these models should or should not have, and which users should have access to them and how. The 2023 Executive Order (EO) on AI will inform some of this oversight by requiring studies to better understand biological risks, and by mandating reporting requirements for certain types of models (see Box 1 for details). However, model regulations are accompanied by technical, philosophical, and logistical questions that will require careful consideration.

One challenge surrounds how to define which outcomes are considered “risky,” and thus must be assessed and mitigated. Assessing biological risk is difficult because it is context-dependent, does not include clear borders between “risky” and “harmless” pathogens, and necessitates making predictions about how pathogens function under different conditions. Assessing whether AI models generate potentially risky biological information, then, is similarly difficult to determine.

* Biological design tools (BDTs) design elements of biological systems, often implying a generative function. However, some biological models are not generative and may not be used for design purposes, but are still relevant to discussions of AI x Bio. For example, a large language model (LLM) trained on DNA sequences or an AI tool that predicts protein structures can be integrated into experimental workflows without generating or designing molecules. To avoid overly restrictive terminology, models trained on biological data are broadly referred to as “AI-enabled biological tools” throughout this report.

An additional consideration is that some types of model safeguards can only be achieved for closed-source models in which developers maintain complete control and surveillance over their systems and users cannot alter, fine-tune, or covertly use them.¹² Many AI-enabled biological tools are already available open source because they were developed by life sciences academic groups, a community that strongly promotes the concept of open science and transparency for peer review and that often does not have adequate resources to host and maintain an application programming interface (API).¹³ If closed-source models become required, these groups may need additional infrastructure, people, or funding to host and maintain secure interfaces.

Finally, it will be important to carefully balance the tradeoffs between safety and security measures and their potential impacts on model performance, capabilities, and beneficial applications. In some cases, it may be acceptable to decrease model performance in pursuit of safety, for example by limiting chatbots from providing potentially dual-use information about historical bioweapons attacks, public health vulnerabilities, and lists of the most dangerous human pathogens. On the other hand, the capabilities that make an AI-enabled biological tool useful for scientific research—the ability to better understand, manipulate, and design biological systems—are the same ones that could be exploited to cause harm. Because these outcomes are technically challenging to separate, safeguards that prevent an AI-enabled biological tool from providing harmful outputs could impede the meaningful scientific advances that these tools were designed to achieve. Such an outcome could widen resource gaps for less-funded researchers who use AI-enabled biological tools to reduce the need for certain expensive equipment and experiments.

The following options and considerations have been proposed as potential bio-relevant AI safeguards, spanning the AI lifecycle from development to use.

During model development:

- **Increase developer awareness of biosafety and biosecurity principles.** Developers can more responsibly design, use, and disseminate biological models if they are aware that biosecurity risks exist in the first place, and know which outcomes to evaluate for and build safeguards against.¹⁴ The United States could support the development of a developer-focused biosafety and biosecurity training module, which could be voluntarily adopted by developers or made a condition of federal research funding or access to other resources such as databases or compute infrastructure.

- **Restrict how chatbots engage with certain topics.** Methods to establish off-limits biological topics for chatbots include filtering training data (as elaborated in the following section), discouraging or penalizing certain outputs, and building additional filters to flag problematic questions and/or block the model from responding to the user.¹⁵ Notably, deciding which topics to restrict, and to what level of detail, is a challenging question that may be approached asymmetrically by developers with different risk tolerances. While guardrails could be removed or circumvented by particularly committed actors through jailbreaking and creative prompt engineering, they may still deter less-motivated users or those who are unaware that they are requesting harmful information. Initial evidence from a red-teaming exercise to plan a hypothetical biological attack indicates that model guardrails require additional time and effort to overcome.¹⁶
- **Identify and remove certain dual-use data from model training.** Models may be less capable of generating harmful outputs if they are not trained on dual-use data. As discussed above, defining which types of data are risky enough to restrict is a challenge because dual-use data have legitimate uses alongside their potentially harmful ones. For a chatbot, this could be information about historical bioterrorism attacks or the highest-consequence public health threats, while for AI-enabled biological tools this could include data such as pathogen genomes or chemical toxicity profiles.¹⁷ Additionally, users can update open-source models with training data that was initially excluded, and developers may be disincentivized from removing training data if it decreases model performance.¹⁸
- **Implement access controls to limit who is able to develop certain models.** Customer screening for sensitive biological datasets and high-performance computing infrastructure, using the measures described in Box 2, is one option that has been proposed to limit model development to approved developers and increase visibility into development pipelines.¹⁹ However, limiting computing access may be less effective as smaller domain-specific models become increasingly capable, and as technical advances lower the needed computational power to develop cutting-edge models.²⁰

- **Assess models before they are released to identify, characterize, and measure potentially harmful outputs.** Assessments can provide insight into whether and how often undesired outputs occur, how well safeguards are working, and if safeguards need to be updated or improved. For example, red-teaming tests whether users are able to break through guardrails in a controlled environment.²¹ But as of October 2024, no standardized evaluation or assessment guidelines define what exactly is considered a biological risk, which specific factors increase risk, or which factors substantiate reporting. Other types of biorisk governance, such as research oversight frameworks or export control lists, may be a useful starting point to identify biological risk factors of concern.

During model use:

- **Implement model access controls.** Models with dual-use potential could be restricted by user log-ins, credentialing, customer screening, structured access, or other forms of “know-your-customer” measures such as those described in Box 2.²² However, deciding which types of AI-enabled biological tools are deserving of access controls is an open question. Some tools have more obvious dual-use potential, for example those that predict a pathogen’s transmissibility or ability to evade the immune system. However, more generalized AI-enabled biological tools, such as those that generate protein structures or predict molecular interactions, could also be misapplied to cause harm.
- **Monitor user behavior and model usage.** Tracking user behavior could help to identify users who input potentially concerning queries, recognize unusual patterns of use within a single model or across multiple models, and aid in attribution efforts if misuse does occur. If models are hosted through structured-access APIs, user behavior could be monitored by maintaining and reviewing logs of prompt queries or building systems to detect and alert certain types of usage.²³
- **Establish harm reporting mechanisms.** Even if careful consideration goes into anticipating potential negative outputs, there is always a possibility that AI systems could generate other, unforeseen consequences. Reporting mechanisms, which could include one or a combination of mandatory, voluntary, and citizen reporting, would allow

developers or other oversight bodies to track and characterize outcomes that may otherwise go unnoticed.²⁴


- **Conduct continuous research on evolving risks and safeguards.** As the state of the art continues to advance, ongoing monitoring, testing, and horizon-scanning of new model capabilities can ensure that safety and security approaches keep pace with these changes. A number of mechanisms to promote this research have been suggested, including creating federal research funding initiatives, forming U.S. consortia or working groups, and convening international forums.²⁵

Box 2. Customer Oversight: Monitoring and Screening

Also called Know-Your-Customer guidelines, customer oversight monitors or screens the people who use a product, tool, or service. Monitoring can help to detect unusual patterns of use or aid in attribution capabilities, for example by requiring a login to access a system that can then collect and store information about users and their activities.²⁶ Screening makes use of that same information to determine whether an individual should be allowed to complete a purchase or access a system.²⁷

Customer oversight generally involves three factors: verifying a customer's identity (is this person who they say they are, and do they actually exist?), whether a user is legitimate (does the customer have the appropriate credentials to access this resource?) and whether the intended use case is legitimate (is there a valid need for the resource?). All three of these elements can be difficult, costly, time-consuming, and subject to individual perceptions of legitimacy.²⁸ For example, verifying a researcher's institutional affiliation is challenging at scale because of the sheer number of research institutions in the world, the lack of name recognition (and associated perception of legitimacy) associated with new startups and small international institutions, and the frequency with which researchers move between institutions.²⁹


Various forms of customer oversight have been proposed for several resources within the life sciences, from databases and AI models to physical and virtual infrastructure, materials, and services. However, experts also recognize challenges, for example the lack of guidance for providers about what constitutes "concerning" or "unusual" use and the need to consider equity, access, and fairness to avoid discouraging public engagement in the life sciences.³⁰ Some have further suggested using a centralized life sciences customer verification system to reduce screening burdens for individual providers, and to enable better detection of concerning patterns by tracking a single customer's activities across multiple services and providers.³¹

- **Policy Option: Encourage responsible research development by raising researcher awareness and training for biosafety and biosecurity.** 

Researchers can play an important role in risk reduction during legitimate scientific research by designing safer research plans that take biosafety and biosecurity into consideration. At present, both the NIH Guidelines and DURC/PEPP Policy require researchers and their institutions to assess safety concerns and develop risk mitigation plans before any experiments can commence (see Box 1 for details). Additional strategies could include engaging with research communities to set norms and empower responsible research, developing and incorporating biosafety and biosecurity training into government-led biotechnology initiatives and federal funding requirements, and generally increasing such training for researchers at every level, from lead investigators to trainees.³²

Step: Acquire Research Funding

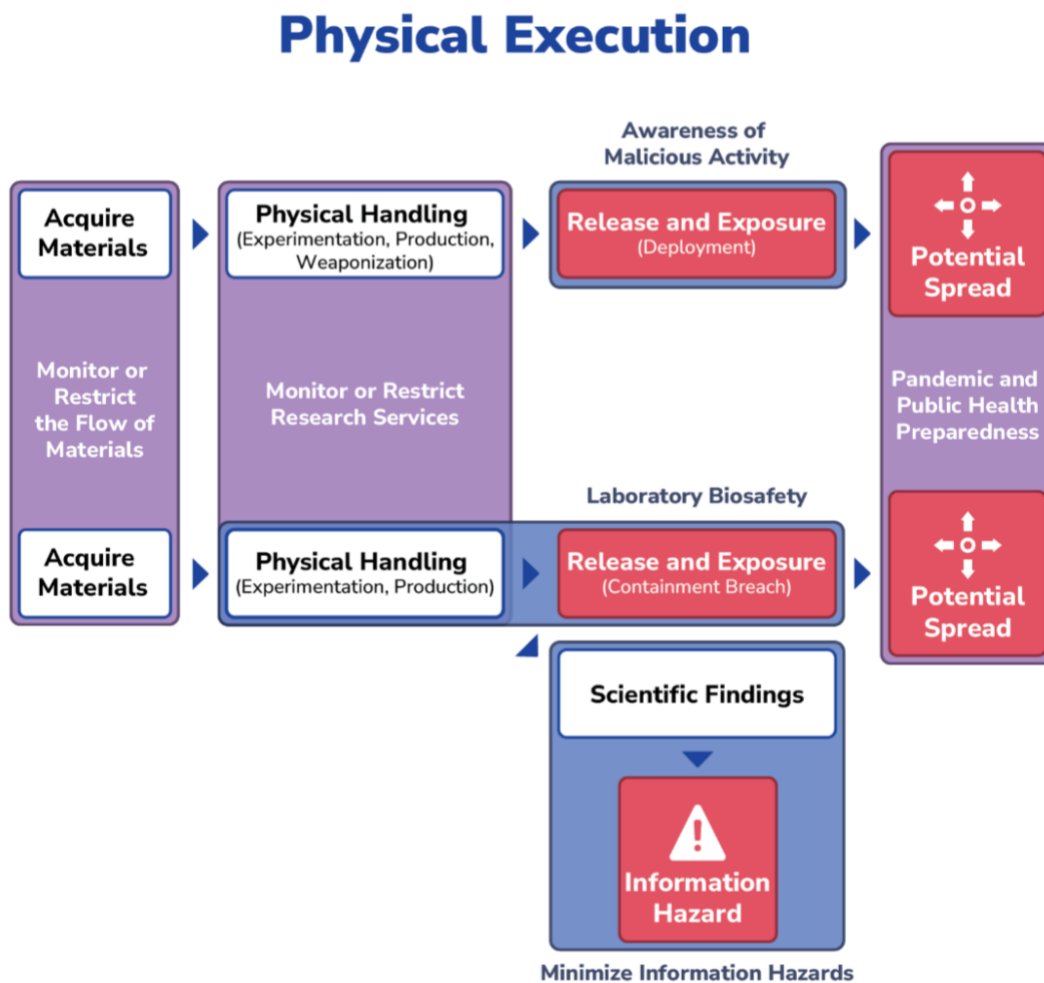
While private companies may fund their own research and development, researchers at academic or nonprofit institutions typically apply for grants to fund their work. In the United States, a large portion of research grant funding comes from the federal government through agencies such as the NIH or the National Science Foundation, although nonprofits, companies, philanthropic, and other organizations also fund research.

- **Policy Option: Include safety and reporting requirements as conditions of research funding.**  Research funding grants can be a major point of oversight because they require researchers to disclose their intended research activities and demonstrate that they are able to meet safety requirements. The NIH Guidelines, DURC/PEPP Policy, and 2023 EO on AI's OSTP Screening Framework all use federal grant funding as the policy lever to require various biosafety and biosecurity measures (see Box 1 for details). Beyond these policies, individual sources of federal research funding can include additional terms and conditions including, for example, requiring adherence to the BMBL.³³ These levers are limited to federally funded research and leave gaps in visibility and oversight for privately funded research. An additional measure could be for non-federal funders to incorporate biosecurity and biosafety reviews. Voluntary efforts and partnerships, for example the newly announced International Bio Funders Compact, could help non-federal funders to incorporate biosafety and biosecurity reviews into their funding approval and oversight processes.³⁴

Phase 2: Physical Execution

Once an actor has decided on a plan to generate a pathogen or toxin, the next series of steps involves translating that plan into a physical product. This phase is where the potential for real-world infections, with physical pathogens and toxins, can arise. Figure 3 illustrates the steps in the Physical Execution phase and their corresponding policy options, with harmful outcomes indicated in red.

Figure 3. Steps and Policy Options in the Physical Execution Phase



Source: CSET. Red boxes denote harmful outcomes.

Note Blue boxes denote safeguards that are specific to intentional misuse or scientific research, while purple boxes indicate safeguards that apply to both scenarios.

The Physical Execution phase can vary widely in terms of the degree of difficulty, accessibility, and cost to complete. Different types of biological agents—like bacteria, viruses, fungi, and toxins—are handled in specialized ways, using different materials, techniques, and resources. Variability can also stem from differing objectives. A researcher whose goal is to achieve replicable results and draw accurate conclusions may place a high importance on using high-quality, validated materials and doing things the “right way,” while a malicious actor may accept less precise but lower-cost alternatives. Because this brief cannot address every possible scenario, it instead describes general required elements within each step of the Physical Execution phase.

Step: Acquire Materials

Regardless of their intent and final desired outcome, actors will need to procure the materials that are relevant to their plan. Examples of possible materials include equipment, laboratory supplies, and biological materials.


- **Equipment:** can include storage, experimentation, or containment equipment (e.g., refrigerators, centrifuges, incubators, biosafety cabinets, and autoclaves). Some types of equipment could help a malicious actor to disseminate pathogens or toxins and are included on regulatory lists of dual-use items, although they also have legitimate uses (e.g., freeze dryers or sprayers).³⁵
- **Lab Supplies:** can include consumables and reagents (e.g., glassware, paper and plastic products, chemicals, and cell culture media).
- **Biological Materials:** can include infectious or non-infectious samples, nucleic acids, and research animals (e.g., bacteria or viruses, cell lines for experimentation or protein production, DNA or RNA, and laboratory mice).

Researchers can access these materials by purchasing them from major suppliers of scientific products, peer-to-peer sharing between labs, or via shared equipment facilities at their institutions.

It is also possible to set up a functioning laboratory without access to these sources. For example, some online providers cater to customers who set up laboratories for other reasons, including high school or community college courses, science education programs, and citizen scientists.* Resellers on eBay and other auction sites often list

* Biotech companies that ship to residential addresses sell polymerase chain reaction (PCR) kits and machines, incubators, ultra-low temperature freezers, genetic engineering kits, protein expression and purification reagents, assay kits, and benchtop nucleic acid synthesizers, among other materials.

used equipment from academic or commercial research laboratories, including benchtop DNA synthesizers and sequencers, full biosafety cabinets, and chromatography equipment. Researchers without large equipment budgets may also look for discounted equipment from these sellers. Community Bio and DIYbio groups also provide a wealth of information on how to set up a home lab, including by using low-cost at-home alternatives.³⁶ For example, home scientists can source specialized chemicals from hardware, pet supply, or craft stores; 3D-print a centrifuge attachment for a Dremel (a DremelFuge); or use an egg incubator, yogurt maker, or diaper warmer as a bacterial incubator.³⁷ Many of these options are especially useful for science outreach and education initiatives that aim to inspire youths to pursue science through hands-on experiences.

- **Policy Option: Monitor or restrict the flow of certain materials to reduce the chance of their misuse.**  Certain materials have higher levels of baseline risk or an increased dual-use potential. Tracking and restricting the flow of these materials is one way to safeguard against their intentional or unintentional misuse. In general, oversight mechanisms can monitor the flow of materials, screen customers, and create regulations to guide who should be allowed to access certain materials and under what conditions. In the United States, the Federal Select Agent Program (see Box 1 for details) and export controls both regulate the possession, use, and transfer of specific high-risk materials, pathogens, and toxins.³⁸ Individual research institutions may also enact processes for accepting and transferring materials to ensure compliance with federal regulations. For synthesized nucleic acids, export controls regulate the global distribution of certain sequences derived from high-risk organisms.³⁹ In addition, several nucleic acid synthesis providers voluntarily screen orders, and the 2023 EO on AI requires federally funded research to purchase nucleic acids from providers that attest to following the OSTP Screening Framework.⁴⁰ Additional measures could include the following actions:

- **Address implementation challenges for nucleic acid synthesis screening.** Sequences of concern (SOCs) can include gene sequences from regulated agents and toxins and other sequences that may contribute to a pathogen or toxin's ability to cause harm.⁴¹ Current methods screen for SOCs from relatively small lists of regulated pathogens, while the OSTP Screening Framework will require providers that attest to following the guidance to expand the sequences they screen for beyond this list by 2026.⁴² Expanding the list and deciding which additional sequences should be considered SOCs, however,

remains a challenge because a sequence's potential risk changes depending on how it is used and because the function of many sequences is unknown.⁴³ A public-private partnership between the National Institute of Standards and Technology (NIST) and the Engineering Biology Research Consortium (EBRC) is currently engaging with industry and public partners to lay out actionable next steps.⁴⁴ Notably, this is an area where AI tools could help. Projects to develop sequence-to-function algorithms that predict “harmful” and “not harmful” sequences, for example the Fun GCAT Program from the Intelligence Advanced Research Project Activity (IARPA), could help providers to flag additional SOCs during screening.⁴⁵ However, questions remain around how harmful a predicted sequence needs to be to trigger a screening flag, and how confident that prediction needs to be. Regardless of how screening frameworks are designed in the near term, long-term success will require them to be adaptable as technology evolves, potentially by including ongoing research on risks, new technical approaches, and iterative stress-testing protocols.⁴⁶ It will also be important to ensure that screening is part of a broader oversight landscape, and to not overly rely on nucleic acid synthesis screening to prevent all potential biological risks.*

- **Expand nucleic acid synthesis screening requirements beyond federally funded research.** The 2023 EO on AI does not require providers to screen orders, but instead only allows federal research funds to be used to purchase from those that do.⁴⁷ Providers may choose not to implement and attest to screening, if they accept forgoing customers who receive federal research funding. Similarly, privately funded research projects or malicious actors may choose to purchase synthesized nucleic acids from a provider that does not screen orders for SOCs. Experts have recognized this challenge, and suggested various international initiatives to set norms and establish baseline practices for universal gene synthesis screening.⁴⁸
- **Expand screening to other types of materials.** In addition to synthesized nucleic acids, customer screening such as that described in Box 2 has

* In addition to commercial synthesis providers, custom DNA can also be obtained from a benchtop synthesizer or, in some cases, by laboratory techniques. Other sequences can be obtained from plasmid repositories or isolated from samples. Additionally, not every method to generate a pathogen uses DNA or RNA, whether synthesized commercially or otherwise. For example, serial passaging can alter a pathogen's transmissibility, virulence, and host range.

been suggested to regulate other types of materials, for example samples of certain pathogens and toxins, nucleic acids from other sources, plasmids or helper viruses to generate viral particles, and equipment such as benchtop DNA synthesizers, in addition to many others.⁴⁹

- **Address resource-sharing practices among researchers.** Sharing resources is a common practice among researchers, government, and industry, and could result in potentially harmful materials being distributed to third parties without sufficient oversight. If safety measures involve monitoring or restricting access to materials, then visibility and oversight would need to extend to all end users and not just the customer who makes the initial purchase. Some experts suggest strengthening resource-sharing frameworks and engaging with the research community to encourage best practices and security norms for resource-sharing.⁵⁰

Step: Physical Handling (Experimentation and Production , Weaponization)

This step encompasses all of the hands-on, laboratory-based physical actions to culture, expand, manipulate, prepare, or otherwise use pathogens or toxins. For a malicious actor, this step may involve generating a large amount of the desired pathogen or toxin. Depending on the intended use case, malicious actors may also prepare the biological agent for use as a weapon by stabilizing it as necessary for later steps, formulating it with other components, or loading it into a dispersal device.


For a researcher, the goal may be to assemble a biological design, test a system, or measure specific biological characteristics. Experimentation is especially likely to loop back to previous phases in the process, as troubleshooting a failed experiment may involve identifying alternative protocols and new techniques, redesigning a biomolecule, or ordering different materials.

Many researchers perform physical handling steps themselves in purpose-designed laboratories at academic, medical, or private-sector research facilities. It is also possible to perform these steps in functional home laboratories outside of these institutions, or at rented space in a community laboratory. Additionally, life sciences service providers such as academic core laboratories, contract research organizations, biofoundries, and cloud laboratories can help to perform or augment some physical handling steps. These services can make research more efficient by outsourcing certain



steps to specialized facilities with dedicated equipment, expertise, and workflows. For example, there have been recent calls to establish a national network of cloud labs because they could minimize resource gaps for researchers without the budget for expensive equipment, increase the pace of potentially groundbreaking wet-lab experimentation, and automate large-scale biological data collection.⁵¹

Advances in automation and autonomous experimental platforms are emerging as another resource to assist with physical handling steps. These systems follow programmed instructions to move samples between various pieces of laboratory equipment using robotic arms and liquid handlers. Automated platforms can help researchers to increase productivity by running many samples in parallel, standardizing certain types of measurements and data collection, and freeing up hands-on time for other pursuits.

Recent examples demonstrated the ability for large language model (LLM)-based “agents” to control automated experimental platforms for chemical synthesis, and have raised concern about their potential to lower capability barriers for malicious actors.⁵² At present, these concerns may be more realistic for chemical synthesis tasks than biological engineering tasks that involve a greater variety of steps and the added complexity of maintaining living cells.⁵³ Mitigating future misuse risks for LLM-controlled experimental platforms can draw on policy options presented previously in this report. For example, model safeguards would apply to the AI component of such a system, while measures that restrict or monitor material acquisition would apply to the system’s connected laboratory equipment and required laboratory supplies.

- **Policy Option: Monitor or restrict access to research services.**  Current safety measures for life sciences service providers are piecemeal and often voluntary. Given the potential benefits that these services can have for scientific research, and the risk of widening resource inequalities among researchers, any new governance should be carefully designed and evaluated to avoid creating unnecessary barriers. Options to enhance oversight could include:
 - **Employ customer screening for life sciences service providers.** Experts have recommended that customer screening, described in Box 2, could be adopted by service providers to verify customers and that their services are being requested for legitimate purposes.⁵⁴ Which providers and types of services meet some threshold of dual-use potential, and thus should screen customers, is an open question.


- **Strengthen biosecurity standards for research services.** Cloud labs, for example, are not currently obligated to meet any biosecurity requirements, although many of the leading cloud lab companies have voluntarily adopted their own practices.⁵⁵ Potential best practices for cloud labs could include human review and approval of requested experiments, and maintaining experimental logs for monitoring and attribution.⁵⁶

Outcome: Release and Exposure (Intentional Deployment  or Accidental Containment Breach )


This is the stage at which a pathogen or toxin comes in contact with a susceptible human, a necessary step to cause physical harm. A biological agent that stays fully contained within a test tube or appropriate laboratory setting will not cause human infection or toxicity; it must be released from containment to do so.

Malicious actors may purposefully deploy a biological agent by intentionally dispersing the agent among chosen targets. The deployment process can vary widely in scale, requirements, and difficulty level based on the chosen targets and objectives. Malicious actors can have many goals, from targeting a few key people with a small amount of material, to large-scale distribution of a toxin, to strategic deployment of infectious diseases to initiate an outbreak, epidemic, or pandemic. Malicious actors may be able to use chatbots to help plan these operational steps, for example by attaining information about previously successful proliferation pathways or general security vulnerabilities to exploit.⁵⁷ While this knowledge can inform deployment, for the purposes of this report it is considered information-gathering and is included in the Planning and Design phase.

Researchers can be exposed to a pathogen or toxin during the course of scientific research if appropriate containment measures and practices are not taken, or if a laboratory accident causes a pathogen or toxin to escape those containment measures. Incidents that could lead to laboratory-acquired infections include procedural errors, splashes or spills, accidental needle punctures, bites from research animals, and engineering failures.⁵⁸

- **Policy Option: Enable law enforcement intervention by enhancing awareness of malicious activity.**  If a malicious actor successfully prepares a pathogen or toxin, the predominant remaining safeguard is for intelligence and law enforcement agencies to become aware of the threat and intervene before it is

deployed. Such intelligence efforts could be improved by elevating and coordinating biological threat intelligence collection. Detecting potential bioterrorists and their plans of deployment requires the collection and integration of many types of intelligence, spanning jurisdictions and agency roles. Expert recommendations to better coordinate such efforts include designating a National Intelligence Manager for Biological Threats and maximizing interagency intelligence-sharing between appropriate entities.⁵⁹

- **Policy Option: Strengthen laboratory biosafety to prevent or lessen the consequences of research accidents.**  Proper biosafety practices, systems, and infrastructure can make research safer by preventing exposures from occurring in the first place, or by responding appropriately to mitigate the harm if they do occur. Effective laboratory biosafety strategies are multilayered to address the many potential points of failure that can occur during research, including failures to follow procedures, handle materials and equipment appropriately, safeguard facilities, and effectively train personnel.


In the United States, the BMBL is the primary document guiding the safety and biocontainment practices for research involving infectious microorganisms and hazardous biological materials (see Box 1 for details). The BMBL recommends standard practices and risk assessments to inform safe research with hazardous biological agents, but is not a regulatory document and cannot itself require researchers to take any particular measures. The NIH Guidelines, DURC/PEPP Policy, and FSAP all involve risk mitigation plans and require evidence that the investigator and research institution can conduct the relevant research safely and securely and can respond to potential lapses (Box 1). Actions that could strengthen laboratory biosafety include:

- **Expand standardized biosafety practices to all research.** Of the described research oversight, only the FSAP is required regardless of a research project or institution's funding source. In contrast, the BMBL provides guidance and best practices but is not a regulatory document, and the NIH Guidelines and DURC/PEPP Policy apply only to research projects that meet specified federal funding criteria. All three regulatory policies (FSAP, NIH Guidelines, and DURC/PEPP Policy) are limited to research with specified pathogens or experimental conditions. As discussed on page 17, adherence to biosafety practices such as those outlined in the BMBL could be made a condition of all research funding, including from private, nongovernmental funders.

- **Increase support and resources for institutional biosafety efforts.** Institutions have unique insight into their facilities' available resources, expertise, and containment infrastructure and are the first line of defense in the event of an incident. Yet institutional biosafety officers are often under-resourced, especially as new duties are formally or informally added to their roles. Experts have suggested a range of solutions including funding for additional biosafety officers, developing a framework to train them, and establishing working groups or networks of IBCs, IREs, and biosafety and biosecurity experts to share and standardize best practices across institutions.⁶⁰
- **Increase research to inform evidence-based biological risk management.** Information about real-world biocontainment failures, efficacy of risk-reduction measures, and experimental outcomes could help to design more effective and efficient biosafety protocols.⁶¹ For example, research on how aging affects a material's biocontainment efficacy could inform requirements for equipment recertification and replacement. Possible solutions include incentivizing research for biosafety measurements and implementing no-fault reporting structures for various types of failures, accidents, or near-accidents.⁶²

Outcome: Potential Spread

Some infections can spread beyond the initial infected individual and lead to an outbreak, epidemic, or pandemic. While a large-scale outbreak can have substantial negative consequences, a single infection or localized outbreak can still be highly impactful and cause public alarm or panic.


- **Policy Option: Improve pandemic and public health preparedness to better respond to a biological event.**  Regardless of how an outbreak starts—whether by intentional misuse, a research accident, or natural causes—robust pandemic and public health infrastructure is an important line of defense. Efforts to quickly detect and monitor an outbreak, develop and manufacture diagnostics and medical countermeasures (MCMs), and secure medical supply chains can all reduce the chance of a local outbreak turning into a global pandemic. Pandemic and public health preparedness spans disciplines, jurisdictions, and federal agencies, and expert suggestions for improvement are too extensive to exhaustively list here. Some select considerations include:

- **Improve and expand global biosurveillance.** Biosurveillance, described on page 9, can aid in pandemic and public health preparedness by quickly detecting disease-causing agents regardless of their origin. For both naturally occurring and engineered pathogens, early detection can inform resource allocation, public health measures, therapy development, and emergency response strategies.⁶³ In addition to the actions to improve U.S. biosurveillance described on page 9, biosurveillance will be most effective as a global endeavor that includes international information and resource-sharing because pathogens do not limit their spread to national borders.
- **Increase research and development to proactively prepare diagnostics and MCMs for potential threats.** The United States could have a more proactive pipeline of research, therapies, and interventions for future potential infectious diseases.⁶⁴ A more forward-thinking approach would prioritize research and development to better understand infectious diseases and to develop MCMs and diagnostic tests that are either broad-spectrum or adaptable to a variety of potential threats.⁶⁵ The United States can be more proactive by increasing funding for government initiatives such as the Biomedical Advanced Research and Development Agency's Project BioShield, creating a market for MCM development, and standardizing a regulatory pathway for adaptable technologies.⁶⁶
- **Investigate AI tools to assist with pandemic preparedness and response.** AI tools can provide information to help public health officials respond to outbreaks more quickly and effectively. For example, AI can identify patterns that indicate an outbreak, predict the impact of emerging pathogen strains, speed up antibody and vaccine design, and optimize supply chain and distribution networks.⁶⁷ Future strategies could build on these advances by incentivizing research to develop further capabilities and assess how to implement them in public health systems.

Outcome: Scientific Findings and Potential Information Hazard

The overarching goal of research is to learn more about a biological system. Researchers share these new insights with their peers by publishing their findings, for example in peer-reviewed scientific journals or online as preprints. Open, transparent, and accessible scientific literature is widely valued in the scientific community because

it allows other researchers to generate new ideas, collectively build upon shared information, and test the reproducibility of other researchers' work. However, some research findings could constitute an information hazard if published and misused by malicious actors, even if the research itself was conducted responsibly.

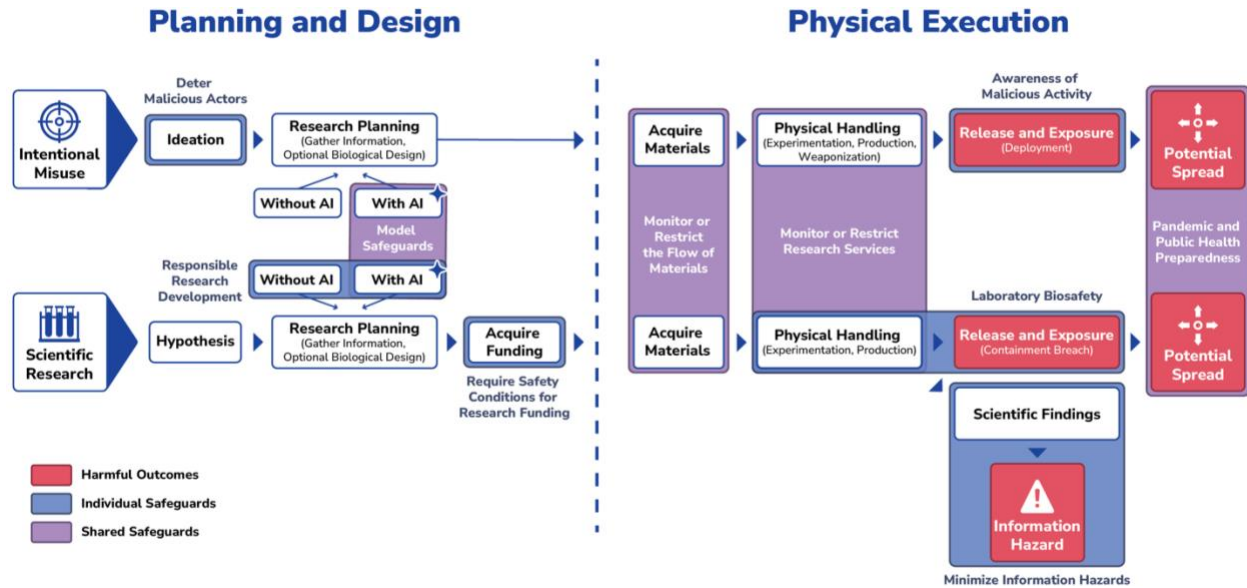
- **Policy Option: Minimize information hazards for potentially dual-use scientific information.**  Deciding how to handle potential information hazards raises challenging questions about how to weigh the risks and benefits of publishing potentially dual-use information. As has been discussed, biology's dual-use nature means that the information, protocols, and materials that could be misused to cause harm also have widespread practical uses. Furthermore, security concerns will need to be balanced against scientific openness and the risk that withholding information could undermine public trust in science.

The DURC/PEPP Policy's Implementation Guidance includes a Risk-Benefit Analysis of Communication to help decide whether, and to what extent, to share findings that pose potential dual-use risks (see Box 1 for details).⁶⁸ Possible measures to consider include:

- **Standardize publication practices for scientific publishers.** Editorial policies for potential dual-use information vary between scientific journals and can be vague. For example, it is unclear whether some journals would allow manuscripts to be published with redacted or incomplete information, and how different regional norms affect global publication practices. Publishing norms could be strengthened through a forum for publishers to share best practices and develop standardized approaches, and by creating or standardizing policies for other publication outlets like preprint servers.⁶⁹
- **Expand the practice of conducting risk-benefit communication analyses.** The risk-benefit communication analysis included in the DURC/PEPP Policy's Implementation Guidance provides a framework to help evaluate risk factors and publication options for dual-use information. It encourages researchers to consider factors such as the time scale and magnitude of potential harm, available mitigation measures, and ease for a malicious actor to carry out.⁷⁰ This or a similar framework could be expanded to research that is not subject to the policy, or it could be made a condition for other types of research funding, or integrated into institutional policies.

Concluding Thoughts

Figure 4. Safety and Security Toolkit for Intentional Misuse and Scientific Research



Source: CSET.

Note: Dashed blue line denotes the Planning-to-Physical transition, and red boxes denote harmful outcomes. Blue boxes denote safeguards that are specific to intentional misuse or scientific research, while purple boxes indicate safeguards that apply to both scenarios.

Figure 4 illustrates the broad spectrum of tools that policymakers could draw upon to mitigate both AI-agnostic and AI-enhanced biological risks, spread across the planning and physical stages. As the United States strengthens its existing biosafety and biosecurity ecosystem, and incorporates additional safeguards for new and emerging threats, it will also be important to keep the following considerations in mind.

First, strategies that reduce risks should not be ignored in favor of strategies that totally remove risks. Total risk mitigation is largely impossible to implement, while risk reduction measures may be more immediately feasible and less likely to impede scientific innovation. For example, list-based biosecurity policies that apply only to specific, predetermined pathogens have widely recognized limitations because they do not include every possible harmful agent, only include existing agents and not future or unknown ones, and can be evaded by savvy actors. Yet oversight that is based on an imperfect list still reduces overall risk by comparison with no oversight at all. In keeping with the adage “perfect is the enemy of good,” the United States should

implement today's imperfect but feasible safeguards while investing in the development of tomorrow's more advanced strategies.

Second, future safety initiatives will need to be balanced against their trade-offs for scientific research. Any measures that restrict information, resources, materials, or experiments in pursuit of risk mitigation also have the potential to hinder legitimate research. This outcome may be acceptable in some cases, especially when the impact on the beneficial application is minor or when the negative consequence is particularly severe. For example, it may be justifiable to limit the types of information that chatbots provide about historical bioweapons attacks even if it limits their utility for biosafety research, or to have stringent safeguards for high-consequence agents such as anthrax or smallpox. In other cases, the chilling effect on scientific research may be deemed to outweigh the safety improvements. These determinations should be informed by engagement with researchers to gauge how regulations would impact their work, and by a biological risk framework to help stakeholders identify and prioritize risk factors.

Finally, future researchers and stakeholders should think more broadly about many types of potential biological harms. This report, like much of the recent discussion about biosafety, biosecurity, and biodefense, has been very narrowly focused on risks from human pathogens and toxins. Such a limited scope leaves out other potential consequences such as those resulting from plant and animal pathogens, human genetic engineering, environmental harms, and others. For emerging technologies such as AI, neglecting to consider these factors now could result in missing proactive opportunities to implement safety measures.


In sum, viewing biological harms as a multistep process introduces a range of options at various stages to mitigate risks. As the United States continues to expand its biosecurity and biodefense capabilities, it will benefit from using each of these policy tools to build a more robust and layered system of safeguards.

Author

Steph Batalis is a research fellow at CSET.

Acknowledgments

The author is grateful to Alexander Titus, James Diggans, and Matthew E. Walsh for reviewing and providing feedback on this report. For additional comments and assistance, special thanks to Igor Mikolic-Torreira, Helen Toner, Vikram Venkatram, Mina Narayanan, Shelton Fitch, and Matt Mahoney.



© 2021 by the Center for Security and Emerging Technology. This work is licensed under a Creative Commons Attribution-Non Commercial 4.0 International License. To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc/4.0/>.

Document Identifier: doi: 10.51593/20240013

Endnotes

¹ U.S. Department of Defense, *2023 Biodefense Posture Review* (August 16, 2023), https://media.defense.gov/2023/Aug/17/2003282337/-1/-1/1/2023_BIODEFENSE_POSTURE_REVIEW.PDF.

² Todd Kuiken, “Artificial Intelligence in the Biological Sciences: Uses, Safety, Security, and Oversight” (Congressional Research Service, November 22, 2023), <https://crsreports.congress.gov/product/pdf/R/R47849>; Christopher East, “The UK’s AI Safety Summit and the Future of AI Bioweapons” (The Council on Strategic Risks, November 6, 2023), <https://councilonstrategicrisks.org/2023/11/06/the-uks-ai-safety-summit-and-the-future-of-ai-bioweapons/>.

³ Subcommittee on Emerging Threats and Spending Oversight, “Advanced Technology: Examining Threats to National Security,” Homeland Security Governmental Affairs, September 19, 2023, <https://www.hsgac.senate.gov/subcommittees/etso/hearings/advanced-technology-examining-threats-to-national-security/>.

⁴ David Gillum, George Poste, Craig Woods, and Rachel Levinson, “Biotech Promises Miracles. But the Risks Call for More Oversight,” *Bulletin of the Atomic Scientists*, August 31, 2023, <https://thebulletin.org/2023/08/biotech-promises-miracles-but-the-risks-call-for-more-oversight/>.

⁵ Department of Homeland Security, Department of Homeland Security Report on Reducing the Risks at the Intersection of Artificial Intelligence and Chemical, Biological, Radiological, and Nuclear Threats (Washington, D.C., April 26, 2024), https://www.dhs.gov/sites/default/files/2024-06/24_0620_cwmd-dhs-cbrn-ai-eo-report-04262024-public-release.pdf; “New Interactive Tool from the Bipartisan Commission on Biodefense Maps the Nation’s Biodefense Enterprise” (Bipartisan Commission on Biodefense, July 10, 2024), <https://biodefensecommission.org/new-interactive-tool-from-the-bipartisan-commission-on-biodefense-maps-the-nations-biodefense-enterprise/>.

⁶ Arielle D’Souza and Janika Schmitt, “Mapping America’s Biosurveillance” (Institute for Progress, April 3, 2024), <https://ifp.org/mapping-americas-biosurveillance/>; Tina Won Sherman, “Weapons of Mass Destruction: DHS Has Made Progress in Some Areas, but Additional Improvements Are Needed,” Testimony to the Subcommittee on Emergency Management and Technology, 118th Congress, March 20, 2024, <https://www.gao.gov/assets/d24107426.pdf>; *Biodefense: DHS Exploring New Methods to Replace BioWatch and Could Benefit from Additional Guidance* (Washington, DC: Government Accountability Office, May 20, 2021), <https://www.gao.gov/products/gao-21-292>.

⁷ *Biodefense: National Biosurveillance Integration Center Has Taken Steps to Address Challenges, But Could Better Assess Results* (Washington, DC: Government Accountability Office, November 29, 2023), <https://www.gao.gov/products/gao-24-106142>; *The National Blueprint for Biodefense* (Bipartisan Commission on Biodefense, April 2024), https://biodefensecommission.org/wp-content/uploads/2024/05/National-Blueprint-for-Biodefense-2024_final_digital.pdf; D’Souza and Schmitt, “Mapping America’s Biosurveillance.”

⁸ Matthew E Walsh and Gigi Kwik Gronvall, “Discussion on the Future Science and Technology of Biological Attribution” (Johns Hopkins Center for Health Security, December 6, 2023), <https://centerforhealthsecurity.org/sites/default/files/2023-02/20230124-bioattribution-mtg-rpt.pdf>; *Pandemic Origins: Technologies and Challenges for Biological Investigations* (Washington, DC: Government Accountability Office, January 27, 2023), <https://www.gao.gov/products/gao-23-105406>.

⁹ National Biodefense Strategy and Implementation Plan for Countering Biological Threats, Enhancing Pandemic Preparedness, and Achieving Global Health Security (Washington, DC: The White House, October 2022), <https://www.whitehouse.gov/wp-content/uploads/2022/10/National-Biodefense-Strategy-and-Implementation-Plan-Final.pdf>; Tracey Rissman and Annette Prieto, “Attributing Biological Weapons Use: Strengthening Department of Defense Capabilities to Investigate Deliberate Biological Incidents” (RAND Corporation, February 6, 2024), https://www.rand.org/pubs/research_reports/RRA2360-1.html; Jaime M. Yassif, Shayna Korol, and Angela Kane, “Guarding Against Catastrophic Biological Risks: Preventing State Biological Weapon Development and Use by Shaping Intentions,” *Health Security* 21, no. 4 (August 2023): 258–65, <https://liebertpub.com/doi/10.1089/hs.2022.0145>; D’Souza and Schmitt, “Mapping America’s Biosurveillance”; Walsh and Gronvall, “Discussion on the Future Science and Technology of Biological Attribution”; *Pandemic Origins: Technologies and Challenges for Biological Investigations* (Washington, DC: Government Accountability Office, January 27, 2023), <https://www.gao.gov/products/gao-23-105406>.

¹⁰ For a more thorough description of the various sources of scientific information, see Steph Batalis, “AI and Biorisk: An Explainer” (Center for Security and Emerging Technology, December 2023), <https://cset.georgetown.edu/publication/ai-and-biorisk-an-explainer/>.

¹¹ Batalis, “AI and Biorisk: An Explainer.”

¹² Kyle Miller, “Open Foundation Models: Implications of Contemporary Artificial Intelligence” (Center for Security and Emerging Technology, March 12, 2024), <https://cset.georgetown.edu/article/open-foundation-models-implications-of-contemporary-artificial-intelligence/>.

¹³ Sarah R. Carter, Nicole E. Wheeler, Sabrina Chwalek et al., “The Convergence of Artificial Intelligence and the Life Sciences,” *Nuclear Threat Initiative*, October 30, 2023, <https://www.nti.org/analysis/articles/the-convergence-of-artificial-intelligence-and-the-life-sciences/>.

¹⁴ Stephanie Batalis, Caroline Schuerger, Gigi Kwik Gronvall, and Matthew E. Walsh, “Safeguarding Mail-Order DNA Synthesis in the Age of Artificial Intelligence,” *Applied Biosafety* 29, no. 2 (December 27, 2023): 79–84, <https://www.liebertpub.com/doi/10.1089/apb.2023.0020>.

¹⁵ Carter, Wheeler, et al., “The Convergence of Artificial Intelligence and the Life Sciences”; Jessica Ji, Josh A Goldstein, and Andrew J Lohn, “Controlling Large Language Model Outputs: A Primer” (Center for Security and Emerging Technology, December 2023), <https://cset.georgetown.edu/publication/controlling-large-language-models-a-primer/>.

- ¹⁶ Christopher A. Moutun, Caleb Lucas, and Ella Guest, “The Operational Risks of AI in Large-Scale Biological Attacks: Results of a Red-Team Study” (RAND Corporation, January 25, 2024), https://www.rand.org/pubs/research_reports/RRA2977-2.html.
- ¹⁷ Natalie R Zelenka, Nina Di Cara, Kieren Sharma, et al., “Data Hazards in Synthetic Biology,” *Synthetic Biology* 9, no. 1 (January 1, 2024), <https://doi.org/10.1093/synbio/ysae010>; Department of Homeland Security, *Report on Reducing the Risks at the Intersection of Artificial Intelligence and CBRN Threats*; Carter, Wheeler, et al., “The Convergence of Artificial Intelligence and the Life Sciences.”
- ¹⁸ Charlie D. Johnson, Wilson Sinclair, and Rebecca Mackelprang, “Security Considerations at the Intersection of Engineering Biology and Artificial Intelligence” (Engineering Biology Research Consortium [EBRC], November 2023), <https://ebrc.org/wp-content/uploads/2023/11/Security-Considerations-at-the-Intersection-of-Engineering-Biology-and-Artificial-Intelligence-EBRC.pdf>; Carter, Wheeler, et al., “The Convergence of Artificial Intelligence and the Life Sciences.”
- ¹⁹ Federation of American Scientists, Johns Hopkins Center for Health Security, Nuclear Threat Initiative Global Biological Policy and Programs, and Scowcroft Institute of International Affairs, “Response to the NSCEB’s Interim Report and AlxBio Policy Options,” comments to the National Security Commission on Emerging Biotechnology, April 9, 2024, <https://centerforhealthsecurity.org/sites/default/files/2024-04/2024-04-09-joint-nsceb-response.pdf>.
- ²⁰ Carter, Wheeler, et al., “The Convergence of Artificial Intelligence and the Life Sciences.”
- ²¹ Jessica Ji, “What Does AI Red-Teaming Actually Mean?” *Center for Security and Emerging Technology* (blog), October 24, 2023, <https://cset.georgetown.edu/article/what-does-ai-red-teaming-actually-mean/>.
- ²² U.S. AI Safety Institute, *Managing Misuse Risk for Dual-Use Foundation Models* (NIST AI 800-1, July 2024), <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.800-1.ipd.pdf>. Includes considerations for model developers when assessing potential deployment options.
- ²³ Toby Shevlane, “Structured Access: An Emerging Paradigm for Safe AI Deployment,” in *The Oxford Handbook of AI Governance*, ed. Justin B. Bullock, Yu-Che Chen, et al. (Oxford University Press, 2024), doi: [10.1093/oxfordhb/9780197579329.013.39](https://doi.org/10.1093/oxfordhb/9780197579329.013.39).
- ²⁴ Owen Daniels and Dewey Murdick, “Enabling Principles for AI Governance” (Center for Security and Emerging Technology), July 2024, <https://cset.georgetown.edu/publication/enabling-principles-for-ai-governance/>; Ren Bin Lee Dixon and Heather Frase, “An Argument for Hybrid AI Incident Reporting” (Center for Security and Emerging Technology, March 2024), <https://cset.georgetown.edu/publication/an-argument-for-hybrid-ai-incident-reporting/>; Department of Homeland Security, *Report on Reducing the Risks at the Intersection of Artificial Intelligence and CBRN Threats*.
- ²⁵ National Security Commission on Emerging Biotechnology, *Interim Report* (December 2023), <https://www.biotech.senate.gov/wp-content/uploads/2024/01/NSCEB-December-2023-Interim->

[Report.pdf](#); Daniels and Murdick, “Enabling Principles for AI Governance”; Johnson, Sinclair, and Mackelprang, “Security Considerations at the Intersection of Engineering Biology and Artificial Intelligence.”

²⁶ Johnson, Sinclair, and Mackelprang, “Security Considerations at the Intersection of Engineering Biology and Artificial Intelligence.”

²⁷ Sarah R. Carter, “Developing a Customer Screening Framework for the Life Sciences” (Blueprint Biosecurity, March 2024), <https://blueprintbiosecurity.org/u/2024/03/KYC-Report--Developing-a-Customer-Screening-Framework-for-the-Life-Sciences.pdf>.

²⁸ Arianne Kane and Michael T. Parker, “Screening State of Play: The Biosecurity Practices of Synthetic DNA Providers,” *Applied Biosafety* 29, no. 2 (June 2024): 85–95, <https://www.liebertpub.com/doi/10.1089/apb.2023.0027>; Sophie Rose, Tessa Alexanian, et al., “Practical Questions for Securing Nucleic Acid Synthesis,” *Applied Biosafety*, March 14, 2024, <https://www.liebertpub.com/doi/10.1089/apb.2023.0028>; Carter, “Developing a Customer Screening Framework for the Life Sciences.”

²⁹ Carter, “Developing a Customer Screening Framework for the Life Sciences.”

³⁰ Johnson, Sinclair, and Mackelprang, “Security Considerations at the Intersection of Engineering Biology and Artificial Intelligence.”

³¹ Carter, Wheeler, et al., “The Convergence of Artificial Intelligence and the Life Sciences.”

³² Department of Homeland Security, Report on Reducing the Risks at the Intersection of Artificial Intelligence and CBRN Threats”

³³ Todd Kuiken, “U.S. Oversight of Laboratory Biosafety and Biosecurity: Current Policies, Recommended Reforms, and Options for Congress” (Congressional Research Service, September 15, 2023), <https://crsreports.congress.gov/product/pdf/R/R47695>.

³⁴ “International Bio Funders Compact,” The Nuclear Threat Initiative, accessed August 12, 2024, <https://www.nti.org/about/programs-projects/project/bio-funders-compact/>.

³⁵ For equipment that is considered to be dual-use, see equipment that is regulated by: The Australia Group Control List of Dual-use Biological Equipment and Related Technology and Software, https://www.dfat.gov.au/publications/minisite/theaustraliagroupnet/site/en/dual_biological.html; The Wassenaar Arrangement List of Dual-Use Goods and Technologies and Munitions List, <https://www.wassenaar.org/app/uploads/2023/12/List-of-Dual-Use-Goods-and-Technologies-Munitions-List-2023-1.pdf>; The European Union List of Dual-Use Items (Annex I to Council Regulation [EC] No. 428/2009), <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:02009R0428-20191231&from=EN#toclid43>; and the U.S. Bureau of Industry and Security, Supplement No. 1 to Part

774, the Commerce Control List, <https://www.bis.gov/ear/title-15/subtitle-b/chapter-vii/subchapter-c/part-774/supplement-no-1-part-774-commerce-control#category1>.

³⁶ Elliot Roth, “A Guide to DIYbio (updated 2019),” Medium, August 17, 2021, <https://thatmre.medium.com/a-guide-to-diybio-updated-2019-abd0956cdf74>.

³⁷ Patrick D’haeseleer, “How to Set Up Your Own DIY Bio Lab,” *Make*, April 11, 2017, <https://makezine.com/article/science/health-science/how-to-set-up-your-own-lab/>.

³⁸ For pathogens and toxins regulated by the Federal Select Agent Program, see: 7 CFR Part 331, 9 CFR Part 121, and 42 CFR Part 73. For pathogens and toxins regulated by the International Traffic in Arms Regulations (ITAR) see: 22 CFR Parts 120-130. For pathogens and toxins regulated by the Export Administration Regulations (EAR) see: 5 CFR Parts 730-774, <https://www.ecfr.gov/>.

³⁹ 22 CFR Parts 120-130; 5 CFR Parts 730-774, <https://www.ecfr.gov/>.

⁴⁰ International Gene Synthesis Consortium, <https://genesynthesisconsortium.org/>; Exec. Order No. 14110, 88 FR 7519 (2023).

⁴¹ Administration for Strategic Preparedness and Response, *Screening Framework Guidance for Providers and Users of Synthetic Nucleic Acids* (U.S. Department of Health and Human Services, October 2023), <https://aspr.hhs.gov/legal/synna/Documents/SynNA-Guidance-2023.pdf>.

⁴² Administration for Strategic Preparedness and Response, *Screening Framework Guidance for Providers and Users of Synthetic Nucleic Acids*.

⁴³ National Science Advisory Board for Biosecurity, “Addressing Biosecurity Concerns Related to the Synthesis of Select Agents” (December 2006), <https://biosecurity.fas.org/resource/documents/NSABB%20guidelines%20synthetic%20bio.pdf>; Becky Mackelprang, “Security Screening in Synthetic DNA Synthesis” (Engineering Biology Research Consortium, April 2022), <https://ebrc.org/wp-content/uploads/2022/04/EBRC-2022-Security-Screening-in-Synthetic-DNA-Synthesis.pdf>.

⁴⁴ National Institute of Standards and Technology, *Nonprofit Research Consortium to Develop Safety Tools for Synthetic Biology to Defend Against Potential Misuse of AI* (U.S. Department of Commerce, February 16, 2024), <https://www.nist.gov/news-events/news/2024/02/nist-nonprofit-research-consortium-develop-safety-tools-synthetic-biology>.

⁴⁵ Intelligence Advanced Research Projects Activity, *Fun GCAT: Functional Genomic and Computational Assessment of Threats* (Office of the Director of National Intelligence, September 2016), <https://www.iarpa.gov/research-programs/fun-gcat>.

⁴⁶ Bridget Williams and Rowan Kane, “Preventing the Misuse of DNA Synthesis” (Institute for Progress, February 15, 2023), <https://ifp.org/preventing-the-misuse-of-dna-synthesis/>; “Strengthening Gene

Synthesis Security,” National Security Commission on Emerging Biotechnology, June 2024, https://www.biotech.senate.gov/wp-content/uploads/2024/06/NSCEB_WP_Strengthening-Gen-Synthesis-Security.pdf; Batalis, Schuerger, Gronvall, and Walsh, “Safeguarding Mail-Order DNA Synthesis in the Age of Artificial Intelligence”; Rose, Alexanian, et al., “Practical Questions for Securing Nucleic Acid Synthesis.”

⁴⁷ Steph Batalis and Vikram Venkatram, “Breaking Down the Biden AI EO: Screening DNA Synthesis and Biorisk” *Center for Security and Emerging Technology* (blog), November 16, 2023, <https://cset.georgetown.edu/article/breaking-down-the-biden-ai-EO-screening-dna-synthesis-and-biorisk/>.

⁴⁸ National Security Commission on Emerging Biotechnology, *Gene Synthesis Security* (NSCEB, April 2024), https://www.biotech.senate.gov/wp-content/uploads/2024/04/NSCEB_WP_Gene-Synthesis-Screening.pdf; World Economic Forum and Nuclear Threat Initiative, “Biosecurity Innovation and Risk Reduction: A Global Framework for Accessible, Safe and Secure DNA Synthesis” (January 2020), <https://www.weforum.org/publications/biosecurity-innovation-and-risk-reduction-a-global-framework-for-accessible-safe-and-secure-dna-synthesis-582d582cd4/>; Nicole E. Wheeler, Sarah R. Carter, et al., “Developing a Common Global Baseline for Nucleic Acid Synthesis Screening,” *Applied Biosafety* 29, no. 2 (June 2024): 71–78, doi: [10.1089/apb.2023.0034](https://doi.org/10.1089/apb.2023.0034).

⁴⁹ Sarah R. Carter, Jaime M. Yassif, and Chris Isaac, “Benchtop DNA Synthesis Devices: Capabilities, Biosecurity Implications, and Governance,” *Nuclear Threat Initiative*, May 10, 2023, <http://www.nti.org/analysis/articles/benchtop-dna-synthesis-devices-capabilities-biosecurity-implications-and-governance/>; Williams and Kane, “Preventing the Misuse of DNA Synthesis”; Rose, Alexanian, et al., “Practical Questions for Securing Nucleic Acid Synthesis”; Carter, “Developing a Customer Screening Framework for the Life Sciences.”

⁵⁰ Carter, Wheeler, et al., “The Convergence of Artificial Intelligence and the Life Sciences.”

⁵¹ National Security Commission on Emerging Biotechnology, *Interim Report*.

⁵² Daniil A. Boiko, Robert MacKnight, and Gabe Gomes, “Emergent Autonomous Scientific Research Capabilities of Large Language Models,” arXiv preprint arXiv:2304.05332 (April 11, 2023), <http://arxiv.org/abs/2304.05332>.

⁵³ Hector G Martin, Tijana Radivojevic, et al., “Perspectives for Self-Driving Labs in Synthetic Biology,” *Current Opinion in Biotechnology* 79 (February 1, 2023): 102881. <https://www.sciencedirect.com/science/article/pii/S0958166922002154>.

⁵⁴ Carter, “Developing a Customer Screening Framework for the Life Sciences.”

⁵⁵ Johnson, Sinclair, and Mackelprang, Security Considerations at the Intersection of Engineering Biology and Artificial Intelligence.

⁵⁶ Filippa Lentzos and Cédric Invernizzi, “Laboratories in the Cloud,” *Bulletin of the Atomic Scientists*, July 3, 2019, <https://thebulletin.org/2019/07/laboratories-in-the-cloud/>; Federation of American Scientists, Johns Hopkins Center for Health Security, Nuclear Threat Initiative Global Biological Policy and Programs, and Scowcroft Institute of International Affairs, “Response to the NSCEB’s Interim Report and AlxBio Policy Options.”

⁵⁷ “GPT-4 System Card” (OpenAI, March 23, 2023), <https://cdn.openai.com/papers/gpt-4-system-card.pdf>.

⁵⁸ Stuart D. Blacksell, Sandhya Dhawan, et al., “Laboratory-Acquired Infections and Pathogen Escapes Worldwide between 2000 and 2021: A Scoping Review,” *The Lancet Microbe* 5, no. 2 (December 12, 2023): e194–202, [https://www.thelancet.com/journals/lanmic/article/PIIS2666-5247\(23\)00319-1/fulltext](https://www.thelancet.com/journals/lanmic/article/PIIS2666-5247(23)00319-1/fulltext).

⁵⁹ DOD, 2023 Biodefense Posture Review; The National Blueprint for Biodefense.

⁶⁰ David Gillum, “The Making of a Biosafety Officer,” *Issues in Science and Technology*, April 19, 2023, <https://issues.org/making-biosafety-officer-gillum/>; Angela L. Rasmussen, Gigi K. Gronvall, et al., “Virology—the Path Forward,” *Journal of Virology* 98, no. 1 (January 3, 2024): e01791-23, <https://journals.asm.org/doi/10.1128/jvi.01791-23>; Daniel Greene, Kathryn Brink, et al., “The Biorisk Management Casebook: Insights into Contemporary Practices,” Stanford Digital Repository, March 30, 2023, <http://www.nti.org/analysis/articles/the-biorisk-management-casebook-insights-into-contemporary-practices/>; *National Biodefense Strategy and Implementation Plan*.

⁶¹ Stuart D. Blacksell, Sandhya Dhawan, et al., “The Biosafety Research Road Map: The Search for Evidence to Support Practices in Human and Veterinary Laboratories,” *Applied Biosafety: Journal of the American Biological Safety Association* 28, no. 2 (June 1, 2023): 64–71, <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10277988/>; *National Biodefense Strategy and Implementation Plan*.

⁶² *Report of the Federal Experts Security Advisory Panel* (Federal Experts Security Advisory Panel, December 2014), <https://www.phe.gov/s3/Documents/fesap.pdf>; “Impact Assessment of Research on Infectious Agents” (workshop summary, American Society for Microbiology, September 2023), <https://asm.org:443/Reports/Impact-Assessment-of-Research-on-Infectious-Agents>; Alec Stapp and Arielle D’Souza, “Congress Can Improve Our Pandemic Preparedness” (Institute for Progress, January 18, 2024), <https://ifp.org/congress-can-improve-our-pandemic-preparedness/>; Clint A. Haines and Gigi Kwik Gronvall, “Improving U.S. Biosafety and Biosecurity: Revisiting Recommendations from the Federal Experts Security Advisory Panel and the Fast Track Action Committee on Select Agent Regulations,” *Applied Biosafety* 28, no. 1 (March 2023): 43–54, <https://www.liebertpub.com/doi/10.1089/apb.2022.0025>; Blacksell, Dhawan, et al., “The Biosafety Research Road Map”; Rasmussen et al., “Virology—the Path Forward.”

⁶³ “Biosurveillance and Pathogen Detection,” National Institute of Science and Technology, updated August 19, 2024, <https://www.nist.gov/programs-projects/biosurveillance-and-pathogen-detection>;

D'Souza and Schmitt, "Mapping America's Biosurveillance"; DOD, *2023 Biodefense Posture Review; National Biodefense Strategy and Implementation Plan*.

⁶⁴ Caroline Schueger, Steph Batalis, et al., "Viral Families and Disease X: A Framework for U.S. Pandemic Preparedness Policy" (Center for Security and Emerging Technology, April 2023), <https://cset.georgetown.edu/publication/viral-families-and-disease-x-a-framework-for-u-s-pandemic-preparedness-policy/>; *The National Blueprint for Biodefense*.

⁶⁵The National Blueprint for Biodefense; National Biodefense Strategy and Implementation Plan.

⁶⁶The National Blueprint for Biodefense.

⁶⁷ Juan Cambeiro, "How AI Can Help Prevent Biosecurity Disasters," Institute for Progress, July 10, 2023, <https://ifp.org/how-ai-can-help-prevent-biosecurity-disasters/>; National Security Commission on Emerging Biotechnology, *Interim Report*; Carter, Wheeler, et al., "The Convergence of Artificial Intelligence and the Life Sciences."

⁶⁸ The White House, Implementation Guidance for the United States Government Policy for Oversight of Dual Use Research of Concern and Pathogens with Enhanced Pandemic Potential, (National Science and Technology Council, May 6, 2024), <https://www.whitehouse.gov/wp-content/uploads/2024/05/USG-DURC-PEPP-Implementation-Guidance.pdf>.

⁶⁹ National Security Commission on Emerging Biotechnology, *AIxBio White Paper 4: Policy Options for AIxBio* (NSCEB White Paper Series on AIxBio, January 2024), <https://www.biotech.senate.gov/press-releases/aixbio-white-paper-4-policy-options-for-aixbio/>.

⁷⁰ The White House, Implementation Guidance for the United States Government Policy for Oversight of Dual Use Research of Concern and Pathogens with Enhanced Pandemic Potential.