JOHN P. TARANGELO, AURELIA ATTAL-JUNCQUA, EVAN DAVISON KOTLER, CASEY AVEGGIO, MARIE BURNS, BARBARA DEL CASTELLO, TYLER BURNS, STEPH GUERRA, SASKIA POPESCU, DAVID GLICKSTEIN

Building Resilience Against Artificial Intelligence–Enabled Biological Threats

Conference Proceedings from a Follow-Up RAND and Helena Workshop, May 2026

This work was independently initiated and made possible by RAND's own resources, including income from operations, unrestricted gifts from RAND supporters, and distributions from its long-term investment fund.

Conference Proceedings

For more information on this publication, visit **www.rand.org/t/CFA4954-2**.

About RAND

RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our research integrity is grounded in RAND's core values of quality and objectivity. Rigorous quality assurance procedures, conflict of interest screening, and transparency in funding ensure that every study is objective and nonpartisan. Learn more at www.rand.org/integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.

© 2026 RAND Corporation
RAND® is a registered trademark.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, visit www.rand.org/about/publishing/permissions.

CF-A4954-2

About These Conference Proceedings

This document summarizes a workshop that took place on May 7, 2026, in Washington, D.C. and was hosted by RAND and Helena. This was the second workshop in a two-part series. The first workshop took place in January 2026, also in Washington, D.C. The January 2026 workshop convened technical experts from the fields of artificial intelligence (AI) research, biotechnology, biosecurity, and academia to determine priority mitigations for reducing the risk of AI-enabled biological threats. A description of the January 2026 workshop and the resulting mitigations are detailed in another publication titled *Building Resilience Against Artificial Intelligence–Enabled Biological Threats: Conference Proceedings from a RAND and Helena Workshop*.[1] The May 2026 workshop convened a group of policy experts with the goal of outlining the implementation of and key policy considerations for the mitigations identified during the January 2026 workshop.

RAND Center on AI, Security, and Technology

RAND Global and Emerging Risks is a division of RAND that delivers rigorous and objective public policy research on the most consequential challenges to civilization and global security. This work was undertaken by the division's Center on AI, Security, and Technology, which aims to examine the opportunities and risks of rapid technological change, focusing on artificial intelligence, security, and biotechnology. For more information, contact cast@rand.org.

Helena

Helena is an independent problem-solving institution integrating for-profit, nonprofit, and legislative action to address pressing societal challenges. Through convenings of scientists, ethicists, domestic and international policymakers, industry leaders, and intergovernmental organizations in the health and biosecurity spaces, Helena has been at the forefront of establishing best practices and guardrails for emerging technologies, enhancing global biosurveillance, and strengthening rapid-response mechanisms against biological threats.

Funding

This effort was jointly initiated and conducted by the Center on AI, Security, and Technology and Helena. Helena's work on this effort was partially supported by the Future of Life Institute. One participant from the workshop is affiliated with this organization; however, the participant attended in a personal, expert capacity. The participant had no role in shaping the workshop agenda, facilitating discussions, or determining the content or conclusions of these proceedings. The Center on AI, Security, and Technology's efforts used income from operations and gifts and grants from philanthropic supporters. A complete list of donors and funders is available at www.rand.org/CAST. RAND clients, donors, and grantors have no influence over research findings or recommendations.

[1] Aurelia Attal-Juncqua, John P. Tarangelo, Evan Davison Kotler, Casey Aveggio, Marie Burns, Barbara Del Castello, Tyler Burns, Steph Guerra, and Saskia Popescu, *Building Resilience Against Artificial Intelligence–Enabled Biological Threats: Conference Proceedings from a RAND and Helena Workshop,* RAND Corporation, CF-A4954-1, 2026.

Acknowledgments

We are grateful to our colleagues Henry Willis, for his thoughtful edits, and Jason Etchegaray, for quality assurance management, as well as Paige Smith, Sean Mintus, Pete Soriano, and Dan Cabrera for their assistance during the quality assurance review and publication processes. We also thank Alana Morro and Emily Garner for their planning and day-of support, as well as the Rockefeller Foundation for providing the venue at which RAND and Helena hosted the workshop. Finally, we thank the experts who attended this workshop for dedicating their time and expertise. The takeaways and recommendations expressed in this document are our views that we based on the discussions that took place during the workshop—and our takeaways and recommendations might or might not reflect the views of the experts who participated. The views that the experts shared during the workshop do not necessarily represent the views of any institution or organization with which the consulted participants are, or were previously, affiliated, including the U.S. government.

Contents

About These Conference Proceedings iii
Figures vii

Building Resilience Against Artificial Intelligence–Enabled Biological Threats: Conference Proceedings from a Follow-Up RAND and Helena Workshop, May 2026 1
January 2026 Workshop Overview 2
May 2026 Workshop 4
Conclusion 20

APPENDIX
Scenario Summaries 21
Scenario 1: Pandemic Disruption with a Virus Scenario 21
Scenario 2: Agroterrorism with a Fungus Scenario 22
Scenario 3: Critical Infrastructure Attack with a Bacteria Scenario 22

Abbreviations 25
References 27
About the Authors 29

Figures

1. Ten Priority Mitigations Across Three Scenarios from the January 2026 Workshop 3
2. How the Crosscutting Themes from Workshop 1 Informed Implementation Considerations During Workshop 2 6
3. Scenario 1: Implementation Considerations for Mitigation 1 8
4. Scenario 1: Implementation Considerations for Mitigation 2 9
5. Scenario 1: Implementation Considerations for Mitigation 3 10
6. Scenario 2: Implementation Considerations for Mitigation 1 12
7. Scenario 2: Implementation Considerations for Mitigation 2 13
8. Scenario 2: Implementation Considerations for Mitigation 3 14
9. Scenario 3: Implementation Considerations for Mitigation 1 16
10. Scenario 3: Implementation Considerations for Mitigation 2 17
11. Scenario 3: Implementation Considerations for Mitigation 3 18
12. Scenario 3: Implementation Considerations for Mitigation 4 19

Building Resilience Against Artificial Intelligence–Enabled Biological Threats: Conference Proceedings from a Follow-Up RAND and Helena Workshop, May 2026

On May 7, 2026, RAND and Helena hosted a workshop in Washington, D.C. This workshop brought together 15 individuals with expertise in biosecurity policy, health security policy, artificial intelligence (AI) policy, national security policy, and foreign policy. RAND and Helena convened this group of experts to assess a list of mitigations to combat AI-enabled biological threats.

The mitigations that were assessed were outputs from a previous RAND and Helena workshop. The previous workshop occurred on January 14 and 15, 2026, and was technically focused, whereas the May 2026 workshop was policy-focused. The January 2026 workshop brought together technical experts in AI research, biotechnology, biosecurity, and academia. The technical experts were charged with developing priority mitigations in the context of three scenarios to reduce the risk of AI-enabled biological threats. The January 2026 workshop resulted in 13 priority mitigations.[1] These mitigations are published in another conference proceedings, titled *Building Resilience Against Artificial Intelligence–Enabled Biological Threats: Conference Proceedings from a RAND and Helena Workshop*.[2] The policy experts convened at the May 2026 workshop were asked to both examine how those mitigations could be implemented and to identify key policy considerations that could affect the mitigations' implementation.

The convergence of AI and the life sciences could yield beneficial scientific breakthroughs.[3] Simultaneously, this convergence creates new opportunities for AI-enabled biology to be misused.[4] RAND and Helena hosted these workshops to (1) identify promising mitigations, (2) explore mitigations' implementation options, and (3) address key policy considerations to combat the potential for AI-enabled biology to be misused to cause harm.

---

[1] Four of the mitigations developed during the January 2026 workshop were closely related and complementary. Therefore, we grouped them under a single mitigation in our first document and presented the mitigation as having four component parts. This mitigation and its four components are described under Priority Mitigation 4 ("Improved Threat-Detection Mechanisms," as Mitigations 4A–4D) in Chapter 5 of Aurelia Attal-Juncqua, John P. Tarangelo, Evan Davison Kotler, Casey Aveggio, Marie Burns, Barbara Del Castello, Tyler Burns, Steph Guerra, and Saskia Popescu, *Building Resilience Against Artificial Intelligence–Enabled Biological Threats: Conference Proceedings from a RAND and Helena Workshop,* RAND Corporation, CF-A4954-1, 2026. Participants in the May 2026 workshop reviewed this as one mitigation with four parts (as presented in Attal-Juncqua et al.) rather than as four individual mitigations. Therefore, this document contains ten mitigations rather than 13.

[2] Attal-Juncqua et al., 2026.

[3] Nicole E. Wheeler, "Responsible AI in Biotechnology: Balancing Discovery, Innovation and Biosecurity Risks," *Frontiers in Bioengineering and Biotechnology*, Vol. 13, February 5, 2025.

[4] David Luckey, Sara Duhachek Muggy, Taylor Frey, David Stebbins, Tracey Rissman, Bianca Espinosa, Daniel Tapia, Greg McKelvey, Jr., Neeti Pokhriyal, Joseph Nelson Dawson, Sara Hughes, Morgan Sandler, Rushil Bakhshi, Marta Kepe, Geoffrey Kirkwood, Sarah W. Denton, David DeSmet, Minami Makino, Ella Guest, Sina Beaghley, Suzanne Genc, Michael Miller, Skye A. Miner, Barbara Del Castello, Forrest W. Crawford, Jeffrey Lee, Clay Strickland, Sunny D. Bhatt, John Vahedi, Lydia Grek, Vanya Barrer, Ramiro Insuasti, Jr., Jack Lashendock, Derek Roberts, Aleksandr Esparza Hartunian, Shannon Walsh, Will Shumate, Elliott Brennan, Tyler Liggett, Kara Jia, Ajay K. Kochhar, James Smith, and James Ryseff, *Mitigating Risks at the Intersection of Artificial Intelligence and Chemical and Biological Weapons*, Homeland Security Operational Analysis Center operated by the RAND Corporation, RR-A2990-1, 2025.

January 2026 Workshop Overview

As mentioned, the results of the January 2026 workshop are detailed in another publication.[5] However, we include a summary here because the context of that workshop and its resulting recommendations are critical for understanding the May 2026 workshop.

During the January 2026 workshop, participants examined risks arising at the intersection of AI and biology (AI x Bio) across the following biological risk chain: **Intention → Ideation→ Design→ Build → Test → Learn → Weaponization**.

The stages of the risk chain are defined in the prior publication and reproduced here:

- *intention*—identifying an actor's initial motivation to pursue a biological solution to a strategic aim (e.g., a state actor's motivation) or grievance (e.g., an insider threat's or terrorist's motivation)
- *ideation*—determining a suitable biological attack plan with a particular emphasis on biological agent selection and converting the abstract intent into a specific technical concept that supports a goal
- *design*—using biology and AI tools to create or select an agent
- *build*—synthesizing biological material or crossing the digital-to-physical divide
- *test*—verifying that the physical agent functions as intended
- *learn*—analyzing in silico, in vivo, or in vitro or testing information to refine the agent
- *weaponization*—stabilizing and scaling the agent for dissemination.

We developed three scenarios to illustrate the development cycle, and we used these scenarios as analytical tools to stress test technologies, institutions, and governance arrangements; surface implicit assumptions; underscore risk tolerance and awareness; and identify mitigation strategies that could be relevant across multiple threat pathways. Twenty-two participants attended the workshop and had backgrounds in AI research, biotechnology, biosecurity, and academia. Participants worked under the Chatham House rule and were organized into small, interdisciplinary teams tasked with developing high-impact mitigation strategies across key stages of the AI x Bio value chain.[6] We divided participants into three groups, and we assigned each group a different scenario.

The scenarios were developed to capture a wide variety of threat targets and threat vectors. The first scenario involved a highly capable nonstate actor intent on causing catastrophic harm to the human population by developing a novel influenza A strain. The second scenario involved a moderately skilled nonstate group that created a modified agricultural fungus, disrupting both the food supply and the economy. Finally, the third scenario involved two individuals who were insider threats and used a bacterium to attack critical infrastructure. We include brief descriptions of the scenarios that participants analyzed. Longer summaries are included in the previous conference proceedings.[7] The scenario summaries from the previous conference proceedings are reproduced in this document in the appendix. In both documents, we intentionally include a summary of the scenarios rather than the full scenario because of infohazard concerns. However, participants at both workshops had access to the full scenarios.

In the January 2026 workshop, the groups were tasked with developing technical mitigations to prevent their assigned scenario from occurring. Each scenario was assumed to take place in the near future, defined as over the next two to five years (2028 to 2031). Groups were *not* responsible for developing implementation plans. Rather, the focus was solely on the technical interventions that each scenario revealed. Each group developed a "Mitigation Roadmap" consisting of a set of interventions that they believed would interrupt

[5] Attal-Juncqua et al., 2026.

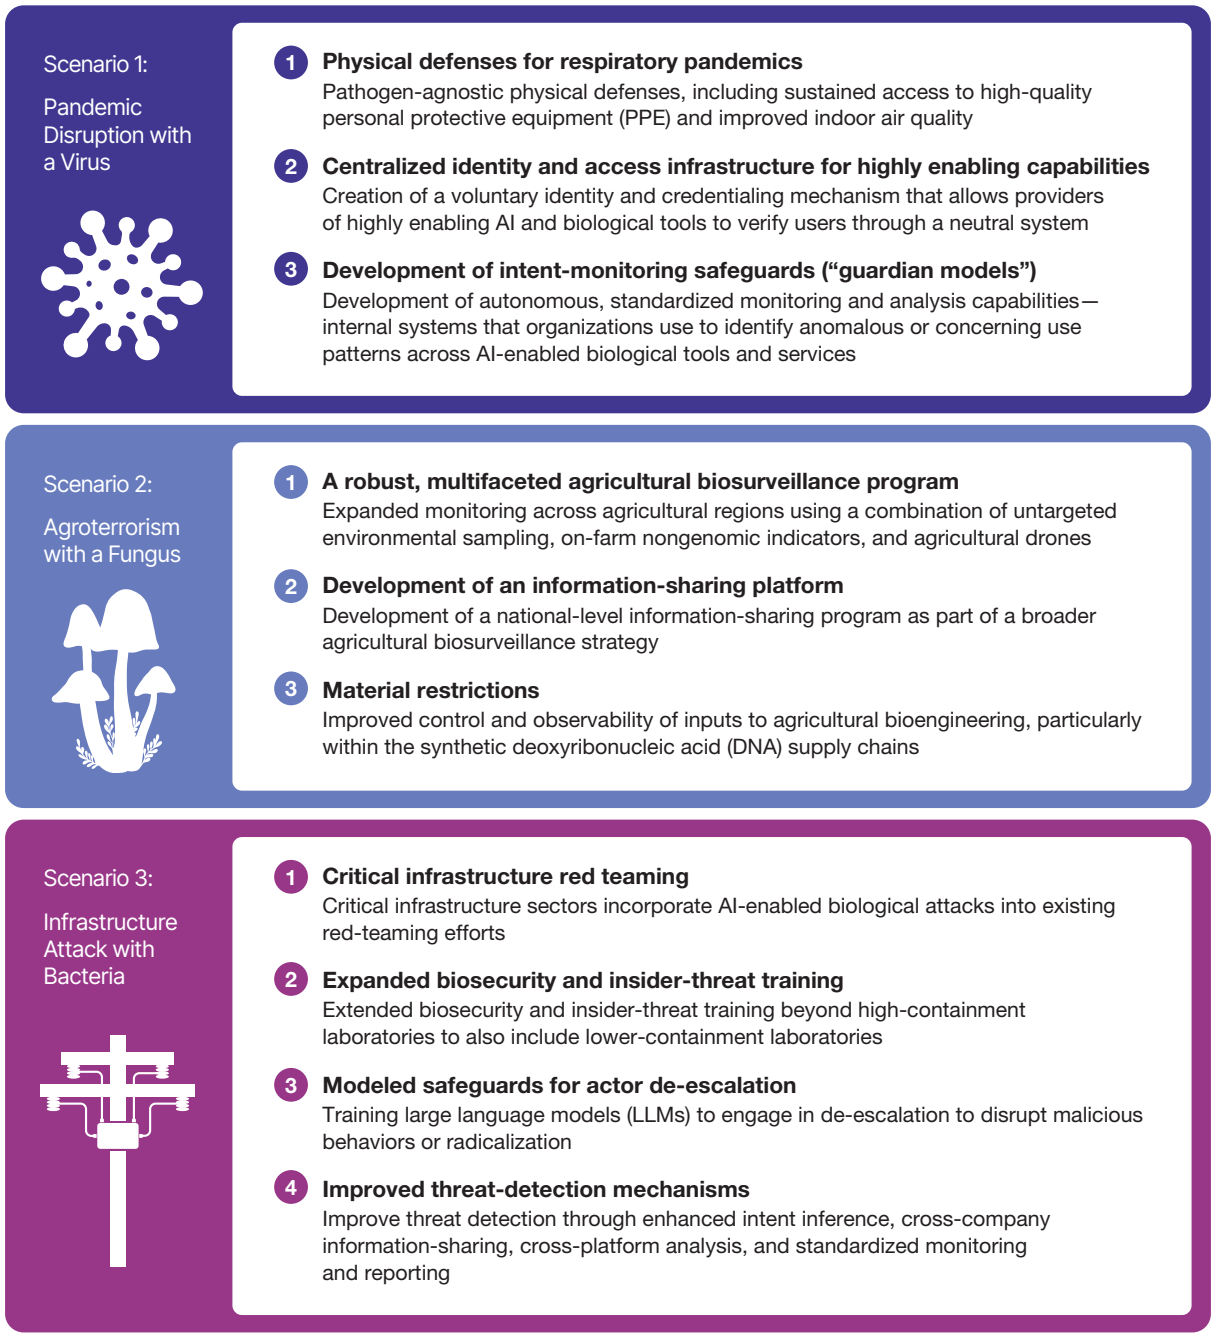
[6] Chatham House, "Chatham House Rule," webpage, undated.

[7] Attal-Juncqua et al., 2026.

the scenario at one or more points along the development cycle. Groups then exchanged roadmaps and conducted a peer review exercise to stress test their proposed mitigations. After receiving feedback, each group refined their Mitigation Roadmap and selected three top priority mitigations to present at the close of the workshop.[8] Figure 1 displays the top priority mitigations that were developed at the January workshop.

FIGURE 1

Ten Priority Mitigations Across Three Scenarios from the January 2026 Workshop



[8] Workshop participants assigned to the third scenario at the January 2026 workshop developed seven priority mitigations. The participants did not identify their top three mitigations, unlike the groups assigned to the first two scenarios. Four of these mitigations were closely related and addressed different components of detecting user intent. Therefore, we elected to

May 2026 Workshop

Workshop Scope

Four months after the initial workshop, a second workshop was held on May 7, 2026. Whereas the January workshop focused on identifying technical mitigations that could interrupt the development pathways described in each scenario, the May workshop was designed to translate those mitigations into robust, practical, and policy-focused implementation plans. Doing so included identifying metrics for success, necessary resources, critical stakeholders, key outcomes, enablers and dependencies, and potential implementation challenges required to move each mitigation from a conceptual technical intervention to an actionable strategy that relevant stakeholders could pursue.

Workshop Attendees

The workshop's topic and goals necessitated participants with specific expertise. Therefore, we used purposive sampling to identify potential participants who had experience working on policy issues related to AI, biosecurity, or both. We reached out to 45 people with AI policy, biosecurity policy, or AI x Bio policy experience. Some potential participants also had additional work experience in adjacent fields, such as national security policy and foreign policy. This interdisciplinary composition ensured that implementation planning would reflect both the mitigations' technical aspects and the broader policy, governance, and operational realities that would shape the mitigations' adoption.

Fifteen people agreed to participate in the May 2026 workshop under the Chatham House rule.[9] Five participants worked in academia, three worked in AI industry roles, three worked at nonprofit biosecurity organizations, and four worked at think tanks. Four of the participants across these groups had previously worked in either AI or biosecurity policy roles for the federal government.

Four participants who attended the January 2026 workshop also attended the May 2026 workshop. We elected to include these participants to help bridge discussions from the earlier workshop. The participants from the first workshop who participated in the second possessed the technical experience required for the first workshop and the policy experience required for the second.

Workshop Design

Participants were again divided into three groups. Group 1 was assigned to the pandemic disruption scenario (scenario 1), group 2 was assigned to the agroterrorism scenario (scenario 2), and group 3 was assigned to the critical infrastructure scenario (scenario 3). Each group had five participants. This included one returning participant in groups 1 and 2. The third group had two returning participants; we designed the group this way intentionally because the third group had a larger set of mitigations to address. Therefore, we viewed additional continuity with the first workshop as valuable.

Each group received the full scenario narrative and the corresponding mitigations that were produced during the first workshop. Rather than asking participants to revisit whether the mitigations were technically

nest these mitigations under an overarching mitigation ("Improved Threat-Detection Mechanisms") in the previous document (Attal-Juncqua et al., 2026). This mitigation and its four components are described in Chapter 5 of Attal-Juncqua et al. (2026), under Priority Mitigation 4 (Mitigations 4A–4D). Participants from the May 2026 workshop reviewed this mitigation as one mitigation with four parts (as presented in Attal-Juncqua et al. [2026]) rather than as four individual mitigations. The findings under the third scenario's Mitigation 4 ("Improved Threat-Detection Mechanisms") in this document reflect this holistic approach to analyzing that mitigation as one mitigation with four components.

[9] Chatham House, undated.

valuable, we asked participants to focus on how the mitigations could be operationalized in practice, including who would need to act, what steps would be required, and what barriers might prevent implementation.

To make the action items as realistic and useful as possible, we asked groups to consider several implementation dimensions. These included identifying the responsible party or parties, defining key milestones and success metrics, assessing resource needs, specifying enabling conditions and dependencies, describing expected outcomes, and anticipating possible implementation challenges. We came up with these categories to encourage participants to think beyond high-level recommendations and instead develop concrete pathways for action.

Over several hours, participants engaged in structured small-group discussions to develop and refine each action item in detail. Discussions centered on feasibility, investment needs, stakeholder responsibilities, policy mechanisms, and the coordination required across government, industry, academia, and other relevant sectors. Participants also considered how different implementation pathways might vary depending on institutional authority, funding availability, technical maturity, and key actors' willingness to participate.

By identifying actionable paths forward, the groups clarified what would be required to move each mitigation from theory to reality. The resulting implementation plans highlighted not only potential policy levers and operational steps but also the dependencies, trade-offs, and coordination challenges that would need to be addressed for each mitigation to succeed. The second workshop built directly on the outputs of the first, transforming the Mitigation Roadmaps from the first workshop from theoretical technical concepts into more-detailed strategies for real-world action. The following pages present in-depth infographics that synthesize the operational and policy actions that were developed during the second workshop, providing a visual summary for each mitigation's implementation plan.

How Priority Actions Fit Together

In the first workshop, participants identified four key takeaways that represented common themes across the mitigations that were developed in response to the three scenarios:

- Exchanging and integrating information are paramount actions for detecting threats.
- Effectively detecting AI-enabled biological threats requires a multisectoral approach.
- Preventing access to sensitive materials and information disrupts malicious actors.
- Addressing concerns about privacy, civil liberties, and intellectual property is important for implementing AI x Bio mitigations.

The May 2026 workshop built on these findings by examining the institutional, legal, operational, and financial conditions needed to put the priority mitigations into practice.[10]

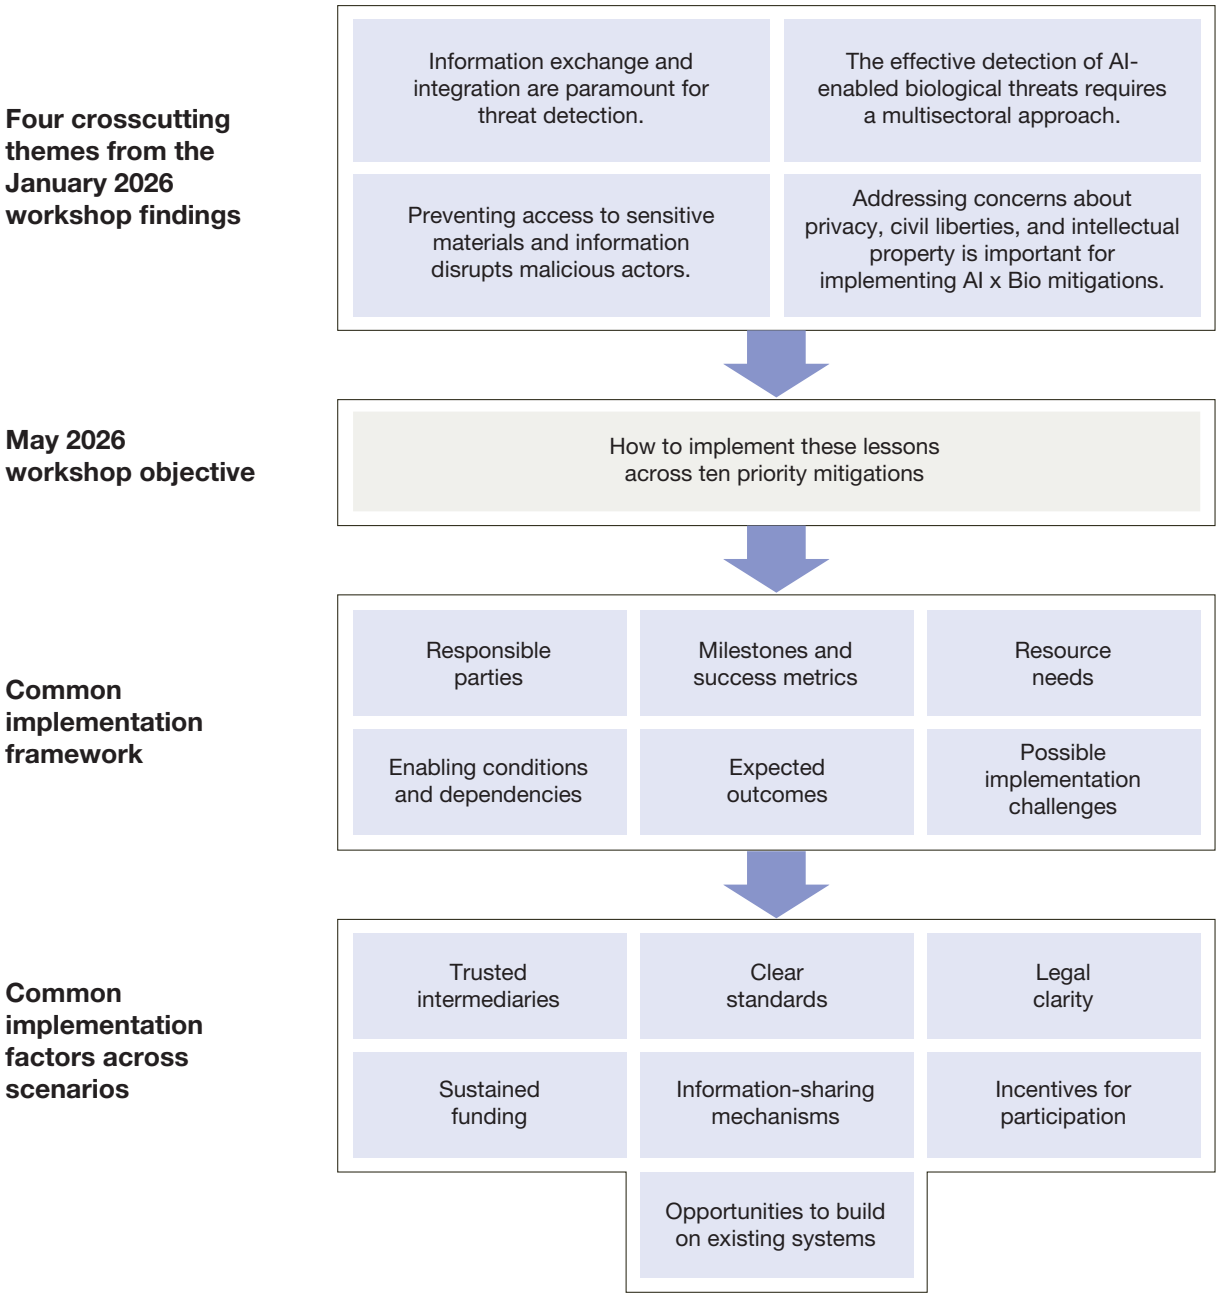
Although the priority mitigations remain organized by scenario because that is how they were developed, the implementation discussions reinforced the relevance of these four takeaways listed in the bullets above. Several mitigations rely on information exchange to verify a user's legitimacy, improve situational awareness, or identify suspicious patterns. Their implementation requires coordination among stakeholders across the AI, biotechnology, research, public health, agriculture, critical infrastructure, law enforcement, and government sectors. Other mitigations seek to disrupt access to sensitive information, tools, or biological materials through managed access, customer screening, or insider-threat awareness. Across these approaches, implementation also requires careful consideration of privacy, civil liberties, intellectual property, and related legal and economic concerns.

[10] Attal-Juncqua et al., 2026.

Participants also identified recurring conditions that could affect implementation, including clear institutional ownership, sustained funding, incentives for adoption, trusted information-sharing arrangements, and appropriate standards and metrics. Figure 2 shows how these common implementation considerations relate to the four key takeaways that participants identified in the first workshop. Figures 3 through 12 present more-detailed implementation and policy considerations for each priority mitigation. An overview of the mitigations is shown in Figure 1.

FIGURE 2

How the Crosscutting Themes from Workshop 1 Informed Implementation Considerations During Workshop 2



Scenario 1: Pandemic Disruption with a Virus

The first scenario examined a well-resourced group creating a virus using AI assistance throughout the process. The group used this virus in a biological attack at a major world event. The group was able to successfully spread the AI-engineered virus across several regions by attacking airports and other locations. We described the group as having similar beliefs and resources as contemporary bioterrorists, such as Aum Shinrikyo.[11] The virus, a new strain of influenza, was modified to enhance transmissibility, reduce vaccine efficacy, and evade the immune system. AI was also used to create disinformation to add to the confusion.

With this scenario, we sought to make participants consider how AI could not only uplift novice actors but also accelerate or magnify the efforts of more-sophisticated groups. The participants developed the resulting mitigations to address areas in which, according to them, the least financial incentive or political will have been invested. These areas were resilience to airborne pathogens (such as clean air initiatives) and AI model development and deployment infrastructure (such as ways to verify biocapable model users, as well as AI agents with more power to detect misuse than existing classifiers). These mitigations all sit at the crossroads of deterrence and response.

Taken together, the scenario's three mitigations illustrate a layered approach: reducing the consequences of respiratory outbreaks, increasing accountability for access to highly enabling AI and biological capabilities, and improving contextual detection of potentially concerning use. The actions address different points in the threat pathway but rely on common implementation conditions, including shared standards, legal clarity, trusted governance arrangements, and sustained public- and private-sector participation.

Mitigation 1: Physical Defenses for Respiratory Pandemics (Personal Protective Equipment and Clean Air)

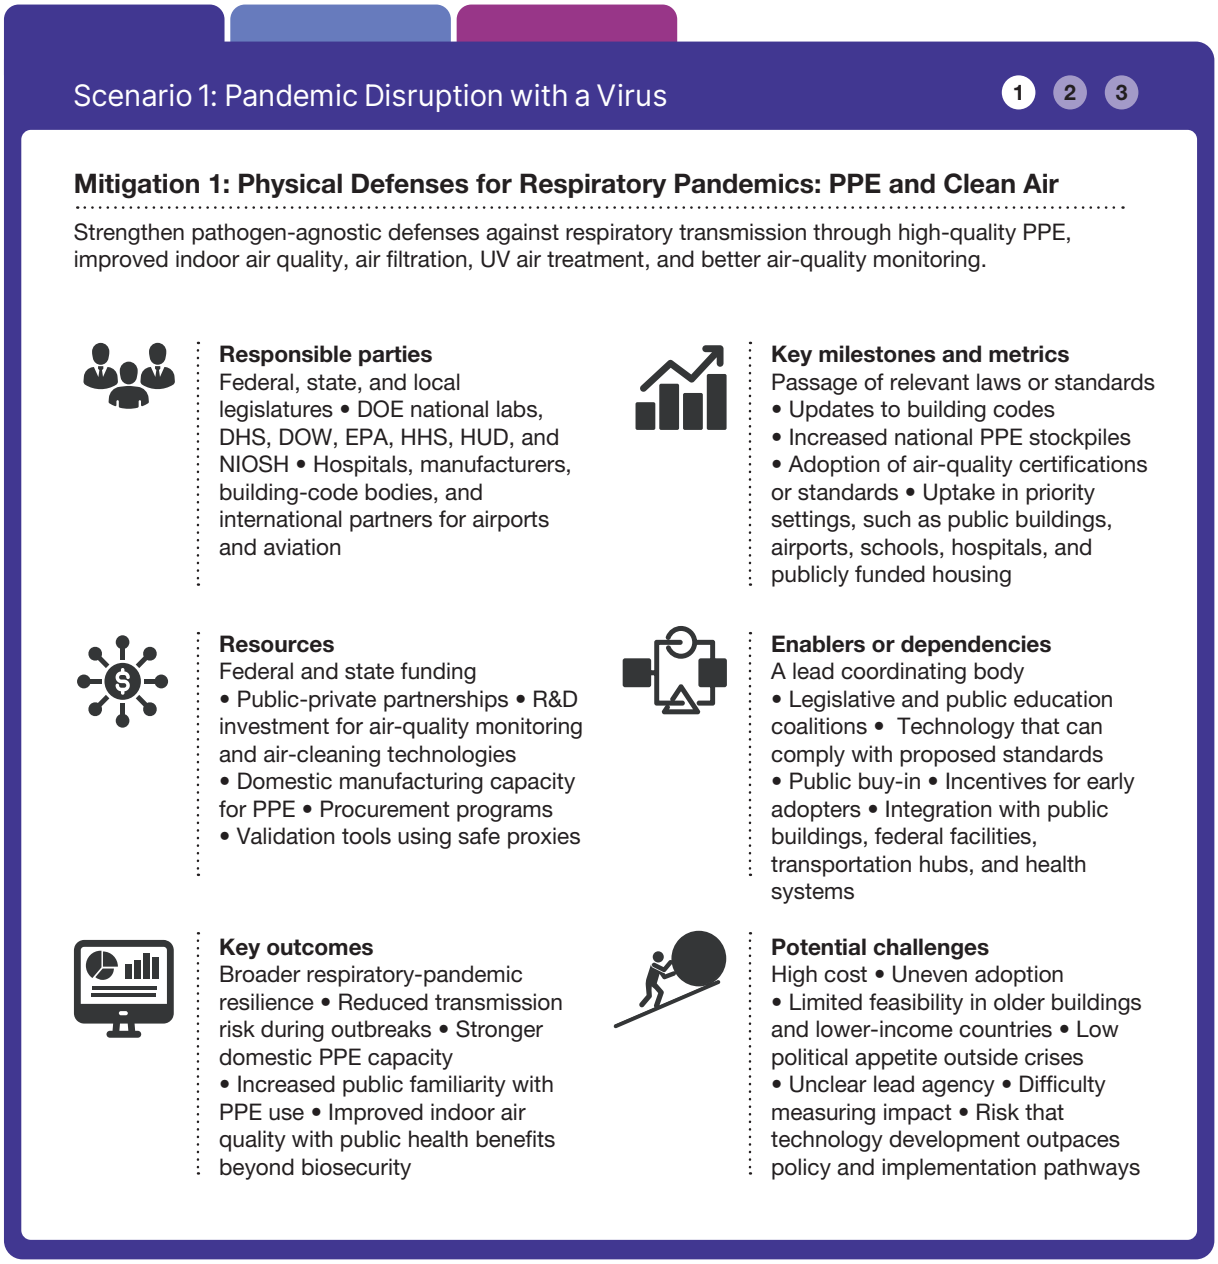
The first of the three proposed mitigations was expanding the availability of PPE and supporting clean air initiatives. These are pathogen-agnostic ways of preventing the spread of airborne pathogens regardless of their origin. The mitigation can serve as a deterrent for malicious actors who would see less appeal in deploying a respiratory disease because infection would not spread at the levels they intended. Participants agreed that PPE and clean air products have significant overlap regarding the funding needed for the research and development and deployment of each. The participating experts saw any expansion of PPE and other clean air technology, such as ultraviolet light products, as critical. However, the participants argued that the clean air technologies would involve more-significant hurdles, such as restrictive building codes, a lack of research funding to further develop technologies, uneven deployment, and government resistance or lack of will. The participants considered the easiest avenues for spreading clean air technologies to be using these technologies in airports and updating building codes to require new or renovated buildings to have such technology. Implementation and policy considerations for this mitigation are displayed in Figure 3.

Mitigation 2: Centralized Identity and Access Infrastructure for Highly Enabling Capabilities

A centralized identity-verification system that enables trusted researchers to access AI models emerged as a novel mitigation concept during the first workshop. Modeled on such existing frameworks as the CLEAR verification program used at airports, the user-verification system would aim to standardize developers' acceptance procedures and reduce friction for researchers navigating individual access requests. The workshop group explored expanding the concept beyond biocapable model access to include cloud labs, contract research organizations, and DNA synthesis companies while also addressing two critical challenges: increasing researcher participation despite privacy concerns about third-party data-sharing and establishing stan-

[11] Aum Shinrikyo was a doomsday cult founded in Japan by Shoko Asahara in the 1980s. It is best known for carrying out the 1995 Tokyo subway sarin gas attack (John Tarangelo, Aurelia Attal-Juncqua, Elika Somani, Derek Roberts, and Kyle Webster, *Protecting Biological Materials and Services from Misuse: Opportunities for Access Monitoring and Control*, RAND Corporation, RR-A4067-1, 2025, p. 6).

FIGURE 3
Scenario 1: Implementation Considerations for Mitigation 1



NOTE: DHS = U.S. Department of Homeland Security; DOE = U.S. Department of Energy; DOW = U.S. Department of War; EPA = Environmental Protection Agency; HHS = U.S. Department of Health and Human Services; HUD = U.S. Department of Housing and Urban Development; NIOSH = National Institute of Occupational Safety and Health; R&D = research and development; UV = ultraviolet. The Department of War is designated the Department of Defense under Public Law 81-216, National Security Act Amendments of 1949.

dardized verification protocols and governance frameworks should multiple companies operate in this space. Participants recognized that, because of the numerous stakeholders involved, an overarching governing body could facilitate cross-company coordination. These layered considerations made this the most extensively developed mitigation that emerged from the scenario discussion. Implementation and policy considerations for this mitigation are displayed in Figure 4.

FIGURE 4
Scenario 1: Implementation Considerations for Mitigation 2



SOURCE: Features information from Open Researcher and Contributor ID (ORCID), homepage, undated.

Mitigation 3: Development of Intent-Monitoring Safeguards ("Guardian Models")

The final mitigation for scenario 1 was the "Guardian Model," an AI agent embedded in AI infrastructure to review prompts, flag misuse or suspicious intentions, and report incidents to developers and enforcement bodies. As an AI agent rather than a classifier, the Guardian can understand context and make decisions, enabling more-effective model security. Participants debated whether to implement a singular Guardian Model with access across all developer platforms or allow individual developers to create their own. A singular model would track individual users across platforms, preventing actors from fragmenting malicious intent across multiple systems—but it would require that companies grant infrastructure access,

raising commercial concerns and questions about adoption rates. Individual Guardians would be tailored to specific models, but doing so would sacrifice the standardization that a centralized system provides. Although this core tension remained unresolved during the workshop, participants recognized the Guardian Model as a potentially powerful safeguarding tool. Implementation and policy considerations for this mitigation are displayed in Figure 5.

FIGURE 5
Scenario 1: Implementation Considerations for Mitigation 3



Scenario 2: Agroterrorism with a Fungus

In this scenario, we examined agricultural biodefense, an often under-resourced area in which biological threats can cause cascading economic and societal harm without directly targeting humans. The scenario focused on the deliberate release of a genetically modified fungal plant pathogen affecting the U.S. wheat supply, and impacts were concentrated in key agricultural regions. The actor in this scenario was characterized as a moderately skilled nonstate group seeking economic disruption and market destabilization rather than mass casualties. Although the group lacked advanced technical expertise, it used AI tools to support genetic modification, pathogen assembly, laboratory workflows, release planning, and disinformation efforts intended to delay detection and response. The scenario resulted in moderate to significant agricultural losses (including reduced crop yields), market instability, and downstream supply chain disruptions.

Participants used the scenario to assess how AI could lower barriers to biological misuse and to identify mitigation strategies for strengthening agricultural biodefense, governance, and information systems. The resulting mitigations focus on agricultural detection, information-sharing, and greater observability of biological supply chains. A crosscutting theme is that many of the mitigation's technical pieces might already exist in some form, but implementation would require policy coordination, farmer adoption, data-sharing agreements, funding, and incentives that make participation worthwhile for agricultural stakeholders.

Together, the three mitigations show how agricultural risk reduction depends on linking early detection, rapid information exchange, and upstream controls on biological inputs. The mitigations' effects are mutually reinforcing—surveillance is more valuable when warning signals can be shared and acted on, while stronger oversight of biological inputs reduces the number of threats that downstream surveillance and response systems must detect. Implementation therefore requires coordination across farmers, researchers, synthesis providers, insurers, agricultural agencies, and other government stakeholders.

Mitigation 1: A Robust, Multifaceted Agricultural Biosurveillance Program

This mitigation would expand agricultural monitoring across agricultural regions using a combination of untargeted environmental sampling, on-farm nongenomic indicators, and drones. Participants emphasized that implementation would depend on several factors, including a well-funded and comprehensive federal policy, state and local participation, crop insurance reform, an image-based notification system, and a "Grand Challenge" to spur innovation and investment. The central challenge obstructing implementation is political resistance to increasing funding and resource allocation and changing crop insurance law. To address this, lobbyists and other proponents of this mitigation can stress the economic case for agricultural biosurveillance, accentuating that less disease leads to higher yields and higher revenues for farmers. Implementation and policy considerations for this mitigation are displayed in Figure 6.

Mitigation 2: Development of an Information-Sharing Platform for Agricultural Threats

This mitigation would create an information-sharing program within a broader agricultural biosurveillance strategy that would rapidly identify and alert stakeholders to emerging agricultural threat events. Participants emphasized that implementation would depend on the development of a request for proposal (RFP) to create the information-sharing platform, and the winning submitter would build the platform and eventually expand it to be available to international partners. The central challenge preventing successful implementation are data issues, including privacy concerns, misinformation and disinformation, and the extensive legal agreements needed at both national and international levels. Implementation and policy considerations for this mitigation are displayed in Figure 7.

FIGURE 6

Scenario 2: Implementation Considerations for Mitigation 1

Scenario 2: Agroterrorism with a Fungus

1 2 3

Mitigation 1: A Robust, Multifaceted Agricultural Biosurveillance Program

Expand agricultural monitoring via environmental sampling, on-farm indicators, and drones; enabled through federal policy, state and local implementation, image-based notifications, crop insurance reform, and a "Grand Challenge" for improved real-time analysis and agnostic threat detection.

Responsible parties
Federal agencies (DOC, USDA, USTR, State Department) • HHS (ASPR, FDA) • White House (DPC, NEC, NSC, and OSTP) • Congress • State departments of agriculture • USDA Risk Management Agency • NASA • Crop Protection Network

Key milestones and metrics
Publication of the federal policy • Positive feedback from relevant stakeholders • Congressional appropriation, potentially through the Farm Bill • Reduction in insurance claims • Reduced crop losses and increased crop yields

Resources
People, time, and effort • State and local agriculture co-ops and extension programs • NASA image feeds • Data-sharing with the Crop Protection Network • Funding for the network to use NASA data to quickly inform farmers • Interstate coordination point of contact

Enablers or dependencies
Farm lobbyists • Existing technical solutions • Funding • Political will to implement

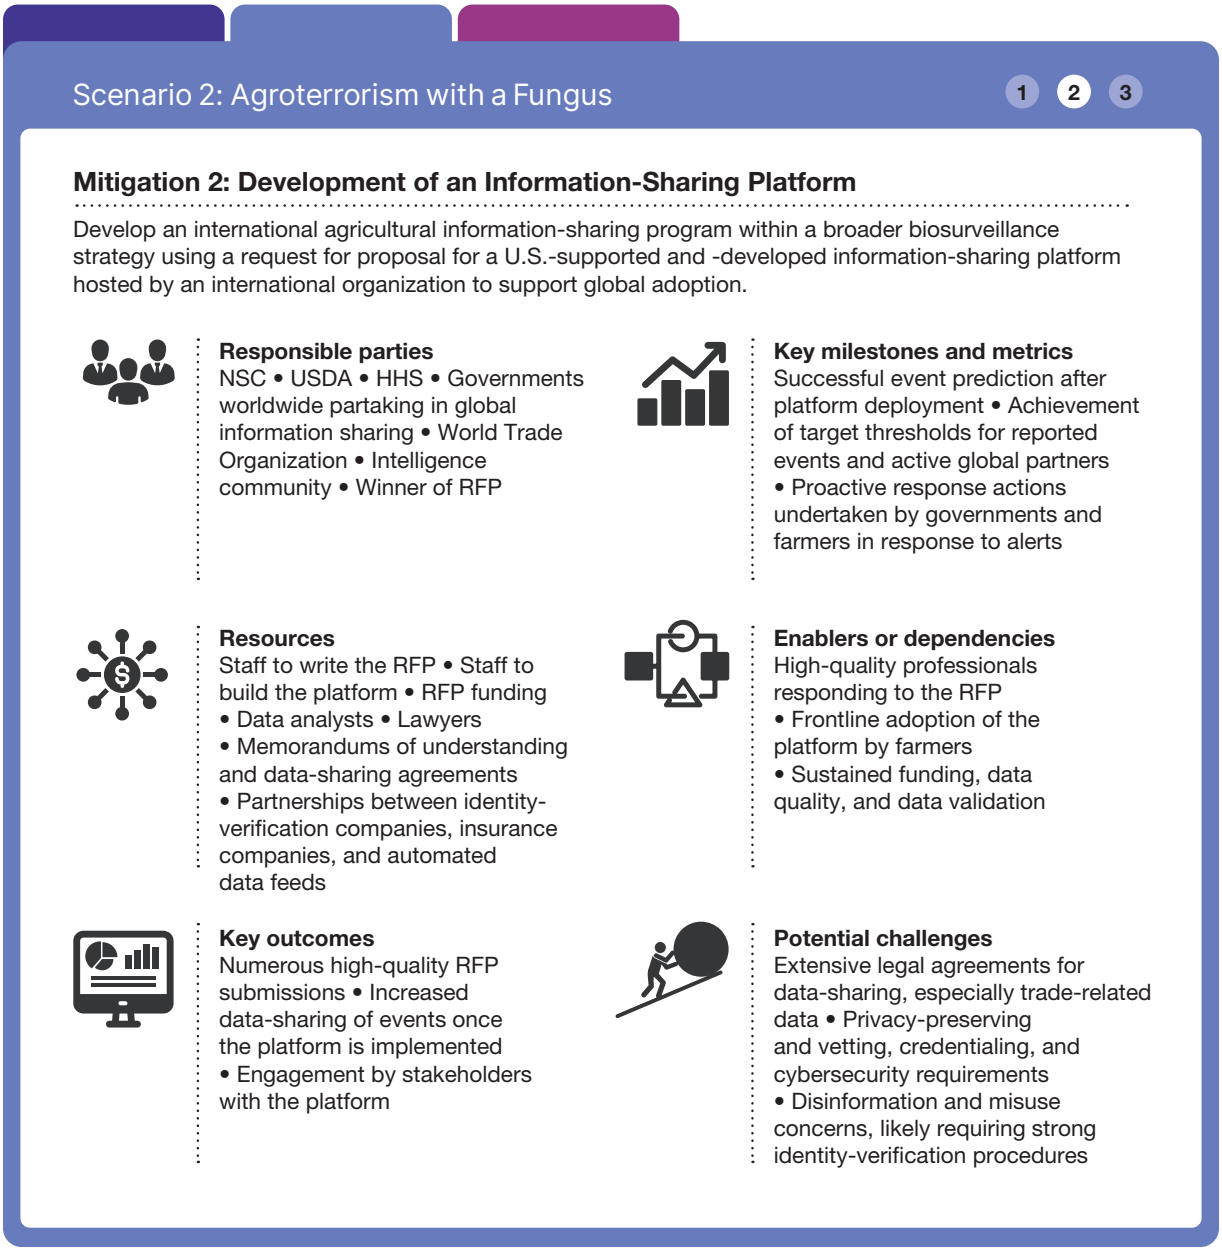
Key outcomes
Reduction in crop loss • Successful state-level implementation • Reduction in insurance claims • Farmers' actions are based on image-based notifications

Potential challenges
Lack of political will to approve funding and resource allocation • Difficulty in coordination and information-sharing • Political resistance to changing insurance law

SOURCE: Features information from Crop Protection Network, homepage, undated.

NOTE: ASPR = Administration for Strategic Preparedness and Response; DOC = U.S. Department of Commerce; DPC = Domestic Policy Council; FDA = U.S. Food and Drug Administration; NASA = National Aeronautics and Space Administration; NEC = National Economic Council; NSC = National Security Council; OSTP = Office of Science and Technology Policy; USDA = U.S. Department of Agriculture; USTR = United States Trade Representative.

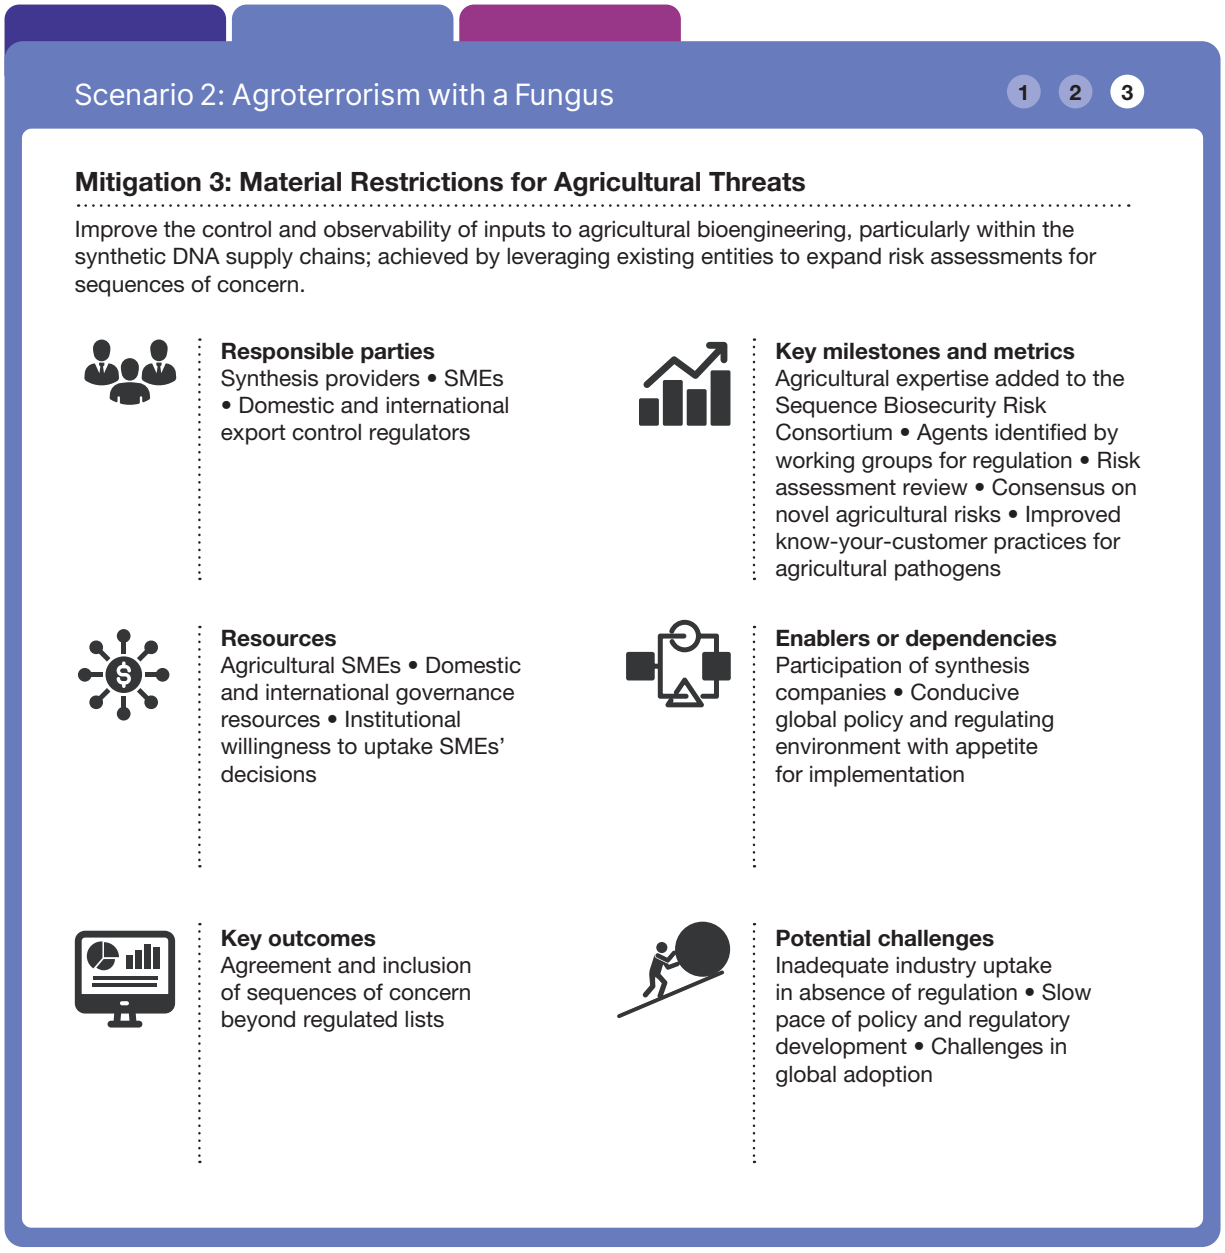
FIGURE 7
Scenario 2: Implementation Considerations for Mitigation 2



Mitigation 3: Material Restrictions for Agricultural Threats

This mitigation would strengthen both the oversight of and visibility into the biological inputs used in agricultural bioengineering, and there would be a particular focus on DNA supply chains. Participants agreed that implementation should build on existing entities, such as the Sequence Biosecurity Risk Consortium, to expand risk assessments for agricultural DNA sequences of concern. Subject-matter experts (SMEs) would contribute to dedicated working groups tasked with identifying, evaluating, and reaching consensus on the agricultural sequences that warrant the greatest scrutiny. The primary implementation challenge for this mitigation is achieving consistent adoption across industry partners, both domestically and internationally, in the absence of regulatory requirements. Implementation and policy considerations for this mitigation are displayed in Figure 8.

FIGURE 8
Scenario 2: Implementation Considerations for Mitigation 3



SOURCE: Features information from Sequence Biosecurity Risk Consortium, homepage, undated.

Scenario 3: Critical Infrastructure Attack with a Bacteria

With this scenario, we explored the potential misuse of AI and biology to target critical infrastructure, specifically the U.S. semiconductor manufacturing sector. The scenario centered on two radicalized insiders who were recruited by a state-sponsored group for financial gain. One insider had access to a semiconductor plant's maintenance systems. The other had master's degree–level training in biology and worked at an academic microbiology lab. The pair leveraged an LLM to efficiently research possible biological agents, deployment methods, and optimal semiconductor plant targets. They also used an LLM to identify a crucial water source in the plant that tolerated some microbial growth, as well as specific design details and temperature

conditions of the water source. After the incident, the state-sponsored group launched a sophisticated disinformation campaign, again aided by LLMs, to exacerbate public mistrust in U.S. manufacturing and divert business to rival national industries. This scenario demonstrated how nonexpert actors with AI assistance could efficiently orchestrate economic disruptions and undermine trust in critical infrastructure through coordinated biosecurity and information attacks.

The mitigations developed in response to this scenario during the January 2026 workshop focused on adapting existing security practices, such as red teaming, insider-threat training, incident reporting, tabletop exercises, and interagency or cross-company coordination. The participants at the May 2026 workshop analyzed these mitigations to develop key implementation features that we present in this section's figures. Participants suggested that the major barriers center more on incentives, legal constraints, liability, trusted information-sharing, and awareness among critical infrastructure stakeholders rather than on the development of entirely new tools.

Together, the four mitigations emphasize the adaptation of existing safety and security systems to account for AI-enabled biological risks. Red teaming can identify institutional vulnerabilities, expanded training can reduce insider risk, de-escalation research can test whether models can interrupt harmful behavior, and improved threat detection can connect suspicious signals across companies and platforms. The mitigations' implementation depends especially on clear responsibility, aligned incentives, legal guidance, trusted information-sharing arrangements, and sustained cooperation among critical-infrastructure operators, AI companies, laboratories, and government agencies.

Mitigation 1: Critical Infrastructure Red Teaming

This mitigation would call for AI-enabled biological threats to be integrated into existing red-teaming exercises within the critical infrastructure sector. Implementation hinges on critical-infrastructure stakeholders understanding the threat that AI-enabled biology poses and recognizing that their sector could be a target. Partner organizations across AI, biotechnology, and security spaces can help relay this messaging, which will help incentivize critical-infrastructure stakeholders to incorporate AI-enabled biology into their red-teaming exercises. These organizations can also support the development of red teaming guidance, inform red-teaming design, and optimize red teaming by tracking lessons learned. Implementation and policy considerations for this mitigation are displayed in Figure 9.

Mitigation 2: Expanded Biosecurity and Insider-Threat Training

This mitigation would expand the biosecurity and insider-threat training that high-containment laboratories use to lower-containment laboratories. Modifying the training for lower-containment laboratories and securing a partner organization to oversee the training's deployment are key for the mitigation's success. Obtaining a financial investment for training modification and deployment is a potential challenge. Additionally, communicating with lower-containment laboratories to secure their buy-in and facilitate their incorporation of the training is a critical implementation factor. Implementation and policy considerations for this mitigation are displayed in Figure 10.

Mitigation 3: Model Safeguards for Actor De-Escalation

This mitigation would explore training AI models to de-escalate malicious and radicalization behaviors. Prior to implementation, studies must first be conducted to identify de-escalation methods and AI models' ability to successfully implement those methods to accomplish de-escalation. Close coordination between AI companies and behavioral health experts would be needed to design these studies and oversee the AI models' implementation if the studies showed promising results. Addressing ethical and liability concerns is paramount if AI models are to be used in de-escalation. Implementation and policy considerations for this mitigation are displayed in Figure 11.

FIGURE 9

Scenario 3: Implementation Considerations for Mitigation 1

Scenario 3: Critical Infrastructure Attack with a Bacteria

1 2 3 4

Mitigation 1: Critical Infrastructure Red Teaming

Enable critical-infrastructure sectors to update existing red-teaming efforts to account for potential AI-enabled biological attacks.



Responsible parties
Government agencies (e.g., Congress, DHS/CISA, DOC/CAISSI, FBI, intelligence community) for national emergency scenario planning
• AI companies to identify security standards • Private security groups working with AI companies • Think tanks to identify global security trends



Key milestones and metrics
Reporting of suspected incidents
• Stakeholder compliance with established security standards
• Recognition by critical-infrastructure stakeholders of AI-enabled biology risks



Resources
Best practices and standards adopted by countries for consideration • Repository of lessons learned



Enablers or dependencies
Strong internal security team protocols are developed and prioritized



Key outcomes
Recognition of AI x Bio critical infrastructure risks • Sectorwide training on risks • Guidelines on integrating biology into red teaming
• Improved detection and industry-law enforcement coordination • Federal guidance for manufacturers and facilities
• Lessons-learned and near-miss library



Potential challenges
Better alignment of incentives to include AI x Bio into red-teaming exercises • Recognition by critical-infrastructure actors that AI x Bio poses a serious threat to their industry

NOTE: CAISSI = Center for Advancing Innovation and Standards for Super Intelligence; CISA = Cybersecurity and Infrastructure Security Agency; FBI = Federal Bureau of Investigation.

FIGURE 10

Scenario 3: Implementation Considerations for Mitigation 2

Scenario 3: Critical Infrastructure Attack with a Bacteria

1 2 3 4

Mitigation 2: Expanded Biosecurity and Insider-Threat Training

Extend biosecurity and insider-threat training beyond high-containment laboratories to include lower-containment laboratories.

Responsible parties
Laboratories • CITI or another accredited body as the distributor of the training • ABSA International to assist with training and training design • DIYbio labs for information-sharing and continuous learning

Key milestones and metrics
Widespread adoption of insider-threat modules across BSL-1 and BSL-2 labs

Resources
Examples from other countries • Modified trainings from higher-containment laboratories

Enablers or dependencies
Ensuring sustained compliance as training scales

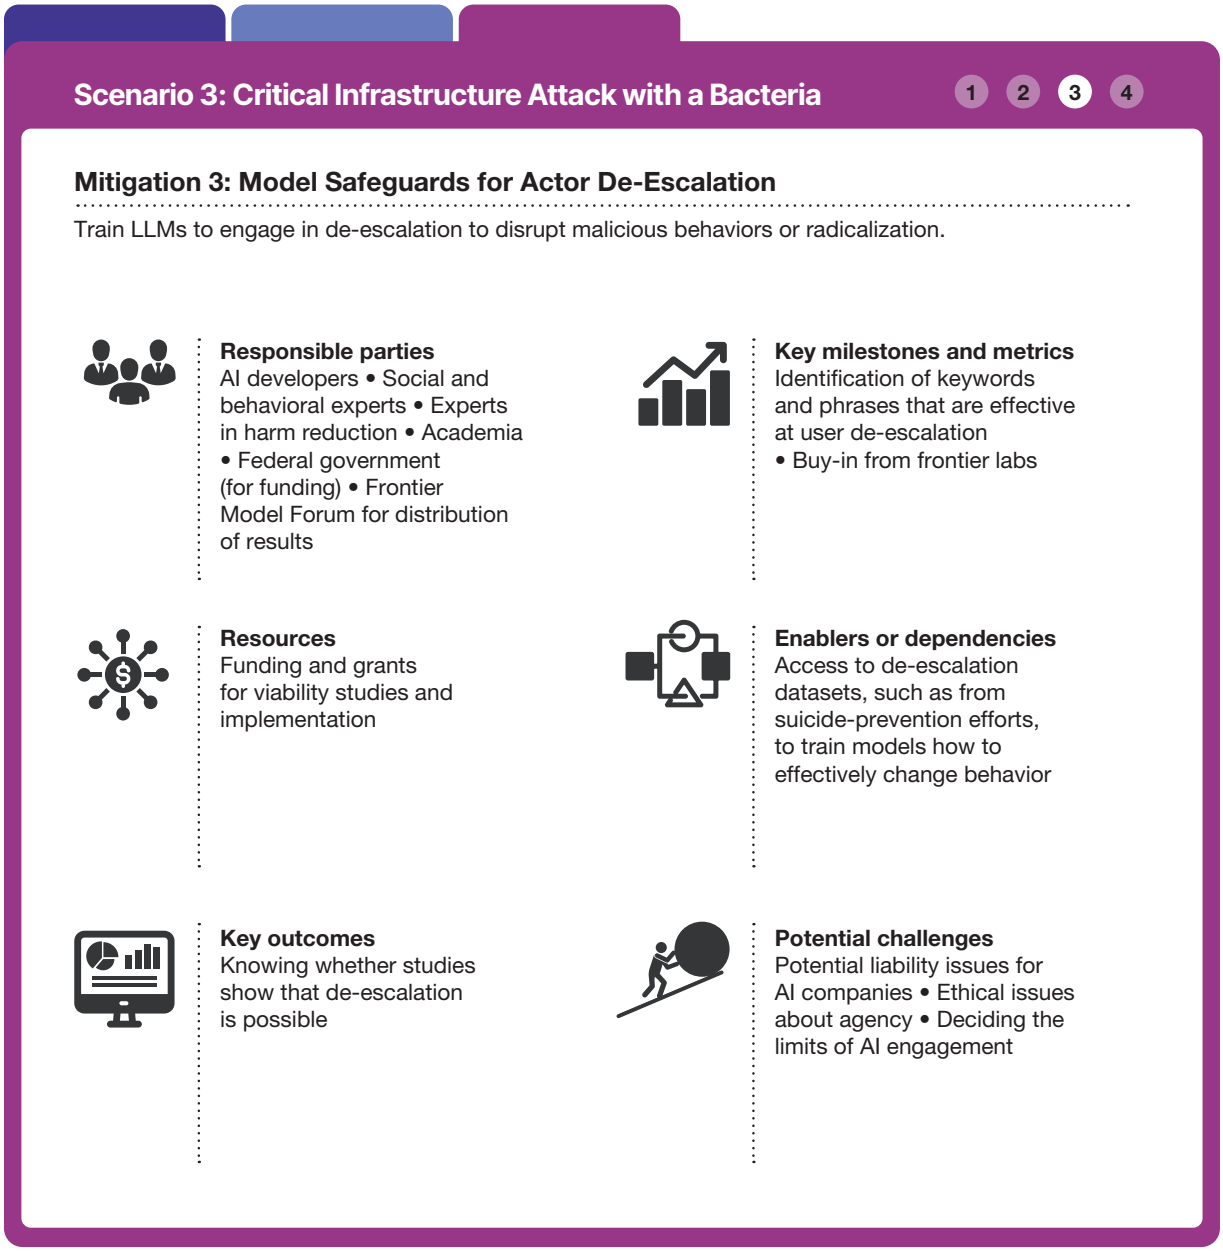
Key outcomes
Incorporating both online and in-person insider-threat training at lower-containment laboratories

Potential challenges
Financial investments for developing training • Implementation of training at BSL-1 and BSL-2 laboratories

SOURCE: Features information from Collaborative Institutional Training Initiative Program, homepage, undated.

NOTE: ABSA = Association for Biosafety and Biosecurity; BSL = biosafety level; CITI = Collaborative Institutional Training Initiative; DIYbio = do-it-yourself biology.

FIGURE 11
Scenario 3: Implementation Considerations for Mitigation 3



SOURCE: Features information from Frontier Model Forum, homepage, undated.

Mitigation 4: Improved Threat-Detection Mechanisms

This mitigation would improve threat detection via four components: (1) inferring intent through subthreshold LLM prompts, (2) exchanging information between LLM companies, (3) aggregating data across platforms and analyzing the data for pattern recognition, and (4) adopting standardized monitoring and reporting criteria across LLM companies. AI companies' adoption of these practices and effective government guidance is essential for the mitigation's success. Additionally, addressing liability, intellectual property, and antitrust concerns related to data-sharing and aggregation is needed for effective implementation. Implementation and policy considerations for this mitigation are displayed in Figure 12.

FIGURE 12

Scenario 3: Implementation Considerations for Mitigation 4

Scenario 3: Critical Infrastructure Attack with a Bacteria

1 2 3 **4**

Mitigation 4: Improved Threat-Detection Mechanisms

Improve threat detection through enhanced intent inference, cross-company information-sharing, cross-platform analysis, and standardized monitoring and reporting.

Responsible parties
AI developers and AI companies • Frontier Model Forum • Civil society, such as think tanks or special interest groups • DOC/CAISSI • FBI • Intelligence community • Congress • DHS/CISA • State governments • Law enforcement

Key milestones and metrics
Definitions of which risks are acceptable of unacceptable • Identification or establishment of a third party for sharing, housing, and analyzing information

Resources
Developers, government agencies, and regulators for tabletop exercises • Experts in government roles • Sustained funding

Enablers or dependencies
Exploring models from other sectors, including the NCMEC CyberTipline + CENP, CAASD, and ICAO ADREP system

Key outcomes
Enhanced information-sharing and threat detection • Enhanced ability to ensure the validation of the detected incident is successful or is set up for success

Potential challenges
Liability and intellectual property concerns • Navigating antitrust laws that affect AI companies' information-sharing practices

SOURCES: Features information from Federal Bureau of Investigation, "Child Exploitation Notification Program," webpage, undated; Frontier Model Forum, undated; International Civil Aviation Organization, "Accident/Incident Data Reporting (ADREP)," webpage, undated; MITRE, "CAASD," webpage, undated; National Center for Missing & Exploited Children, "CyberTipline," webpage, undated.

NOTE: ADREP = Accident/Incident Data Reporting; CAASD = Center for Advanced Aviation System Development; CENP = Child Exploitation Notification Program; ICAO = International Civil Aviation Organization; NCMEC = National Center for Missing & Exploited Children.

Conclusion

Participants' discussions of the three scenarios underscored that many AI-enabled biological risk mitigations depend as much on governance, incentives, and adoption as they do on technical capability. Although the mitigations that participants developed during the first workshop were ways to potentially interrupt AI-enabled biological threat pathways, the second workshop highlighted the institutional, legal, operational, and financial conditions that would be necessary to make those mitigations a reality. Despite each scenario's specific differences, several key findings emerged that were broadly applicable across all three threat pathways, and these findings point to concrete actions that society can take to reduce these risks. Participants repeatedly identified the need for trusted intermediaries, clear standards, legal clarity, sustained funding, information-sharing mechanisms, and participation incentives in threat-reduction programs.

The discussions also showed that the most practical pathways are likely to build on existing systems rather than rely solely on new institutions or stand-alone technical solutions. Participants identified specific opportunities to integrate AI-enabled biological risk mitigations into public health infrastructure, agricultural biosurveillance networks, crop insurance mechanisms, and AI safety evaluations, among others. These existing systems provide useful foundations for implementation because they already involve relevant stakeholders, authorities, operational practices, or incentive structures, lowering both the institutional barriers and initial costs to act on this workshop's outputs.

However, participants recognized that several crosscutting barriers could limit implementation, including unclear institutional responsibility, uneven adoption, privacy and civil liberties concerns, liability and intellectual property issues, resource constraints, global coordination challenges, and limited political will outside periods of crisis. The discussions of the implementation enablers and barriers suggest that moving the mitigations toward implementation will require stakeholders to identify responsible parties, build on relevant existing systems, establish appropriate standards and success metrics, address legal and information-sharing barriers, and align funding and incentives for participation. The importance and form of these requirements will vary across mitigations and institutional contexts.

We do not establish a single implementation sequence or suggest that all mitigations should be pursued simultaneously. Rather, we provide a basis for policymakers, industry actors, research institutions, and other stakeholders to (1) assess which approaches are most feasible within their authorities and operating environments, (2) identify opportunities for coordination across mitigations, and (3) develop more-detailed implementation plans. Future work should assess the relative feasibility, expected benefits, trade-offs, resource requirements, and stakeholder incentives associated with each mitigation and identify concrete near- and farther-term opportunities for action.

APPENDIX

Scenario Summaries

Earlier in these conference proceedings, we provided a short description of the scenarios from the January 2026 workshop. These scenarios informed the technical mitigations that participants developed at that workshop. Implementation and policy considerations for those mitigations were assessed during the May 2026 workshop, which is the focus of this document.

To provide further context about the scenarios, we reproduce the scenario descriptions from our January 2026 workshop conference proceedings in this appendix.[1] These scenarios are summaries of the scenarios that we presented to participants. The scenarios are summarized in the prior conference proceedings and in this document because of infohazard concerns. However, participants at both the January 2026 and May 2026 workshops had access to the scenarios' full versions.

Scenario 1: Pandemic Disruption with a Virus Scenario

We designed a scenario to explore the risks posed by nonstate actors seeking to leverage AI to cause mass casualty harm with a biological agent, particularly by blurring the line between a natural outbreak and a deliberate attack. The scenario focused on the ways in which AI-enabled tools might lower traditional barriers to misuse across the biological threat life cycle while simultaneously complicating the detection of the threat, as well as the investigation into, attribution of, and response to an outbreak.

The scenario centered on the emergence of a novel influenza A strain in the early 2030s. Although we considered the hypothetical strain to be genetically similar to seasonal influenza A strains in circulation, this strain exhibited characteristics that significantly challenged containment efforts, including enhanced transmissibility, atypical disease dynamics, and reduced responsiveness to existing medical countermeasures. In the scenario, the near-simultaneous appearance of severe outbreaks across multiple global locations created uncertainty about the origin of the event, complicating early decisionmaking and communication with the public.

We characterized the scenario's actor as a highly capable, well-resourced nonstate group motivated by a millenarian ideology and intent on causing catastrophic physical, economic, and psychological harm. Although the group was not state-sponsored, we told participants to assume that the group had access to advanced expertise and infrastructure through decentralized networks. AI played a central enabling role by accelerating the group's research and development activities, supporting operational coordination, and amplifying disinformation efforts designed to undermine trust in institutions and exacerbate public panic.

Participants examined how BTs [biological tools] and AI-enabled data access and automation could compress timelines and expand the feasible threat space for nonstate actors, all while AI-driven information operations further degraded situational awareness during the crisis. The results of the outbreaks were severe global health impacts, overwhelmed medical systems, cascading economic disruptions, and long-term political

[1] Attal-Juncqua et al., 2026.

instability. The workshop group used this scenario to identify mitigation strategies focused on governance, information access controls, biosecurity oversight, and resilience measures that could reduce the likelihood or impact of similar AI-enabled biological threats in the future.

Scenario 2: Agroterrorism with a Fungus Scenario

We designed this scenario to examine an often overlooked and under-resourced domain of biodefense: threats to agricultural systems, which can generate cascading economic and societal harm even when humans are not the direct target. With this scenario, we explored how disruptions to critical food supplies could produce widespread instability, and we highlighted systemic vulnerabilities that sit outside traditional public health preparedness frameworks. The scenario centered on the deliberate release of a genetically modified fungal plant pathogen that possessed enhanced resistance relative to naturally occurring strains, enabling it to overcome common crop defenses and spread across multiple agricultural sites. The target was the U.S. wheat supply, and impacts were concentrated in key agricultural regions.

We characterized the actor as a moderately skilled nonstate group motivated by economic disruption and the destabilization of markets rather than mass casualties. Although the group lacked elite technical capacity, it leveraged AI tools to compensate for gaps in expertise, streamline laboratory processes, and support decisionmaking across the biological risk chain. For this scenario, the actor used AI to assist with genetic modification and pathogen assembly, improve efficiency in laboratory workflows, inform the timing and placement of release sites, and amplify disinformation intended to delay detection and response.

The scenario's outcome was moderate to significant agricultural losses leading to broader economic disruption through reduced crop yields, market instability, and downstream supply chain effects. Participants used this scenario to explore how AI-enabled capabilities could lower traditional barriers to misuse. Participants also identified mitigation strategies aimed at strengthening resilience across agricultural biodefense, governance, and information systems.

Scenario 3: Critical Infrastructure Attack with a Bacteria Scenario

We designed this scenario to investigate how AI x bio could be used to attack critical infrastructure. With this scenario, we explored how actors with access to biological materials and a desired target site could execute a smaller-scale, targeted attack that triggers economic disruptions and distrust in U.S. manufacturing. The scenario involved two radicalized nonstate insider threats. These two individuals were recruited by a state-sponsored group in exchange for financial compensation. One insider worked at a semiconductor manufacturing plant. The other had master's level training in biology and worked in an environmental microbiology laboratory at a medium-sized academic institution.

Although one of the scenario's actors had a moderate skill set in biology and the other had access to the semiconductor facility's maintenance systems, they decided to use generative AI to help them more efficiently parse through options for the agent, deployment method, and target site within the plant. Through this process, the actors identified a water source within the plant that is critical to the plant's manufacturing processes and tolerates a degree of microbial growth. The actors also learned other key facts about the design of the towers holding the water, where the optimal exposure point would be, and the water's temperature. The actor with a background in biology knew that several bacteria within his workplace readily form biofilms at temperatures similar to those in the towers. He used an LLM to confirm his knowledge and swiftly compared the agents available to him, leading him to select a biofilm-forming bacteria with a BSL-1 (i.e., the bacteria are safer for humans to handle than bacteria with other BSLs). The actor clandestinely procured the agent

from the workplace, cultured it at home using basic laboratory materials purchased online, and provided the cultured bacteria to the insider working at the semiconductor manufacturing plant.

The insider at the plant then contaminated several towers housing the target water source. Although biosensors ultimately alerted the plant's staff that microbial growth exceeded permitted levels, the simultaneous notification across several towers resulted in the plant ceasing operations for 24 to 48 hours while the incidents were investigated. This resulted in millions of dollars lost and an entire batch of semiconductors being destroyed. As word of the incident became public, the state-sponsored group that recruited the two insiders deployed a disinformation campaign developed with an LLM's assistance, which included giving strategies for spreading disinformation on social media. The campaign was designed to cause wider economic disruptions by sowing distrust in U.S. industries and driving business to U.S. competitors.

Abbreviations

ABSA	Association for Biosafety and Biosecurity
ADREP	Accident/Incident Data Reporting
AI	artificial intelligence
AI x Bio	the intersection of artificial intelligence and biology
ASPR	Administration for Strategic Preparedness and Response
BSL	biosafety level
BT	biological tools
CAASD	Center for Advanced Aviation System Development
CAISSI	Center for Advancing Innovation and Standards for Super Intelligence
CENP	Child Exploitation Notification Program
CISA	Cybersecurity and Infrastructure Security Agency
CITI	Collaborative Institutional Training Initiative Program
DHS	U.S. Department of Homeland Security
DIYbio	do-it-yourself biology
DNA	deoxyribonucleic acid
DOC	U.S. Department of Commerce
DOE	U.S. Department of Energy
DOW	U.S. Department of War
DPC	Domestic Policy Council
EPA	Environmental Protection Agency
FBI	Federal Bureau of Investigation
FDA	U.S. Food and Drug Administration
HHS	U.S. Department of Health and Human Services
HUD	U.S. Department of Housing and Urban Development
ICAO	International Civil Aviation Organization
LLM	large language model
NASA	National Aeronautics and Space Administration
NEC	National Economic Council
NCMEC	National Center for Missing & Exploited Children
NSC	National Security Council
NIOSH	National Institute of Occupational Safety and Health
OSTP	Office of Science and Technology Policy
ORCID	Open Researcher and Contributor ID
PPE	personal protective equipment
R&D	research and development
RFP	request for proposal
SME	subject-matter expert
USDA	U.S. Department of Agriculture
USTR	United States Trade Representative

References

Attal-Juncqua, Aurelia, John P. Tarangelo, Evan Davison Kotler, Casey Aveggio, Marie Burns, Barbara Del Castello, Tyler Burns, Steph Guerra, and Saskia Popescu, *Building Resilience Against Artificial Intelligence–Enabled Biological Threats: Conference Proceedings from a RAND and Helena Workshop*, RAND Corporation, CF-A4954-1, 2026. As of July 31, 2026:
https://www.rand.org/pubs/conf_proceedings/CFA4954-1.html

Chatham House, "Chatham House Rule," webpage, undated. As of October 4, 2026:
https://www.chathamhouse.org/about-us/chatham-house-rule

CITI Program, homepage, undated. As of August 13, 2026:
https://about.citiprogram.org/

Crop Protection Network, homepage, undated. As of August 13, 2026:
https://cropprotectionnetwork.org/

Federal Bureau of Investigation, "Child Exploitation Notification Program," webpage, undated. As of August 13, 2026:
https://www.fbi.gov/how-we-can-help-you/victim-services/cenp

Frontier Model Forum, homepage, undated. As of August 13, 2026:
https://www.frontiermodelforum.org/

International Civil Aviation Organization, "Accident/Incident Data Reporting (ADREP)," webpage, undated. As of August 13, 2026:
https://www.icao.int/safety/AIG/reporting-and-adrep-data

Luckey, David, Sara Duhachek Muggy, Taylor Frey, David Stebbins, Tracey Rissman, Bianca Espinosa, Daniel Tapia, Greg McKelvey, Jr., Neeti Pokhriyal, Joseph Nelson Dawson, Sara Hughes, Morgan Sandler, Rushil Bakhshi, Marta Kepe, Geoffrey Kirkwood, Sarah W. Denton, David DeSmet, Minami Makino, Ella Guest, Sina Beaghley, Suzanne Genc, Michael Miller, Skye A. Miner, Barbara Del Castello, Forrest W. Crawford, Jeffrey Lee, Clay Strickland, Sunny D. Bhatt, John Vahedi, Lydia Grek, Vanya Barrer, Ramiro Insuasti, Jr., Jack Lashendock, Derek Roberts, Aleksandr Esparza Hartunian, Shannon Walsh, Will Shumate, Elliott Brennan, Tyler Liggett, Kara Jia, Ajay K. Kochhar, James Smith, and James Ryseff, *Mitigating Risks at the Intersection of Artificial Intelligence and Chemical and Biological Weapons*, Homeland Security Operational Analysis Center operated by the RAND Corporation, RR-A2990-1, 2025. As of August 25, 2026:
https://www.rand.org/pubs/research_reports/RRA2990-1.html

MITRE, "CAASD," webpage, undated. As of August 13, 2026:
https://www.mitre.org/our-impact/rd-centers/center-advanced-aviation-system-development

National Center for Missing & Exploited Children, "CyberTipline," webpage, undated. As of August 13, 2026:
https://www.missingkids.org/gethelpnow/cybertipline

ORCID, homepage, undated. As of August 13, 2026:
https://orcid.org/

Sequence Biosecurity Risk Consortium, homepage, undated. As of August 13, 2026:
https://sbrc.bio/

Tarangelo, John, Aurelia Attal-Juncqua, Elika Somani, Derek Roberts, and Kyle Webster, *Protecting Biological Materials and Services from Misuse: Opportunities for Access Monitoring and Control*, RAND Corporation, RR-A4067-1, 2025. As of October 4, 2026:
https://www.rand.org/pubs/research_reports/RRA4067-1.html

Wheeler, Nicole E., "Responsible AI in Biotechnology: Balancing Discovery, Innovation and Biosecurity Risks," *Frontiers in Bioengineering and Biotechnology*, Vol. 13, February 5, 2025.

About the Authors

John P. Tarangelo is a senior technical analyst at RAND. His research focuses on health security, biodefense, and biosecurity policy. Tarangelo holds an M.S. in public health microbiology and emerging infectious diseases.

Aurelia Attal-Juncqua was a senior policy researcher at RAND when this document was drafted. Her research focuses on biosecurity and emerging technology. Attal-Juncqua holds a Dr.P.H. in health security.

Evan Davison-Kotler is the director of research and health security at Helena. His work focuses on biosecurity policy, biological risk enabled by artificial intelligence and pandemic preparedness with an emphasis on bridging technical analysis and U.S. and international governance efforts. Davison-Kotler holds a B.Sc. (Hons.) in biology and is completing a Ph.D. in biomedical engineering.

Marie Burns is the communications lead for biosecurity at Helena. Her work focuses on biosecurity and risks associated with emerging technologies. Burns holds a Ph.D. in immunology.

Casey Aveggio is a technical analyst at RAND. Her work focuses on global health security, zoonotic spillover at the human-wildlife interface, and reducing the risks of deliberate biological threats. Aveggio holds an M.S. in biohazardous threat agents and emerging infectious diseases.

Barbara Del Castello is an associate scientist at RAND. Her work focuses on biosecurity and examining how emerging technologies influence biological risk, bioterrorism, and science diplomacy. Del Castello holds a Ph.D. in genetics.

Tyler Burns is the research lead at Helena. His work focuses on the intersection between AI and biology. Burns holds a Ph.D. in cancer biology.

Steph Guerra is a senior research resident at RAND. She leads research at the intersection of AI, biotechnology, and national security. Guerra holds a Ph.D. in biological and biomedical sciences.

Saskia Popescu is the policy lead for health security at Helena. Her work focuses on the intersection of science, security, and policy and the enhancement of bioresilience, readiness to infectious disease outbreaks, and global health security. Popescu holds a Ph.D. in biodefense.

David Glickstein is research communications analyst at RAND. He supports research on emerging risks related to technology, biology, AI, national security, and space. He holds an M.P.P. with a concentration in science and technology policy.