

# PUBLIC SUBMISSION

|  |
|--|
| <b>Received:</b> May 29, 2025<br><b>Tracking No.</b> mb9-xhwn-27va<br><b>Comments Due:</b> May 28,<br>2025 <b>Submission Type:</b> API |
|--|

**Docket:** NSF-2025-OGC-0001  
NITRD\_FRDOC\_0001

**Comment On:** NSF-2025-OGC-0001-0001  
Request for Information: Development of a 2025 National Artificial Intelligence Research and Development Strategic Plan

**Document:** NSF-2025-OGC-0001-DRAFT-0277  
Comment on FR Doc # 2025-07332

---

## Submitter Information

**Organization:** The MITRE Corporation

---

## General Comment

MITRE is pleased to respond to the RFI. Please see the uploaded PDF (MITRE Response to the National AI Research and Dev Plan RFI.pdf) for our detailed response.

---

## Attachments

MITRE Response to the National AI Research and Dev Plan RFI



# **MITRE's Response to the OSTP RFI on a 2025 National Artificial Intelligence (AI) Research and Development (R&D) Strategic Plan**

**May 29, 2025**

For additional information about this response, please contact:

Katie Enos  
Michael Garriss  
The MITRE Corporation  
202 Burlington Road  
Bedford, MA 01730-1420

[governmentrelations@mitre.org](mailto:governmentrelations@mitre.org)

“This document is approved for public dissemination. The document contains no business proprietary or confidential information. Document contents may be reused by the government in developing the 2025 National AI R&D Strategic Plan and associated documents without attribution.”

## Introduction

The United States leads the world in artificial intelligence (AI) innovation, models and tools. These advances are already central to our economic strength, as they underpin global market leaders, and AI has the promise of fueling scientific breakthroughs that will supercharge our economy going forward. However, the United States' leadership is increasingly vulnerable due to intense global competition, and a coordinated National level campaign is needed to sustain and strengthen our position. *MITRE's Response to OSTP RFI on AI Action Plan*<sup>1</sup> highlights the foundational importance of both basic and applied research and development (R&D) in accelerating the advancement of AI technologies. In the paper *Partnerships to Accelerate Advancement of Priority S&T*,<sup>2</sup> MITRE points out that investment and public-private partnership in R&D is essential to fostering invention and driving innovation. The government plays a key role in reducing the institutional risk associated with research, particularly in areas with long-term national benefits. To this end, MITRE emphasizes the need for the government to invest in AI research Grand Challenges and recommends enabling investments in Federal Frontier Labs, AI assurance research for trusted innovation, and securing American AI research.

## National Grand Challenges in AI Research

The Strategic Plan should establish AI Grand Challenge problems of U.S. strategic importance where private investment is unlikely to advance our leadership. These challenges will align federal research investments with national interests, and bring together government, industry, and academia. Grand Challenge problems can serve as a powerful mechanism to unite capable stakeholders around critical missions, and federal action is necessary for addressing those challenges to maximize impact and ensure alignment with broader national priorities.<sup>3</sup>

Execution priorities should include rapidly convening stakeholders, addressing competitive pressures and regulatory hurdles, establishing norms and incentives for collaboration, protecting individual interests (e.g., reputation, confidentiality, and intellectual property), and developing solutions that exceed the capabilities of any single entity.<sup>4</sup> Leveraging existing mission-focused entities including National Labs and Federally Funded Research and Development Centers (FFRDCs) that are deeply grounded in serving National interests and already partner with industry and academia can serve as catalysts.

Research focus areas should include:

- **Biotechnology:** Using AI to accelerate drug discovery, biomanufacturing, personalized medicine, genomic research, and the response to health crises such as chronic diseases and pandemics. This challenge should create transformative solutions that enhance biosecurity, optimize clinical trials, and build widespread adoption of a learning health

---

<sup>1</sup> MITRE's Response to OSTP RFI on AI Action Plan. 2025. MITRE. <https://www.mitre.org/news-insights/publication/mitres-response-ostp-rfi-ai-action-plan>.

<sup>2</sup> Partnerships to Accelerate Advancement of Priority S&T. 2023. MITRE. <https://www.mitre.org/sites/default/files/2023-09/PR-23-02057-05-Partnerships-to-Accelerate-Advancement-of-Priority-S-T.pdf>.

<sup>3</sup> Use of Grand Challenges in the Federal Government. 2019. IDA. <https://www.ida.org/-/media/feature/publications/u/us/use-of-grand-challenges-in-the-federal-government/d10699final.ashx>.

<sup>4</sup> See <https://assemble.mitre.org/>.

architecture.<sup>5</sup> Additionally, the challenge should focus on systematic design practices to promote safe and secure applications of biotechnology while characterizing and mitigating emerging biological threats.

- **Semiconductors:** Harnessing AI to revolutionize chip design, improve manufacturing efficiency, and strengthen domestic supply chains—key steps toward bolstering national economic competitiveness. Special emphasis should be given to scalable AI solutions that enhance supply chain security and safeguard the integrity of microelectronics in critical systems. This includes the creation of robust semiconductor security frameworks to defend against adversarial threats. The challenge should also aim to drive innovation in semiconductor manufacturing, design security, and analytical capabilities to support both legacy systems and emerging technology nodes. Additionally, the initiative should encourage the adoption of memory-safe programming languages, AI-powered design tools, and cutting-edge co-design approaches.
- **Domestic Energy:** Leveraging AI to optimize energy production, improve grid reliability, and accelerate the development of next-generation energy sources like nuclear and renewables, ensuring energy independence and affordability. This challenge should leverage AI in meeting the nation's growing energy demands while ensuring resilience, sustainability, and security. The development of resilient energy architectures capable of supporting the rapid expansion of data centers, AI-driven systems, and other high-demand sectors should be prioritized, while maintaining grid stability and affordability.
- **Advanced Manufacturing:** Integrating AI to streamline production processes, enhance quality control, and drive innovation in materials science, reinforcing U.S. leadership in global markets. This initiative should enhance supply chain resilience and accelerate the adoption of advanced technologies such as additive manufacturing and robotics. Using AI may reduce production costs, improve quality assurance, and enable real-time adaptability in manufacturing systems. Additionally, research should explore transformative approaches to workforce development, empowering human-machine collaboration and addressing critical skill gaps in the sector.
- **Cybersecurity:** Harnessing AI to detect threats, protect critical infrastructure and essential services, and respond to evolving risks in real-time, safeguarding both public and private sectors. The challenge should prioritize the development of zero-trust architectures, secure software supply chains, and post-quantum cryptography to address emerging vulnerabilities and safeguard critical assets. It should also accelerate the adoption of automated cybersecurity workflows, enabling real-time threat mitigation and closing workforce gaps in the cyber domain. Additionally, the initiative should explore transformative approaches to integrating AI into DevSecOps practices, improving software assurance, and advancing cyber resilience engineering for complex systems.
- **National Security:** Appropriately leveraging AI to maintain a decisive technological edge in defense/intelligence applications and deter adversaries. AI research should be

---

<sup>5</sup> A learning health architecture is a framework that enables continuous improvement in health outcomes by systematically integrating data from clinical practice, research, and patient experiences to generate actionable insights and inform decision-making. See Foley T, Vale L. A framework for understanding, designing, developing and evaluating learning health systems. *Learn Health Syst.* 2022. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9835047/>.

applied to enhance assured communications, optimize command and control systems, and enable real-time decision making across domains such as border security, air traffic control, air and missile defense, autonomous vehicles, and military operations. AI research should also identify stressors such as population growth, energy demands, and increasingly sophisticated threats from near-peer adversaries and transnational criminal organizations.

- **Efficient and Effective Government:** Optimizing government functions including acquisition, workflow management, and services administration. AI research should aim to improve the delivery of public benefits, such as healthcare, Social Security, taxation, and veterans' services—enabling data-informed policymaking to enhance population outcomes while managing costs and ensuring long-term sustainability.

Investing in Grand Challenges that span strategic application areas will accelerate U.S. technological innovation, advance government missions, and drive economic growth. Large-scale, multi-sector collaboration will ensure that AI innovation strengthens national security, delivers tangible benefits to all Americans, and solidifies the nation's global leadership in science, technology, and innovation.

The remainder of this RFI response focuses on strategic areas for action that can serve as “enablers” for carrying out successful AI research Grand Challenge problems.

## **Federal Frontier Labs**

The United States faces a pivotal moment in the global race for AI leadership, where competition with China will shape the future of economic power, national security, and technological dominance in the 21st century. While the United States has an early lead through pioneering commercial frontier labs (CFLs), it is imperative that we not only maintain but accelerate our lead in AI innovation. Where the fruits of today's commercial innovation are available equally to all nations, it is also important to create a distinctly American approach to AI innovation that capitalizes on our nation's exclusive, mission-critical data and distributes those benefits throughout our economy.

To secure our competitive edge, the federal government must make decisive investments in advanced AI research that not only build upon the momentum of private-sector initiatives but also tailor innovations to meet critical government mission needs. Leveraging public-private models that enable translational engineering can unlock new markets while delivering breakthrough, mission-focused AI solutions. Federal Frontier Labs (FFL) can accomplish this through structured collaboration with industry and academia, responsible utilization of government data, and dedicated and secure infrastructure.

### **FFLs as Public-Private AI Partnerships**

FFLs are envisioned as hubs for public-private partnerships, designed to position the United States as a leader in next-generation AI innovation. These labs can drive domain-specific breakthroughs by integrating commercial frontier lab expertise, government data and mission alignment, academic research, and high-performance computing (HPC) infrastructure. Through thematic research initiatives like Grand Challenges, FFLs can advance the use of commercially developed general-purpose foundation models, train domain-specific models, and accelerate the development and adoption of entirely new technical paradigms beyond transformer architectures.

General-purpose foundation models, developed by commercial frontier labs, will continue to advance universal AI capabilities. FFLs can partner with CFLs to adapt and extend these models for federal missions. A core role for FFLs is leveraging federal data to train domain-specific foundation models in application areas such as healthcare, defense, and intelligence. Success will depend on partnerships that combine government mission and data understanding with expertise in foundation model development and deployment. This collaborative approach not only strengthens mission outcomes but also translates technical breakthroughs into new commercial opportunities.

FFLs can also facilitate academic partnerships focused on research critical to U.S. long-term competitiveness, such as technical paradigms like scaling neuromorphic or brain-inspired algorithms that go beyond transformer-based models, developing low-SWaP (size, weight, and power) models for edge deployments, advancing federated learning for privacy-preserving training, and enhancing AI security through encryption and watermarking. These efforts are essential to ensure AI innovation aligns with federal priorities and remains resilient against adversarial threats.

### **Focus on Government Missions**

FFLs should prioritize three key areas: open science, defense, and intelligence. In open science, FFLs could expand initiatives like the Department of Energy's Frontiers in Artificial Intelligence for Science, Security, and Technology<sup>6</sup> and Los Alamos National Laboratory's collaboration with OpenAI<sup>7</sup> to advance research in physics, healthcare, and biology. In defense, FFLs could improve supply chains, predictive maintenance, and cyber operations. In intelligence, FFLs could enhance capabilities in open-source, imagery, and signals intelligence. By securely leveraging government data at scale, FFLs can address societal challenges and mission-critical needs, particularly in areas with limited commercial interest.

### **Dedicated Infrastructure**

FFLs can expedite federal government access to at-scale, dedicated AI infrastructure, which is essential for training and deploying advanced AI models. While private sector partnerships in projects such as Stargate aim to support CFLs in advancing AI research and development, with FFLs, the government can form public-private partnerships to access scaled data center and compute resources in support of its specific AI needs. Deploying commercially packaged, exascale AI-computing pods to FFLs and HPC institutions can create government innovation clusters within 12 months of authorization. These pods, installed and configured to support different levels of sensitivity—from open, to controlled unclassified and classified research—would provide secure environments for vetted researchers<sup>8</sup> to advance cutting-edge AI capabilities with government data.<sup>9</sup>

---

<sup>6</sup> See <https://www.energy.gov/fasst>.

<sup>7</sup> Los Alamos National Laboratory partners with OpenAI to advance national security. 2025. LANL. <https://www.lanl.gov/media/news/0130-open-ai>.

<sup>8</sup> There is need of commercial security clearances requiring federal guidelines and implementation support on suitability and counterintelligence vetting for researchers, similar to security clearances in the national security space (but geared for staff outside government employees and contractors) as an extension to NSPM-33.

<sup>9</sup> See, e.g., MITRE's Federal AI Sandbox. 2024. MITRE. <https://www.mitre.org/news-insights/fact-sheet/mitres-federal-ai-sandbox>

By strategically investing in FFLs, the United States can maintain not just its global leadership in AI innovation but also in impactful AI adoption. Leveraging our own federal data assets for mission-critical AI development in FFLs will help our nation retain a competitive edge, especially over our adversaries, in addressing the national grand challenges outlined in the previous section.

## **AI Assurance Research for Trusted Innovation**

AI is rapidly transforming industries, with sector-specific regulators navigating the challenges of overseeing its adoption. While AI offers immense potential, it also introduces risks related to factors such as robustness, explainability, reliability, and security. Whether validating the effectiveness of AI-driven medical devices or preventing accidents in autonomous vehicles, regulators face the pressing need for standards and frameworks to identify and address risks across the full life cycle of AI systems—spanning design, development, testing, acquisition, deployment, and ongoing monitoring. The federal government should leverage existing sector-specific regulatory agencies and mechanisms to identify and mitigate AI-related risks in mission spaces. Generalizable assurance research—an engineering process for discovering, assessing, and managing risks throughout the life cycle of AI-enabled systems<sup>10</sup>—should inform and strengthen sector-specific regulatory approaches, ensuring that oversight is tailored to the unique demands of each industry while maintaining consistency in addressing AI risks.

To unlock the full potential of AI technologies while mitigating risks, the federal government should prioritize research into AI assurance ensuring systems operate effectively, exhibit intended behaviors, and generate valid outputs that empower humans to achieve their mission objectives. Advancing this research is critical for fostering trust in AI applications, particularly in high-stakes government functions, and driving responsible development across sectors.

AI assurance research builds on the National Institute of Standards and Technology's AI Risk Management Framework and focuses on developing repeatable engineering approaches to analyze and assure AI systems. These approaches must adapt to diverse missions and applications, ensuring systems meet rigorous standards for reliability, security, and trustworthiness. Further, AI assurance approaches themselves are increasingly becoming AI-enabled. By advancing AI assurance methodologies, the United States can enhance public trust, improve national security, and accelerate AI adoption in critical domains.

AI assurance research should focus on developing reusable and efficient methods for:

1. **Discovering Assurance Needs:** Understanding the coupling between mission problems and the proposed AI solutions to identify actionable assurance requirements. This involves assessing the system's intended use, potential risks, and alignment with stakeholder goals.
2. **Characterizing and Prioritizing Risks:** Systematically identifying and ranking risks based on their potential impact to mission, likelihood, and relevance to the system's operational context.

---

<sup>10</sup> Robbins, Eris, Kapusta, Booker, and Ward, AI Assurance: A Repeatable Process for Assuring AI-enabled Systems. 2024. MITRE. <https://www.mitre.org/news-insights/publication/ai-assurance-repeatable-process-assuring-ai-enabled-systems>.



3. **Evaluating Risks:** Measuring, testing, and assessing current and emerging AI capabilities and risks. This includes evaluating an AI system's ability to generate valid outputs, adapt to dynamic environments, and resist adversarial attacks.
4. **Managing Risks:** Mitigating identified risks, including rapid development of novel mitigation approaches, ensuring ongoing system reliability, and maintaining trust throughout the system's life cycle.

AI assurance research also needs to be supported with laboratory infrastructure and capabilities that can leverage a variety of physical and digital resources, depending on what AI-enabled system is being assured and for what mission. Some of those capabilities will be targeted and specific to the assurance case under consideration (e.g., how to best mitigate a specific type of AI assurance risk) while others will be more general and can serve as reusable resources in any AI assurance consideration.

Accordingly, establishing a network of sector-specific AI assurance labs would provide reusable capabilities to rigorously test and validate AI systems for trustworthiness, resilience, and mission readiness and foster knowledge transfer across domains. These labs would provide advanced resources with reach-across capabilities for public-private collaboration such as secure data and mission-relevant simulation environments tailored to specific operational needs and living AI assurance knowledge bases that facilitate information sharing. Such a network of assurance labs can be integrated with FFLs to strengthen collaboration among government, industry, and academia.

## Securing American AI Research

As the United States continues to lead in the research, development, and deployment of AI, it is imperative to secure the United States' research ecosystem, including the FFLs as proposed above, against a wide range of threats as emphasized in National Security Presidential Memorandum-33. Adversaries such as China and Russia are actively engaged in efforts to steal, poison/disrupt, or destroy data, AI intellectual property, and critical infrastructure to gain military and economic advantages, often employing aggressive tactics that threaten global technological leadership and security.<sup>11,12,13</sup> By implementing comprehensive security measures, we can protect AI research and advancements—particularly those aligned with national security and economic interests.

Suggested security measures are:

- **Infrastructure Cybersecurity:** Implement robust cybersecurity to safeguard AI research, data, and emerging intellectual property from cyber attacks, breaches, and theft. This means deploying advanced threat detection, incident sharing, rapid incident response, and

---

<sup>11</sup> Ford, Clancy, and Blackburn. A "Horizon Strategy" Framework for Science and Technology Policy for the U.S. Innovation Economy and America's Competitive Success. 2021. MITRE.  
<https://www.mitre.org/sites/default/files/2021-11/prs-21-1440-horizon-strategy-framework-science-technology-policy.pdf>.

<sup>12</sup> A Vision for Competitiveness: Mid-Decade Opportunities for Strategic Victory. 2023. Special Competitive Studies Project. <https://www.scsip.ai/wp-content/uploads/2023/04/Vision-for-Competitiveness-1-1.pdf>.

<sup>13</sup> Jin. A Policymaker's Guide to China's Technology Security Strategy. 2025. Information Technology & Innovation Foundation. <https://itif.org/publications/2025/02/18/a-policymakers-guide-to-chinas-technology-security-strategy/>.



continuous monitoring to keep AI research secure.<sup>14</sup> Additionally, promoting research in security-by-design measures—developing new ways of building security directly into AI models—is important to making AI technology inherently more secure.

- **Infrastructure Physical Security:** Physically secure infrastructure (e.g., data centers, communication lines, and power and water sources) to protect large facility investments underpinning at-scale frontier AI research. Strong physical security is paramount to preventing unauthorized physical access or disruption. Data centers, in particular, need secure access controls, regular audits, and operational resilience to safeguard computational and data resources powering AI.
- **Intelligence Collection and Assessment:** We must understand how our adversaries may be exfiltrating and exploiting the United States' AI research ecosystem and the threats such activities/practices pose to U.S. global competitiveness and national security.<sup>15</sup> This requires the collection of science and technology intelligence on adversary AI programs (and counterintelligence on adversary activities) to compare against a baseline of U.S. AI programs. Specifically, this entails continuous monitoring of adversary development through open source and other means as well as monitoring of critical infrastructure for attacks. Further, use of AI red teaming of our own infrastructure calibrates risks and enables assessments of risks to our national security. Therefore, the government should enable research to:
  - Monitor and evaluate AI capabilities of adversaries, as well as collect and report on adversarial tradecraft of these capabilities' use.
  - Characterize the “reach” of such adversary capabilities into U.S. public and commercial AI infrastructure and operations.
  - Characterize and assess the threat such reach poses on our national security.
  - Develop novel mitigations and, depending on threat criticality, targeted regulation.
  - Provide continuous red teaming of U.S. public and commercial AI infrastructure and operations.

These research efforts should encompass testbed environments to provide indications and warnings through objective metrics for advances in AI, its capacity for autonomous effects, and growth in intelligence, for the purpose of accelerating industrial competition and providing our nation with the means to outpace such emerging adversarial capabilities. Finally, the collected intelligence should be used to establish and enforce active defense goals around AI diffusion to inform export controls, sanctions, and non-/counter-proliferation regimes.

- **Information Sharing:** Learning from the relatively slow ramp up on cybersecurity in the United States, and recognizing the rapid technological advancement in AI, we must accelerate our understanding of threats, vulnerabilities, and risks to AI technology adoption for consequential use. Moreover, many real-world AI incidents are happening in

---

<sup>14</sup> Research facility operators can participate in AI incident reporting and response through MITRE's Adversarial Threat Landscape for Artificial-Intelligence Systems (ATLAS™). See <https://atlas.mitre.org/>.

<sup>15</sup> Ford and Meyerriecks. Science and Technology Net Assessment and Competitive Strategy. 2022. MITRE. <https://www.mitre.org/sites/default/files/2023-01/PR-22-02024-5-Science-and-Technology-Net-Assessment-and-Competitive-Strategy.pdf>.

the public sector. In order to share incidents of adversarial exfiltration and exploitation of AI-enabled systems, the government should fund research that enables:

- Information Sharing and Analysis Centers (ISACs) and their interactions with government agencies so that AI risk and threat information sharing keeps pace with rapidly evolving AI technology developments.
- Further development and adoption of AI threat and mitigation sharing frameworks such as ATLAS™.<sup>16</sup>

By prioritizing these security measures, U.S. research institutions (industry, academic, and government) can control when and what research findings are shared publicly thus protecting U.S. research investments while continuing to support the research community at large, all while driving economic growth and national strength.

## **About MITRE**

MITRE is a not-for-profit organization that operates in the public interest, providing objective, data-driven insights to address the nation's most pressing challenges. As the operator of multiple FFRDCs, MITRE brings a unique, independent perspective to its work with federal agencies, free from political or commercial pressures. With over five decades of experience applying AI and machine learning to advance critical government missions, MITRE is well-positioned to contribute to the development of a robust National AI Research and Development Strategic Plan. Our expertise spans the entire AI life cycle, enabling us to anticipate and address emerging research needs that are essential to the nation's security, economic leadership, and public well-being.

---

<sup>16</sup> See <https://atlas.mitre.org/>.