Matthew F. McKnight, MPP, MBA
General Manager, Ginkgo Biosecurity
Ginkgo Bioworks
27 Drydock Avenue, Floor 8
Boston, MA 02210

Additional Questions for the Record

The Honorable John Joyce (R-PA)

1. **Is the Federal Select Agents and Toxins List sufficient in overseeing the possession, use, and transfer of biological agents that have the potential to pose a severe threat to the public, animal, or plant health given that AI-enabled biothreats may not even be imagined yet?**

   No. Biological design tools, particularly AI-driven biological design models, are increasingly sophisticated and able to design around blacklists such as the Select Agents and Toxin list. Those with the necessary skills can disguise hazardous sequences, both at the protein and gene levels, so they no longer match the specific entries on a static regulatory list.

   a. **What criteria should be used to determine which DNA sequences, or parts of sequences, should be regulated as agents of concern?**

      We are engaged in an unending arms race against biological design tools. Any static set of criteria will become obsolete, just as the Select Agents and Toxins list has. To stay left of boom, the Government could fund an evolving portfolio of approaches for threat detection in an ongoing manner.

      Today, we recommend a layered approach, noting that new layers will need to be added, and old layers potentially retired, as biological design tools advance:

      - Updated Blacklists: At the lowest level, blacklists, such as the Select Agents and Toxins list should be updated and expanded on a regular cadence. Doing so will enable sensitive, rapid detection of blacklisted threats, with near-zero false negative detections.
      - Functional Prediction: Develop tools (including AI models, and in particular DNA or protein language models) to predict the actual function or behavior of a DNA or protein sequence, rather than just checking its name against a list. Doing so would enable detection of many threats that have been designed to evade blacklists.
        - Specialized Modeling: Distinct models may be needed to fully cover different types of threats, such as catalytic protein toxins (e.g. ricin); non-catalytic proteins that cause toxicity via binding (e.g. α-amanitin, conotoxins); short synthetic "minibinder" peptides that may cause toxicity via a variety of mechanisms; and proteins involved in synthesis of small-molecule toxins (e.g. saxitoxins).

Please note: While there is no Commercial Off The Shelf (COTS) solution for functional screening today, one could be built using existing models, appropriate fine-tuning, test data, and bioinformatics software.

- Addressing Capability Gaps: Currently, screening very short DNA sequences for threats is scientifically challenging and subject to high false negative and false positive rates. The Government could fund a program or programs to close this gap.

2. Is the suitability screening of personnel to be in the Federal Select Agent Program sufficient to protect against an insider threat at a government laboratory?

Screening of personnel is an essential and useful layer in a multi-layered "Swiss-cheese" model of security. However, it may not be sufficient to preempt all types of insider threat. For example, well-meaning insiders may inadvertently create and/or release a dangerous biological sequence or organism.

To address these vulnerabilities, security protocols must evolve beyond personnel vetting. We strongly recommend coupling this approach with Biothreat Radar, the pervasive, persistent monitoring of nucleic acids circulating within and/or escaping from a facility via human or environmental vectors. This can include targeted monitoring of air filtration systems, wastewater, and surfaces, as well as direct samples (e.g. nasal swabs) from lab team members to maintain a high-resolution audit of biological materials in circulation.

By identifying anomalous sequences before they exit the facility or spread unchecked, this surveillance provides a critical early warning of risk. This allows for not only the rapid containment of accidental leaks, but also provides the necessary oversight to halt specific research projects that are found to pose a grave risk to public safety. Ultimately, Biothreat Radar should be deployed routinely and globally, leveraging environmental monitoring tools to detect novel protein and other biological sequences of concern, whether known or unknown.

3. What government auditing or compliance mechanisms could verify that DNA synthesis companies are screening effectively?

Current screening relies on vendor self-regulation rather than Independent Verification & Validation (IV&V). Without independent teams to stress-test these systems, we have no guarantee they can actually catch the threats they claim to block. To ensure these systems actually work, we recommend:

- Independent Red-Teaming: The government could oversee a system where independent teams test a company's defenses by placing "dummy" orders for dangerous sequences to see if they are flagged.
- Detecting Split-Orders: Audits should test if software can catch hazardous sequences that have been broken into smaller pieces and spread across different orders or providers to avoid detection.
- Verification Standards: Implementing a "Know Your Customer" (KYC) framework to verify the legitimacy of the purchaser. This involves multi-tiered identity verification and

monitoring to ensure that the individual and their affiliated institution have a valid research justification for acquiring specific genetic material.
- Governance Sandboxes: Creating collaborative testing grounds to develop and validate technical biosecurity standards, ensuring that compliance measures are both scientifically robust and practically effective against emerging threats before they are implemented as industry requirements.
- Federal Incentives: Compliance could be tied to federal funding and market access, ensuring that a company's ability to do business with the federal government is contingent upon successfully passing periodic audits and maintaining high performance in testing environments.

Note: Existing legislation – the Biosecurity Modernization and Innovation Act (S. 3741)– addresses several of these elements.

4. What steps are AI companies like OpenAI taking to mitigate risks and protect their platforms from misuse related to biosecurity and terrorism?

Ginkgo Biosecurity is not involved with the internal operations of AI companies and does not have unique insight into the confidential practices of organizations like OpenAI. However, through our work in building global biothreat detection infrastructure and our engagement with the broader security community, we are aware of several industry-standard frameworks and research initiatives designed to address the intersection of AI and biological risks.

Leading AI developers have adopted structured protocols to prevent their models from being used to facilitate biological harm. These frameworks focus on identifying and neutralizing catastrophic risks before models are deployed to the public.

- OpenAI's Preparedness Framework: Establishes Tracked Categories for biological and chemical capabilities, utilizing rigorous evaluations to detect "High" or "Critical" thresholds for "severe harm," defined as injury or death to thousands of people or billions of dollars in economic damage. Before deployment, the framework requires implementing safeguards and security controls—such as refusal training and usage monitoring—to sufficiently minimize the risk of a model providing uplift for the creation or use of biological weapons. Final deployment decisions for models are overseen by a cross-functional Safety Advisory Group (SAG), which ensures that models do not exceed risk levels that could lead to severe harm.
- Anthropic's Responsible Scaling Policy (RSP): Anthropic's Responsible Scaling Policy mitigates biosecurity and terrorism risks through a framework of AI Safety Levels (ASL), which mandates increasingly stringent safety and security protocols as models gain the potential for "catastrophic misuse." Under this policy, models reaching ASL-3—those that substantially increase the risk of creating bioweapons compared to search engines—require "unusually strong security" and a commitment not to deploy the models if they show meaningful risk under adversarial red-teaming. This system ensures that scaling is paused if it outstrips the company's ability to comply with the necessary safety procedures while also incentivizing the company to solve complex safety issues.
- Technical Red Teaming: Companies employ red teams consisting of PhD-level biologists to stress-test models. These experts attempt to bypass safeguards to see if a model can

provide “tacit knowledge”—the practical, hands-on troubleshooting advice traditionally only learned in a physical lab.

a. Are these steps effective? If not, what would you recommend AI companies do to improve risk mitigation and biosecurity monitoring?

Based on recent studies and industry frameworks, the current steps taken by AI companies are partially effective but face significant challenges as AI capabilities advance. The effectiveness of current mitigations is best understood through the lens of “capability uplift” research. Below are a few examples of these assessments:

- RAND: Contemporary Foundation AI Models Increase Biological Weapons Risk
  - Scope: Tactical Planning
  - Key Biosecurity Findings: Contends that current AI safety evaluations significantly underestimate biological weapons risk by incorrectly assuming that complex technical tasks require incommunicable “tacit knowledge”. By demonstrating that frontier models can articulate detailed instructions for sophisticated procedures like viral recovery, the authors warn that AI lowers the technical barrier for malicious actors, necessitating stricter regulatory oversight and improved safety benchmarks.
- RAND: Bridging the Digital to Physical Divide
  - Scope: DNA Design and Acquisition
  - Key Biosecurity Findings: The newest generation of AI models (GPT-5, Opus 4.5, and Gemini 3 Pro) demonstrated a significant breakthrough by reliably designing biologically coherent DNA segments, overcoming a major error-prone step that previously caused designs to fail when moved from a computer to a lab. By successfully automating the digital-to-physical transition—including the selection of correct biological tools and the generation of accurate lab protocols—these models proved capable of guiding a non-expert through the creation of functional biological products in a real-world laboratory.
- Los Alamos: Measuring skill-based uplift from AI in a real biological laboratory
  - Scope: Lab Proficiency
  - Key Biosecurity Findings: The pilot study at Los Alamos National Laboratory found that access to ChatGPT-o1 increased the success rate of novices performing complex wet-lab tasks from 20% to 60% on the first attempt, demonstrating that AI provides a measurable “real uplift” in practical biological skills.

To improve risk mitigation and biosecurity monitoring, we recommend that AI companies:

- Promote and Support Biothreat Radar Deployment: Because no prevention mechanism is guaranteed to be ironclad, AI companies could actively support the deployment of Biothreat Radar—global, persistent genomic surveillance infrastructure. Because digital guardrails can be bypassed via fine-tuning or “jailbreaking,” a physical sensor grid can act as a final backstop, detecting

pathogens in the environment as early as possible, regardless of their origin. AI companies can then provide the computational power and model-based anomaly detection required to analyze the massive metagenomic data streams generated by the Biothreat Radar, helping separate signal from noise and informing response.

- Invest in Downstream Mitigation and Countermeasures: AI companies could dedicate a portion of their research to "defensive AI," using their models to accelerate the development of antibodies, vaccines, and diagnostic tools. The goal should be to outpace the design of new threats with the rapid creation of new protections.

5. How can AI interpret the DNA of a pathogen and trace its origins?

Biological AI tools are excellent at many sequence interpretation tasks when given DNA, protein encoded by DNA, or both, as inputs. A full list would be exhaustive, but to name several relevant and very recent examples:

- Gaia (Nishant et al., 2025) is able to accurately locate and predict the functions of individual genes, as well as emergent functions of groups of genes, within a prokaryotic genome. This tool could be used to identify various signals of danger, such as pathogenicity islands, gene clusters encoding toxins, or single-gene toxins.
- Given a protein-encoding sequence from a genome, models like Chai-1 (Chai Discovery Team et al., 2024) and Boltz-2 (Passaro et al., 2025) are capable of predicting protein structure and identifying their binding partners, two key elements of accurately predicting function. These tools could be used as the basis of a workflow to identify novel or heavily obscured toxin sequences, and infer their mechanisms of action.
- Evo 2 (Brixi et al., 2025), a foundation model, predicts negative effects conferred by multiple types of genomic mutations with surprising accuracy, in coding as well as non-coding regions. Evo 2 could very plausibly be adapted to a variety of new prediction tasks, such as predicting the effect of a mutation on a virus's virulence, via fine-tuning.
- Provided a mammalian DNA sequence, AlphaGenome (Avsec et al., 2026) is able to predict a broad variety of functional properties beyond protein encoding, such as transcription factor binding, splice site usage, and histone modifications, as well as the effects of mutations on these properties. With fine-tuning, this tool could be reasonably adapted to understand the effects of intergenic mutations in pathogens, viral, prokaryotic, and eukaryotic.

Regarding tracing pathogen origins, we typically break this down into: (1) determining the geographic location where a threat likely originated, (2) finding its closest known relatives, and (3) inferring its evolutionary history. There are a variety of AI-enabled tools that are directed at these questions, but today, the most accurate tools for attribution use phylogenetic (statistical evolutionary) approaches rather than AI. These phylogenetic approaches work by comparing genetic sequences to reconstruct a family tree, allowing researchers to infer how a pathogen may have mutated over time. By analyzing the types, numbers, and locations of mutations, scientists can assess the probability that the genome arose from a natural ancestor under natural evolutionary processes. High probabilities suggest natural origins, while low probabilities suggest the pathogen may have been manipulated via gain-of-function research

or otherwise genetically engineered. While phylogenetics currently offers the best toolkit for making such inferences, we anticipate AI-enabled tools will improve and consider it likely that they may displace phylogenetic methods as state of the art.

To go beyond sequence alone, the most powerful attribution workflows would correlate genomic signals with non-biological intelligence streams to add context and constrain hypotheses. For example, AI could integrate phylogeographic outputs with airline and ground mobility data, trade and agricultural shipment records, climate and vector suitability models, satellite imagery of relevant facilities or outbreak conditions, procurement and supply-chain signals (e.g., unusual reagent purchases), digital epidemiology (search trends, social chatter), and even geopolitical or conflict data. Multimodal models can weight these heterogeneous inputs to generate holistic origin assessments, flag inconsistencies (e.g., a lineage's inferred migration path that contradicts known travel flows), and continuously update assessments as new data arrives. In practice, this looks like an iterative fusion layer on top of genomic analysis—turning "closest genetic relatives" into a defensible narrative about where, how, and under what conditions a pathogen emerged—while preserving auditability and uncertainty bounds appropriate for high-stakes biosecurity decisions.