
OSINT BioGovernance Tool

Zhamilia Klycheva, Ph.D. Edidiong James Shubhankar Dharmadhikari
SPAR SPAR SPAR

Erik Leklem

SPAR

With
Apart Research

Abstract

This project proposes tools for faster identification, classification, and response to AI biosecurity “warning shots”—near-miss events that reveal catastrophic risk—so as to translate those events into governance action by relevant policy, intelligence, and other actors. We use a formal definition of warning shots and associated criterion to analyze 21 global biosecurity events, finding that warning shots are typically recognized but rarely converted into binding governance. We propose analysts and governance actors use a Governance Conversion Framework to improve future AI biosecurity warning shot responses. We develop an AI biosecurity event dashboard and set of analytic tools to help biosecurity stakeholders monitor world events, categorize emergent biosecurity risks, and trigger faster response with relevant government and industry actors.

1. Introduction

The term "warning shot" is gaining traction in the AI safety community. As Stuart Russell noted at the NDTV Ind.AI Summit, a "Chernobyl-scale disaster" may be the best-case scenario — an incident alarming enough to trigger regulatory action without crossing into full catastrophe. Yet no consensus exists on what a warning shot actually is, or how to identify one before it escalates.

Identifying warning shots matters because they carry the potential to drive binding regulatory action. Hence, there are several key questions this project aims to address:

1. Are there any incidents that regulators in the biosecurity space should pay attention to?
 - a. How dangerous are such incidents?
2. Are these incidents likely to prompt a regulatory action?
 - a. Is the current government response sufficient?
 - b. What regulatory gaps do such incidents reveal?

Two problems motivate this work: analysts often lack the tools to recognize AIxBio events as warning shots before they escalate, and even when they do, the governance response frequently stalls or fails.

Our project addresses this directly: by combining a *warning shot classification methodology*, a *biosecurity risk score*, and the *Governance Conversion Framework (GCF)*, it equips analysts and policymakers to identify AIxBio warning shots and assess whether the current governance response is sufficient and if not, why it would fail. It also helps individuals and organizations improve their ability to understand, process, and categorize AIxBio events in order to perform their roles and duties more quickly, more clearly, and with greater chance of addressing the emerging risk.

Our main contributions are:

1. Warning Shot definition & criteria
2. Biosecurity Risk Assessment Score & criteria
3. Governance Convergence Framework
4. Biosecurity Governance Tracker Dashboard

Target Audience

This tool serves a broad range of stakeholders — from U.S. executive agencies and legislative oversight bodies to frontier AI companies, international organizations, and civil society researchers — each of whom engages with AIxBio events differently depending on their role and authority. Tier 1 events primarily concern actors with statutory response responsibilities, such as DHS or WHO, while Tier 2 and 3 events are most actionable for intelligence communities, early warning systems, and industry actors positioned to intervene before a signal becomes a warning shot.

2. Related Work

We are not aware of a state-of-the-art method or tool that allows people to determine whether an event is a warning shot or not. While some methods (e.g. open source intelligence augmented by AI systems) help identify people, capability break-outs, violations of laws or regulations, these do not appear to be able to scale insights through a prism of warning shots that could necessitate expert and government interaction and response. Additionally, we are not aware of any methods or tools that provide parsimonious ways to track, monitor, and report on warning shots that are in the governance convergence framework process.

Additionally, our tools help policy-makers learn from historical warning shot cases, with a specific distillation of lessons at each failure mode in the governance conversion framework. These tools help identify what to avoid at each phase, or what to do, in order for a warning shot to be progressed to the next set of actors in the “chain of conversion”.

3. Methods

This project heavily relies on the [SPAR program Warning Shots](#) project methodological foundation that is currently being finalized by the authors and is not yet available as a publication. It then tailors it specifically for the biosecurity domain and develops a governance analysis tool based on the Open Source Intelligence (OSINT) data.

First, we conduct a search according to the warning shot methodology (C1-C5) developed within the SPAR project and classified as tier 1, 2 and 3 cases.

Second, each individual warning shot incident is then assigned a risk score (0-100) that was developed specifically for the biosecurity domain within the scope of this hackathon.

Third, each incident is then passed through a Governance Conversion Framework (GCF) and analyzed, whether it passes through each stage and results in successful governance. Figure 1 below demonstrates this process.

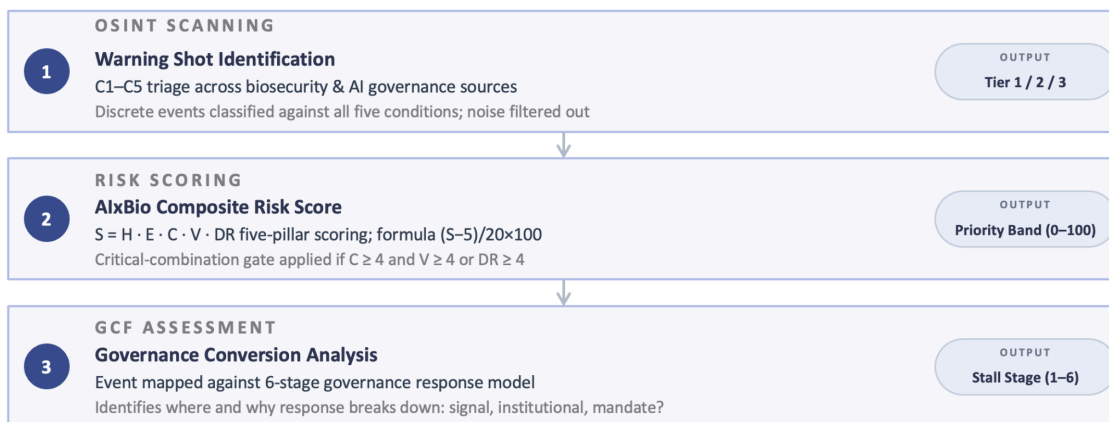


Figure 1, OSINT BioGovernance Scanning Process

Scope of work

The full dashboard, data generation and synching and the Biosecurity Risk Assessment Score were created during the hackathon weekend. The governance conversion framework, Warning Shot definition (incl. C1-C5) are the existing work from the SPAR project.

Warning Shot Definition

A biosecurity warning shot is a specific event that reveals a pre-existing vulnerability in a system, organization, or society. To qualify as a Tier 1 warning shot, an incident must satisfy five conditions: it must be a discrete, bounded event with an identifiable onset and terminus (C1), expose a pre-existing structural weakness in screening, biosafety, governance, or surveillance (C2), present an identifiable, step-by-step pathway by which biological catastrophe could plausibly occur (C3), have had that catastrophic pathway epistemically accessible at the time of the event (C4), and stop short of the catastrophe it reveals as possible (C5). A Tier 2 warning signal satisfies C1 and at least two further conditions but fails one — most commonly C3 or C5. A Tier 3 contextual indicator fails C1 or multiple conditions simultaneously and does not qualify as a warning shot, but shapes the structural landscape in which warning shots occur. Full condition definitions and classification criteria are provided in the Appendix.

Risk Assessment Score

The biosecurity risk assessment score measures urgency across five pillars, each scored 1–5: Hazard (H), Exposure (E), Capability (C), Vulnerability (V), and Detection & Response Deficit (DR).

Composite score

$S = H + E + C + V + DR$, then normalized to a 0–100 scale via $Score = (S - 5) / 20 \times 100$, producing four priority bands:

- Low (0–29, routine review)
- Moderate (30–59, active surveillance)
- High (60–84, priority monitoring)
- Critical (85–100, immediate escalation)

A critical-combination gate applies when $C \geq 4$ AND $V \geq 4$, or $C \geq 4$ AND $DR \geq 4$, flooring the composite score at 60 to ensure that dangerous AI-capability configurations paired with governance or detection failures are never understated by the additive sum alone. Full pillar definitions and scoring rubrics are provided in the Appendix.

Governance Conversion Framework (GCF)

The Governance Conversion Framework (GCF) was developed by analyzing 13 Tier 1 historical warning shot incidents across biosecurity, cybersecurity, military, and nuclear domains to map how an incident translates into effective governance action. The framework identifies six sequential stages: (1) Incident, (2) Signal Recognition, (3) Institutional Engagement, (4) Governance Output, (5) Binding Mandate, and (6) Effective Governance — the rarest outcome, where a mandate is actively enforced and the structural vulnerability durably addressed. Two pathway modifiers can interrupt or redirect conversion: M1 (Competition), where race dynamics between states or industry actors delay binding mandates at the Stage 4→5 and 5→6 transitions; and M2 (Adversarial Attribution), where attribution to a state adversary shifts governance toward deterrence rather than civilian regulation. See Figure 2 below:



Figure 2, Governance Conversion Framework Stages

4. Results

The BioGovernance dashboard enables structured analysis of 21 biosecurity-relevant events spanning 1967 to 2026, integrating classification, risk assessment, and governance progression into a unified analytical interface. Beyond static visualization, the system supports comparative and pattern-based analysis, allowing users to examine how emerging signals relate to historical warning shot events and how they are likely to evolve. Several key findings emerge.

First, Tier 1 warning shots remain relatively rare, with most events classified as Tier 2 (warning signals) or Tier 3 (contextual indicators). However, the dataset suggests that many Tier 2 events exhibit partial characteristics of warning shots, indicating that escalation pathways may be under-recognized rather than absent.

Second, the risk distribution shows a concentration in moderate and high categories, with critical cases typically emerging where capability and vulnerability intersect. This is particularly visible in cases involving synthetic biology, dual-use research, and AI-enabled bio risks.

Third, and most significantly, the Governance Conversion Framework reveals a consistent structural failure pattern. The majority of cases stall at Stage 3 (Institutional Engagement). While signals are identified and investigated, they rarely progress into binding mandates or durable governance outcomes.

Importantly, the dashboard enables cross-case comparison and similarity matching, allowing users to identify how new or emerging signals align with historical events. This supports a deeper insight: many contemporary signals are not novel, but rather recurrent patterns of previously observed failures.

The system also highlights how Tier 2 and Tier 3 events evolve over time, suggesting that early-stage indicators could be used more effectively for preemptive intervention. However, current governance systems do not systematically act at these earlier stages.

Figure 3 (Dashboard Overview) illustrates the distribution of cases across time, tier classification, and risk level, showing increased clustering in recent years. Figure 4 (Case Detail View) demonstrates how individual cases are evaluated across warning shot criteria, risk scoring dimensions, and governance progression, enabling traceability of analytical decisions. Figure 5 (GCF Distribution) highlights the concentration of cases at Stage 3, reinforcing the conclusion that governance breakdown occurs not at detection, but at conversion.

Overall, the results indicate that:

- signals are increasingly detectable
- patterns are historically consistent
- the core failure lies in translating detection into action

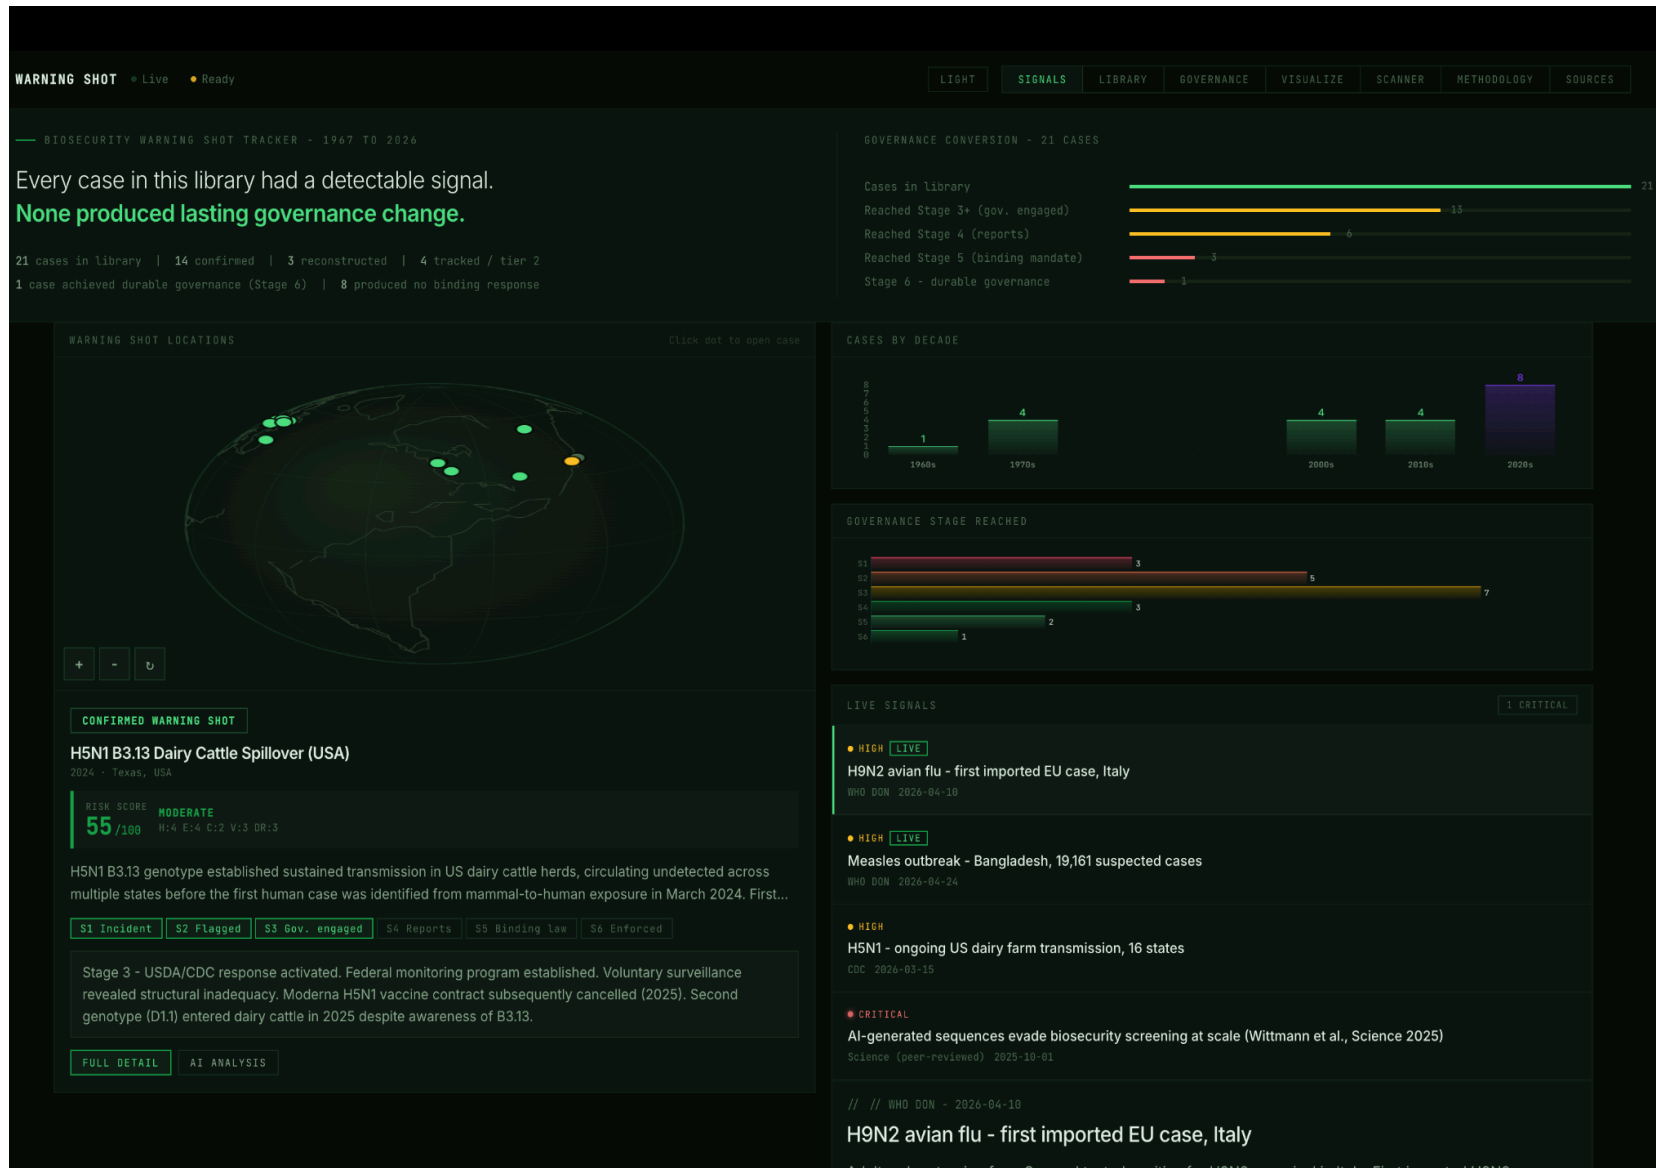


Figure 3, BioGovernance Warning Shot Dashboard

MegaSyn AI Dual-Use Chemical Design (2022)

CONFIRMED WARNING SHOT AI NEXUS: PRIMARY

RISK SCORE **HIGH**
75 / 100 H:4 E:3 C:4 V:5 DR:4

LOCATION
 North Carolina, USA

SUMMARY
 Collaborations Pharmaceuticals repurposed drug-discovery AI (MegaSyn) to generate novel chemical weapon candidates in under 6 hours. Generated VX analogues with predicted high toxicity. Published as a deliberate warning. Zero binding governance output followed.

C1-C5 ASSESSMENT

C1 - Discrete event PASS	C2 - Structural vulnerability PASS
C3 - Escalation pathway PASS	C4 - Prospective accessibility PASS
C5 - Sub-catastrophic ceiling PASS	

GOVERNANCE PIPELINE - STAGE 2/6 REACHED

S1 Incident S2 Flagged S3 Gov. engaged S4 Reports S5 Binding law S6 Enforced

Signal recognized but zero binding governance output. No regulator covers AI-enabled dual-use design. Paradigm case where all three blocking factors reinforce simultaneously.

THREE-FACTOR MODEL
 Signal clarity: **strong** Domestic victim: **weak** Regulatory home: **absent**

AI ANALYSIS PRIMARY SOURCE

Figure 4, Detailed Warning Shot Description

WARNING SHOT Live Ready

LIBRARY SIGNALS LIBRARY GOVERNANCE VISUALIZE SCANNER METHODOLOGY SOURCES

SIX STAGES FROM WARNING SHOT TO LASTING POLICY

STAGE 1 Incident occurs	STAGE 2 Experts flag systemic cause Fails if classified	STAGE 3 Government body engages Fails if no institution	STAGE 4 Reports published advisory only	STAGE 5 - CRITICAL BARRIER Binding law issued reversed under pressure	STAGE 6 Enforced and durable achieved once: 1978
----------------------------	---	---	---	---	--

HOW EACH CASE ENDED Click for detail and key lesson

- STAGE 4/6 - DURABLE GOVERNANCE ACHIEVED
Birmingham Smallpox Lab Accident (1978)
 Stage 6 achieved. Smallpox stocks consolidated to two WHO reference labs (CDC Atlanta and VECTOR Novosibirsk). All three...
- STAGE 5/6 - ADVISORY INSTITUTION CREATED (NO ENFORCEMENT)
IL-4 Mousepox Enhancement (2001)
 NSABB created (2004) but advisory only, not enforcement. Achieved governance response without domestic victims - substit...
- STAGE 5/6 - REPORTS PUBLISHED, GOVERNANCE REVERSED
HSN1 Gain-of-Function Ferret Studies (2011)
 NSABB recommendation reversed. GoF funding moratorium issued (2014) then lifted (2017). DURC/PEPP governance rescinded b...
- STAGE 3/6 - GOVERNMENT ENGAGED, LIMITED BINDING RULE
Marburg Virus Lab Outbreak (1967)
 International import regulations for primates strengthened. First case establishing that exotic pathogens could travel w...
- STAGE 3/6 - GOVERNMENT ENGAGED, NO BINDING RULE
Horsepox De Novo Synthesis (2017)
 WHO assessed but no binding regulation and no...

PROJECT A NEW EVENT

DESCRIBE THE BIOSECURITY EVENT
 Describe the warning shot or emerging technology...

REGULATORY HOME? COMPETITIVE DYNAMICS?

Yes, with enforcement authority No

PROJECT OUTCOME FULL AI ANALYSIS

Figure 5, Governance Conversion Framework

5. Conclusion

This project demonstrates that the biosecurity governance challenge is not primarily one of detection — warning shots are increasingly identifiable, and the patterns they form are historically consistent. The critical failure occurs downstream, at the point of conversion. Across 21 biosecurity events spanning nearly six decades, the majority stall at *Stage 3* of the Governance Conversion Framework, where *institutional engagement fails to produce binding mandates*. Many of the signals present in today's AIBio landscape are not novel; they are recurrences of previously observed structural failures that were recognized but not acted upon.

The practical implication is that better tools for identification are necessary but insufficient on their own. What the field lacks — and what this project begins to address — is a structured methodology for moving from signal recognition to governance action, with explicit diagnosis of where and why that conversion breaks down. The warning shot classification criteria, risk scoring framework, and GCF together provide a unified analytical language that analysts, policymakers, and oversight bodies can use to track emerging threats, communicate urgency across institutional boundaries, and intervene before indicators become warning shots and warning shots become disasters. As AIBio risks continue to converge and accelerate, the cost of stalling at Stage 3 will only increase.

6. Discussion and Limitations

Our findings suggest that even well-recognized warning shots rarely produce durable regulatory outcomes, which is particularly concerning as AI continues to lower the capability threshold for biological threat actors. The consistent Stage 3 stall pattern implies that the bottleneck is not awareness but institutional will and structural readiness — problems that will not self-correct as the threat environment accelerates. Closing this gap requires not only better early warning tools, but deliberate investment in the governance infrastructure needed to convert signals into binding action before the next warning shot becomes something worse.

There are several limitations to our project in the context of a time-limited hackathon, and due to the nature of open-sourced data analysis for threat indications and warning. First, we recognize that government, industry, and other proprietary or sensitive information will not be captured through open-sourced data analysis. In this regard, our method has information input limitations. Second, we make informed assumptions about gaps in information tracking systems available to government, legislative, and other key audiences (e.g. that a Governance Conversion Framework analysis & tracking system is not readily available or currently used). These assumptions bear further testing with envisioned users of our decision-analysis tools. Third, within our team we lacked deep biosecurity or biothreat experience, and acknowledge that in future iterations of our work, these methods and frameworks would be improved from additional review and feedback from biosecurity experts, especially those with context on policy-making, legislative oversight, or global monitoring and reporting of biosecurity incidents. Finally, due to the computational limits and time constraints, the dashboard contains pre-populated data, spanning from 1967 to 2026 and is not running live updates.

7. Future Work

The natural project extension is to improve the data quality and expand the dataset. This includes both timespan and the resource availability, as already mentioned in the *Limitations* section. The current methodology is heavily grounded in qualitative research such as adding a quantitative validation for the classification criteria. Another addition would be adding the escalation likelihood between Tier 3 and 2 events leading up to Tier 1 events (serving as an early warning trigger). Longer term, the warning shot BioGovernance scanner would be formalized into a continuously updated monitoring system, with automated triage feeding into the risk scoring dashboard in near real-time and custom data feeds.

8. Code and Data

Webpage: <https://a-ix-bio-warning-shots-scanner.vercel.app>

The system is implemented as a lightweight, browser-based dashboard using a structured dataset of biosecurity events derived from open-source intelligence (OSINT), including academic publications, institutional reports, and verified historical records. Each case is encoded using a consistent schema capturing warning shot criteria, tier classification, risk scoring, and governance progression. The underlying codebase and dataset are not publicly released at this stage due to the early-stage nature of the project and the need for further validation, particularly given the biosecurity and dual-use context. Materials can be made available upon request for research or policy purposes. Further details on data collection, structure, and implementation are provided in the Appendix.

9. Author Contributions

The dashboard was created by Edidiong James. The report was prepared by Zhamilia Klycheva, Edidiong James and Erik Leklem, and the data was provided by Shubhankar Dharmadhikari.

10. LLM Usage Statement

We used Claude to generate the use cases and help estimate the scores. All results and claims were independently verified by our team members. Claude was also used in a limited and supportive role during dashboard development, primarily for implementation-related tasks. The conceptual design, data structure, and analytical framework were independently developed, and all outputs were subject to human review and validation.

11. References

1. Adamson, Georgia and Allen, Gregory C., August 2025, "Opportunities to Strengthen U.S. Biosecurity from AI-Enabled Bioterrorism: What Policymakers Should Know," Center for Strategic and International Studies, <https://www.csis.org/analysis/opportunities-strengthen-us-biosecurity-ai-enabled-bioterrorism-what-policymakers-should>. (Accessed Apr 25, 2026)
2. Carter, J., Wheeler, N., Chwalek, S., Isaac, C., and Yassif, J. The Convergence of Artificial Intelligence and the Life Sciences. Washington, D.C.: Nuclear Threat Initiative, October 2023. Used for: Vulnerability pillar — biosurveillance gap (p. 24), screening gap (p. 28), regulatory gap (p. 40).
3. Klycheva Z., Dharmadhikari S., James E., Leklem E., 2026, From Near Miss to Pause: Historical Warning Shots and AI Governance, SPAR project, <https://sparai.org/projects/sp26/recJpcWbGJpuOXUs9>
4. RAND Europe and Centre for Long-Term Resilience. Global Risk Index for AI-Enabled Biological Tools. Cambridge: RAND Europe, September 2025. EP-71093. Used for: Capability pillar anchors, Appendix A, p. 73, Table 16.
5. World Health Organization. Joint External Evaluation Tool: International Health Regulations (2005), Third Edition. Geneva: WHO, 2022. Used for: Vulnerability pillar (Indicators P.6.1–P.6.2) and Detection & Response pillar (pp. 55, 74).
6. World Health Organization. Rapid Risk Assessment of Acute Public Health Events. Geneva: WHO, 2012. WHO/HSE/GAR/ARO/2012.1. Used for: Hazard (p. 12) and Exposure (p. 14) pillar anchors.

Appendix

Warning Shot Definition

C1 - Discrete bounded event The event is temporally identifiable with a discernible onset, duration and terminus. It must be distinguishable from background drift, structural trend or gradual risk accumulation.

C2 - Structural vulnerability Such event must expose a pre-existing weakness in screening, biosafety, governance of surveillance. False alarms, misinterpretations, or events that expose no real weakness do not qualify.

C3 - Escalation pathway indicate identifiable mechanisms by which biological catastrophe could occur. A plausible counterfactual harm pathway must be reconstructable step by step.

C4 - Prospective signal accessibility The catastrophic pathway revealed by the event must be epistemically accessible at the time of the event. Events classified at the time of the event do not qualify as Tier 1 warning shots, although can be considered as historically reconstructed warning shots

C5 - Sub-catastrophic ceiling The event must not constitute a catastrophe whose plausibility it reveals. While realized harm ranges from minimal to severe, it remains below the catastrophic outcome it demonstrates as possible.

A Tier 2 warning signal *satisfies C1 and at least two further conditions* but fails one, and is classified by its failure mode — most commonly C3, where a structural vulnerability is confirmed but the specific circumstances do not support a credible catastrophic escalation pathway, or C5, where the catastrophe actually occurred and the event is a disaster case rather than a near-miss.

A Tier 3 contextual indicator *fails C1 or fails multiple conditions* simultaneously. Tier 3 events do not qualify as warning shots but shape the structural landscape in which warning shots occur and are analytically essential for interpreting what Tier 1 and 2 events mean.

Risk Assessment Score

Hazard captures the intrinsic biological danger of an agent independent of context, such as transmissibility, virulence, immune evasion, and drug resistance, derived from the WHO's Rapid Risk Assessment methodology (2012, p. 12).

Exposure measures how realistically that hazard reaches a population, drawing from the same WHO framework (p. 14) and capturing pathways including lab release, dual-use publication, and environmental spillover.

Capability assesses the ease of creating, modifying, or deploying the hazard, with AI uplift as the primary variable — scored against RAND/CLTR's Global Risk Index for AI-Enabled Biological Tools (2025, Appendix A, p. 73, Table 16).

Vulnerability measures gaps in the systems meant to prevent or detect misuse — specifically the three governance gaps identified by NTI (The Convergence of AI and the Life Sciences, 2023, pp. 24, 28, 40): DNA synthesis screening, regulatory oversight, and biosurveillance — combined with biosafety indicators from the WHO Joint External Evaluation Tool (3rd ed., 2022, Indicators P.6.1–P.6.2).

Detection and Response is scored as a deficit - higher means weaker — drawing from WHO's Joint External Evaluation capacity levels (pp. 55, 74), from real-time detection and rapid coordinated response at 1 to no detection system and no response plan at 5.

Governance Conversion Framework (GCF)

Stage 1 Incident Warning shot occurs, a discrete, bounded, sub-catastrophic event revealing a structural vulnerability with credible escalation pathway.

Stage 2 Signal Recognition is framed as evidence of a broader structural problem. Key actors translate the event into a governance-relevant claim requiring institutional response.

Stage 3 Institutional Engagement A body with some authority takes up the signal, initiates investigation, convenes experts, opens regulatory proceedings.

Stage 4 Governance Output Institution produces concrete recommendations, advisory framework, or report. While this stage looks like governance, it is *not yet binding*.

Stage 5 Binding Mandate Legally or institutionally binding requirement issued — regulation, executive order, treaty obligation. Non-compliance carries consequences. It is not yet the final stage, as the governance may still be reversed.

Stage 6 Effective Governance Mandate actively enforced, compliance monitored, structural vulnerability durably addressed. Persists through political cycles. The rarest outcome and a true success.

M1 — Competition: Competitive race dynamics between actors — states, companies, or research institutions — can delay or reverse binding mandates, particularly at the Stage 4→5 and 5→6 transitions where political will is most fragile.

M2 — Adversarial Attribution: When an event is attributed to a state adversary, governance tends to shift from regulatory enforcement toward deterrence, which constrains or redirects the conversion pathway without producing binding civilian regulation.

Data Collection and Sources

The dataset was constructed through systematic open-source intelligence (OSINT) collection. Sources included:

- Peer-reviewed scientific literature in virology, biosafety, and dual-use research
- Institutional and policy reports (e.g., WHO, NTI, RAND)
- Verified investigative journalism and historical case documentation
- Public policy, regulatory, and government records

Cases were selected based on relevance to biosecurity risks and their alignment with the warning shot framework. Where possible, each case includes a verifiable source URL to support traceability.

Data Structure and Schema

Each event in the dataset is encoded using a consistent schema to enable structured analysis. The key fields include:

- Event metadata: ID, name, year, location
- Classification: Tier (1, 2, or 3), verdict category
- Warning shot criteria: C1–C5 (PASS, FAIL, or CONTESTED)
- Risk scoring inputs: Hazard, Exposure, Capability, Vulnerability, Detection and Response
- Governance analysis: GCF stage reached and outcome description
- Qualitative descriptors: Signal clarity, regulatory context, and domain classification
- Source reference: URL or citation

This structure allows for consistent comparison across cases and supports filtering, aggregation, and visualization within the dashboard.

Case Encoding Example

Below is a simplified example of how a case is represented within the dataset:

- Event: Marburg Virus Lab Outbreak (1967)
- Tier: 1
- Criteria: C1–C5 all satisfied
- Domain: Laboratory biosafety
- GCF Stage: 3 (Institutional Engagement)
- Outcome: Strengthened import regulations, no binding global framework
- Signal Clarity: Strong
- Source: CDC archival report

This encoding ensures that each case is both analytically structured and traceable to source material.

System Implementation

The dashboard is implemented as a single-page web application using HTML, CSS, and JavaScript, with no external dependencies or build process required. It is deployed as a static application via Vercel, enabling rapid access and ease of use.

The system consists of:

- A structured case library for historical events
- A visualization layer for filtering and comparing cases
- A scanning interface for evaluating new signals against the warning shot criteria

Signal Processing Workflow

The analytical workflow implemented in the dashboard follows a structured pipeline:

- Signal input → Criteria evaluation (C1–C5) → Tier classification → Risk scoring → Governance Conversion Framework mapping → Comparative analysis with historical cases

This process enables consistent classification of events and supports identification of patterns across cases.

Limitations of the Dataset

The dataset is manually curated and relies on open-source information, which introduces several constraints:

- Potential selection bias in case inclusion
- Limited visibility into classified or underreported incidents
- Subjectivity in interpreting classification criteria
- Static dataset with no real-time updates

These limitations should be considered when interpreting results.

Availability

The interactive dashboard is publicly accessible via the deployed interface.

The underlying dataset and codebase are available upon request for academic, research, or policy-oriented use.