MITRE's Response to the OSTP RFI on AI Action Plan

March 13, 2025

For additional information about this response, please contact:

Duane Blackburn
Center for Data-Driven Policy
The MITRE Corporation
7596 Colshire Drive
McLean, VA 22102-7539

policy@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 AI Action Plan and associated documents without attribution."

About MITRE

MITRE is a not-for-profit company that works in the public interest to tackle difficult problems that challenge the safety, stability, security, and well-being of our nation. We operate multiple federally funded research and development centers (FFRDCs), participate in public-private partnerships across national security and civilian agency missions, and maintain an independent technology research program in areas such as artificial intelligence (AI), intuitive data science, quantum information science, health informatics, policy and economic expertise, trustworthy autonomy, cyber threat sharing, and cyber resilience. MITRE's 9,000(+) employees work in the public interest to solve problems for a safer world, with scientific integrity being fundamental to our existence. We are prohibited from lobbying, do not develop or sell products, have no owners or shareholders, and do not compete with industry—allowing MITRE's efforts to be truly objective and data-driven. Our multidisciplinary teams (including engineers, scientists, data analysts, organizational change specialists, policy professionals, and more) are thus free to dig into problems from all angles, with no political or commercial pressures to influence our decision-making, technical findings, or policy recommendations.

MITRE has over 50 years of history of partnering with federal agencies to apply artificial intelligence (AI) and machine learning (ML) to support and accelerate agency missions. Our team's deep experience with meaningful AI adoption across the AI/ML life cycle in critical agency mission spaces allows us to anticipate and address future needs that are vital to the efficient functioning of our government and innovation leadership of our economy.

Executive Summary

The United States must take decisive action to unleash the transformative opportunities of AI for the American people while addressing critical challenges posed by countries like China that are aggressively advancing AI to control data, promote state-led initiatives, influence markets, and create military advantage. It is imperative for the U.S. government to foster accelerated AI innovation and strategic deployment, rather than stifle it, by optimizing the regulatory framework to ensure these technological capabilities align with American interests and values. Now is the time for the United States to secure its position as the global leader in this transformative field by promoting AI innovation, capability development, and investment in partnership with industry and academia.

AI is a broad and evolving field, with capabilities supporting various potential applications at different stages of progression from invention to marketplace. The different stages of maturity of specific AI technologies, such as generative AI, neuromorphic computing, multimodal sensing and processing, autonomous systems, natural language processing, and agentic AI play a crucial role in determining AI's readiness to support different operational needs in sectors like defense, manufacturing, financial services, and healthcare. MITRE has identified four categories of action, or "Acceleration Levers" that can accelerate technological advancement from the lab to the marketplace, each being applied at different stages of a technology's evolutionary timeline. The AI Action Plan should incorporate an assessment of the maturity of various AI technologies and their proximity to market readiness, then promote the appropriate technology-accelerating lever(s).

In our response, MITRE further focuses on four high-level, interconnected key areas:

- **Accelerate AI innovation with public-private partnerships**: Foster AI innovation through robust public-private partnerships that accelerate the development and deployment of AI technologies. By leveraging the combined strengths of industry, government, and academia, the United States can prioritize the creation of national AI infrastructures to support large-scale AI advancements. These efforts would specifically address the drivers of our nation's economic competitiveness and the needs of government operations. This collaborative approach should be guided by shared visions articulated through Grand Challenge[1] problems, which would be collaboratively developed by industry, government, and other stakeholders. These challenges serve to unite diverse participants in pursuit of common goals.
- **Lower adoption barriers so that AI can be leveraged to transform industries**: Invest in and promote AI assurance[2] as an enabler so that the full potential of AI can be unlocked for the public benefit. Encourage early adoption of AI technologies by offering incentives and establishing legal frameworks that provide support and protection for early adopters, ensuring they can innovate without undue risk. Identify and reduce policy barriers that are overbearing and stifle innovation, such as overly restrictive data management policies leading to unnecessary data silos within the government and industry, and outdated procurement rules that lead to unresponsive acquisition cycles, while establishing robust and enabling policies in areas lacking regulation to prevent gridlock and foster innovation. The federal government can also play a pivotal role by supporting pilot projects that generate valuable insights and interest; for example, in use cases where there is not a clear market incentive for the private sector. These pilots can demonstrate AI's transformative potential and encourage broader adoption.
- **Secure American AI:** Secure innovation requires a comprehensive approach to ensuring physical security for AI data centers and implementing robust cybersecurity safeguards for AI data, models, and intellectual property. Such a comprehensive approach also involves partnering with trusted suppliers to ensure the supply of critical components such as chips, software, and communication infrastructure. Additionally, security features must be integrated into chip architecture and manufacturing processes. Expanding commercial security clearances in key industries is crucial for securing American AI innovation. It ensures a vetted workforce capable of safeguarding sensitive technologies and intellectual property, extending protection beyond federal employees and contractors. To promote effective and trusted AI innovation, the United States should bolster AI security practices with timely information sharing of potential AI system vulnerabilities and adversarial attacks. This also requires facilitating collaboration (e.g., incident sharing) among industry players who are otherwise in healthy competition.
- **Build the American workforce to drive and harness AI innovation opportunities:** To prepare the workforce for transformative AI, the government should collaborate with

[1] Grand Challenges are ambitious, large-scale goals set by the federal government to drive innovation, collaboration, and measurable progress in addressing complex national problems.

[2] AI assurance involves the assessment and monitoring of AI models to evaluate their trustworthiness and risks. It plays a crucial role in avoiding missteps and maintaining trust and accountability, thereby accelerating adoption. Additionally, AI assurance helps protect U.S. innovations from adversarial attacks and sets a standard for auditability and transparency that other nations can invest in and follow.

industry and academia to cultivate a multi-faceted talent pool.[3] This includes: 1) developing and implementing comprehensive AI strategies that speak to the sector-specific business needs of organizations; 2) developing deep expertise for groundbreaking AI innovation beyond data and computing scale; 3) equipping manufacturing workers with practical AI skills and hands-on experience with AI; 4) training business leaders to navigate and drive AI technology transitions; and 5) empowering all Americans to understand and leverage AI applications. Targeted training and retraining programs, built through public-private partnerships, will address these emerging skill needs, ensuring a workforce ready to lead in critical AI domains. These steps will reinforce U.S. dominance in global science and technology, preparing citizens for an AI-integrated future.

The United States must build on its momentum to lead the world in AI, driving economic growth and strengthening national security through rapid, market-driven innovation. By accelerating AI innovation through public-private partnerships, lowering adoption barriers, securing the innovation ecosystem, and building a skilled workforce, the United States can maintain its global AI leadership. The U.S. must also not fall pretty to risk-aversion, and reward pioneering AI applications where informed adoption decisions have been made in light of justified risk assessments. These interconnected actions will not only drive technological advancement but also ensure that AI development aligns with our national interests. By fostering collaboration among government, industry, and academia, and by clearing obstacles to innovation while applying minimal, targeted rules to ensure AI's practical success, the United States can unleash AI's power to drive economic growth and secure its leadership in the global technology race.

Accelerating the Advancement of AI

In the paper *Partnerships to Accelerate Advancement of Priority S&T,*[4] MITRE analyzed the evolutionary timeline of Critical and Emerging Technologies (CETs) such as AI. MITRE identified four categories of actions (Stimulate, Mobilize, Demonstrate, and Engage), referred to as "levers." Execution of these levers strategically depends on public-private collaboration to rapidly accelerate a CET's advancement, such as AI technologies.

[3] See Executive Order 14179, "Removing Barriers to American Leadership in Artificial Intelligence" (January 23, 2025), which prioritizes U.S. AI dominance through deregulation and public-private collaboration.

[4] 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.

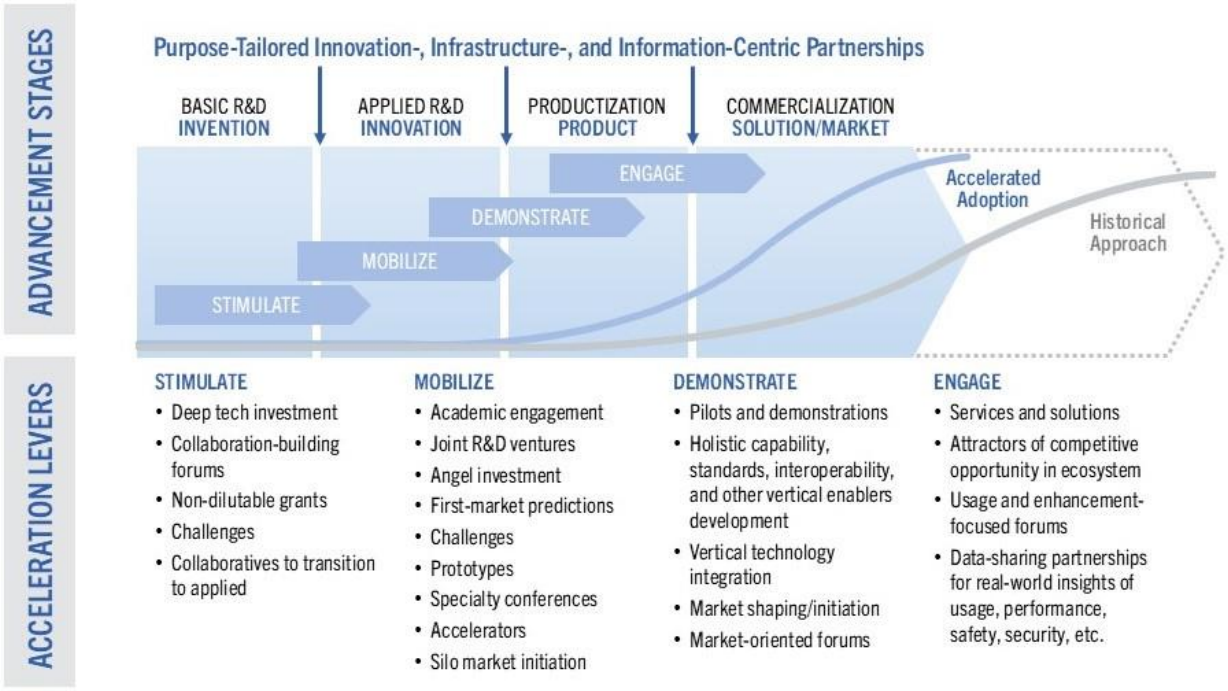

As depicted in the figure above, CETs evolve and mature along a timeline of stages spanning invention to the marketplace. Various AI technologies are at different stages of maturity along this timeline. These Advancement Stages are:

- **Basic Research and Development (R&D) (Invention):** This stage involves researching cutting-edge areas like neuromorphic computing and quantum algorithms, where the fundamental concepts are still being explored or theoretical work is ongoing.
- **Applied R&D (Innovation):** Here, research begins to translate into practical applications, such as developing AI models to predict the efficacy of potential drug candidates.
- **Productization (Product):** Many AI technologies have reached this phase, where they are being turned into tangible products. Examples include generative AI for natural language processing within customer service bots, AI-enhanced business intelligence tools, or machine vision systems for quality control in manufacturing.
- **Commercialization (Solution/Market):** AI technologies like recommendation engines in e-commerce, AI for financial market analysis, and diagnostic tools in healthcare are well into this stage, experiencing widespread market adoption and integration into existing systems.

AI technologies thus span multiple stages of the technology industry evolutionary timeline, reflecting AI's multidimensional, crosscutting impact across various sectors. The AI Action Plan needs to include a focus on Acceleration Lever activities, as described below, that are specific to accelerating the advancement of AI technologies.

Lever 1: Stimulate Research and Create Interest

(a) Lever Objective: To reduce the institutional risk associated with high-risk R&D by incentivizing community-wide investment and creating a fertile ground for new ideas in AI.

(b) Relevant Example AI Research Areas:

- Quantum AI: Exploring how quantum computing can enhance AI capabilities.
- Neuromorphic Computing and other alternatives to the transformer-based architectures: Mimicking biological neural networks for more efficient AI processing.

(c) Needed Public-Private Lever Activities:

- Federal Grants and Tax Incentives: Incentivize specific target industries to engage in basic research in high-reward areas like quantum AI, positioning them for early adoption through targeted tax benefits and grants.
- Research Challenge Problems: Organize national Grand Challenges through collaborative partnerships involving industry, government, and academia to spur innovation in complex AI areas.

(d) Federal Government's Role:

- Investment in Research: Increase funding for basic AI research, particularly in areas with long-term national benefits.
- Policy Support: Create policies that facilitate collaboration between academia, industry, and government labs, ensuring intellectual property rights are protected to encourage participation.

Lever 2: Mobilize a Network (Active Ecosystem)

(a) Lever Objective: To transition AI from research to real-world applications by engaging users early in the development process.

(b) Relevant Example AI Subindustries:

- AI in Drug Discovery: Developing AI models to predict drug efficacy and accelerate pharmaceutical innovation.
- Autonomous Vehicles (AVs) and Drones: Developing AI-enabled systems for navigation, safety, and decision making.
- AI in Cybersecurity: Enhancing threat detection and response with AI.

(c) Needed Public-Private Lever Activities:

- Innovation Hubs: Establish AI-focused ecosystems where startups, tech companies, and government can connect and collaborate on practical implementations.
- Workshops and Conferences: Facilitate knowledge exchange to align market needs with AI innovations.

(d) Federal Government's Role:

- Facilitation of Ecosystems: Support and fund innovation hubs that provide test environments for rapid prototyping and explore collaborative policy frameworks.
- Regulatory Clarity: Provide clear regulatory guidelines to help industries rapidly integrate AI solutions, particularly in sensitive applications such as AVs.

Lever 3: Demonstrate Impactful Solutions

(a) Lever Objective: To speed up market entry by demonstrating AI's effectiveness in practical settings.

(b) Relevant Example AI Subindustries:

- AI for Defense and National Security: Applications ranging from operational planning to autonomous systems in defense.
- Healthcare AI: Diagnostic tools, personalized medicine, and patient management systems.

(c) Needed Public-Private Lever Activities:

- Government-Led Pilot Projects: Use federal assets for real-world demonstrations of AI technologies in sectors such as defense, healthcare, manufacturing, and energy.
- Standards Development: Collaborate on standards that ensure AI's interoperability and effectiveness, accelerating innovation, improving efficiency, and protecting American innovation by establishing a reliable framework for development and deployment.

(d) Federal Government's Role:

- Access to Resources: Provide data sets, testing environments, and pilot opportunities for AI solutions to industry, academia, and research institutions.
- Publicity and Adoption: Promote successful demonstrations to encourage broader adoption and investment.

Lever 4: Increase Business/Industry Engagement

(a) Lever Objective: To facilitate broad market adoption as AI technologies mature, focusing on operational capabilities and market expansion.

(b) Relevant Example AI Subindustries:

- Manufacturing AI: Encourage the adoption of AI-driven automation and predictive maintenance to enhance productivity and efficiency in manufacturing processes.
- Financial Services AI: AI for fraud detection, trading algorithms, and customer service.

(c) Needed Public-Private Lever Activities:

- Policy and Market Access: Develop initiatives that ease AI integration into business operations, such as support programs that provide subject matter experts to assist businesses in effectively implementing AI technologies.
- Industry Standards: Collaborate with industry to set standards that facilitate AI's integration into existing systems, recognizing that standards are a battleground in strategic competition. These standards can influence industry investment, accelerate time to profitability, and shape the global competitive landscape.

(d) Federal Government's Role:

- Policy and Regulatory Frameworks: Develop flexible policies and regulations that support AI adoption, protect innovation, and ensure appropriate oversight. These frameworks should adapt to evolving technological landscapes and address intellectual property and data rights effectively.

- Economic Incentives: Offer targeted incentives for businesses to adopt AI, such as tax credits for research and development, grants for AI integration projects, and subsidies for workforce training. These incentives aim to support national Grand Challenges and enhance economic growth and competitiveness by encouraging widespread AI adoption.

Coordinating federal support for these public-private activities will be crucial for their success. By aligning resources, policies, and strategic initiatives, the federal government can effectively stimulate AI innovation, mobilize networks, demonstrate impactful solutions, and increase business engagement. This coordination ensures that AI advancements are not only accelerated but also strategically aligned with national interests. The National Science and Technology Council (NSTC) has historically played a role in similar coordination efforts, but reimagining this NSTC role could enhance its capacity to support these activities even further. This is discussed in the MITRE paper *A National Science and Technology Council for the 21st Century*.[5]

Accelerate AI Innovation with Public-Private Partnerships

To maintain its leadership in the rapidly evolving field of AI, the United States must foster robust public-private partnerships that accelerate the development and deployment of AI technologies. These partnerships are essential for leveraging the combined strengths of government, industry, and academia, ensuring that AI advancements align with national interests and values. These efforts will ensure the United States outpaces adversaries like China, which are leveraging state-led AI initiatives to dominate data and military applications.

Infrastructure Investment and Market Innovation

To accelerate AI innovation, it is essential to invest in the necessary infrastructure, including large-scale compute and data resources vital for training and deploying advanced AI models. Project Stargate represents a significant private sector investment in data center and compute resources for its partners, and there are other emerging investments in gigawatt-scale data center campuses by other companies. These private investments aim to support commercial frontier labs in advancing research and development toward artificial superintelligence.

While Stargate builds its infrastructure, there is an immediate, near-term need for the U.S. government to access these resources at scale. Investing in commercially packaged, self-contained, exascale AI-computing pods could address this need. These computing pods can be integrated into high-performance computing (HPC) institutions and environments to create functioning innovation clusters within 12 months of authorization. Such investments would have significant impacts, including enhancing national security by supporting defense research, protecting critical infrastructure, and safeguarding American intellectual property. Additionally, they could stimulate manufacturing related to national security and contribute to advancements in space exploration and other strategic sectors. Potential investments could include:

- Open science clusters at Department of Energy (DOE)national labs
- Restricted (controlled unclassified) clusters at Department of Defense (DOD) service labs or FFRDCs and University Affiliated Research Center

[5] A National Science and Technology Council for the 21st Century. 2021. MITRE, https://www.mitre.org/sites/default/files/2021-09/pr-21-2388-national-science-technology-council.pdf.

- Classified clusters at national security data centers

To further catalyze AI advancement, the United States must also provide enabling infrastructure for industry—such as scalable platforms, high-quality data resources, and collaborative ecosystems—while fostering new markets and innovations through standards and consortia. By establishing shared AI development platforms and data repositories, the federal government can empower industry to build and scale cutting-edge solutions. Simultaneously, consortia of government, industry, and academic leaders can drive the creation of interoperable standards that reduce market entry barriers and unlock novel applications, such as AI-driven sustainability solutions or next-generation healthcare diagnostics. These efforts will not only accelerate innovation but also position American industries to lead in emerging global markets.

Federal Frontier Labs

The federal government should prioritize the establishment of a robust, dedicated research program focused on advancing the state of the art beyond current, commercially viable models. To this end, MITRE believes that the establishment of Federal Frontier Labs (FFLs) can be a cornerstone of public-private partnerships in AI. These labs should integrate commercial frontier labs and existing government and academic expertise in AI and HPC to serve as innovation hubs. FFLs would tackle problems relevant to society and the U.S. government, particularly in areas with less immediate commercial interest or opportunity.

FFLs should focus on three primary areas: open science, defense, and intelligence. In open science, FFLs can build on initiatives like the DOE Frontiers in Artificial Intelligence for Science, Security, and Technology (FASST)[6] to advance use of AI for research in physics, healthcare, biology, and more. In defense, FFLs can enhance capabilities in logistics, predictive maintenance, and cyber operations. In intelligence, FFLs can drive advancements in open-source intelligence, imagery intelligence, and signals intelligence.

FFLs should also tackle research topics critical to federal needs. This includes developing low-SWaP[7] models for edge and embedded deployments, exploring federated learning for protecting intellectual property and privacy-preserving training, and enhancing security for AI models through encryption and watermarking. By addressing these topics, FFLs will ensure that AI advancements meet specific federal use cases and maintain U.S. leadership in AI.

Strategic Applications and Collaborative Vision

A strategic focus on AI applications is essential for guiding AI innovation and ensuring its alignment with national priorities. By identifying and prioritizing key applications, FFLs can direct their efforts toward areas that offer the greatest potential for impact. These applications should be informed by a collaborative vision that unites government, industry, and academia in pursuit of common goals.

MITRE has advanced proven methods for quickly uniting parties, even amidst competitive tensions and regulatory challenges, to establish shared expectations for mutually beneficial

[6] See https://www.energy.gov/fasst.

[7] Size, Weight, and Power (SWaP).

collaboration that protects their interests, and to create solutions that surpass what any single party could achieve independently.[8]
Grand Challenge problems serve as a unifying force, bringing together capable stakeholders to address mission-relevant areas.[9] To maximize impact, and integrate this effort with other administration priorities, we recommend prioritizing applications such as:

- Biotechnology: Leveraging AI to accelerate drug discovery, biomanufacturing, personalized medicine, and genomic research, addressing critical health challenges like chronic diseases and pandemics
- Semiconductors: Using AI to optimize chip design, improve manufacturing efficiency, and strengthen domestic supply chains, bolstering national security and economic competitiveness
- Domestic Energy: Applying AI to optimize energy production, enhance grid reliability, and accelerate the development of next-generation energy sources like nuclear and renewables, ensuring energy independence 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
- Cybersecurity: Harnessing AI to detect threats, protect critical infrastructure, and respond to evolving risks in real time, safeguarding both public and private sectors

These Grand Challenge problems can drive progress and innovation, ensuring that American AI innovation supercharges the American economy and advances national security. Proven methods for multi-sector collaboration at scale will not only accelerate technological advancement but also ensure that AI development benefits all Americans.

Mission Engineering Integration

To ensure AI technologies effectively meet mission-critical needs, MITRE emphasizes the importance of mission engineering as a structured approach to bridge the gap between the mission demand perspective and the technology impact perspective. This process involves analyzing agency-specific requirements—such as national security, public health, or infrastructure resilience—and translating them into actionable technology development priorities that address operational mission needs. By applying mission engineering within Federal Frontier Labs and public-private partnerships, we can ensure that AI solutions are not only technologically advanced but also directly responsive to the operational demands of federal missions, maximizing their societal and strategic impact.

High-Quality Data Sets

Creating, collecting, and curating high-quality data sets is essential for advancing AI capabilities. These data sets will support the development of AI models that are more accurate, reliable, and

[8] See https://assemble.mitre.org/.

[9] 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.

applicable to real-world scenarios. Public-private partnerships can facilitate the sharing of data resources, ensuring that AI research is grounded in robust and application-focused data.

Lower Adoption Barriers So That AI Can Be Leveraged to Transform Industries

To fully harness AI's transformative potential, the United States must lower barriers to adoption by tackling the lack of expertise in AI operations, clarifying its business impacts—like workforce shifts and public trust—and minimizing regulatory roadblocks. Through strategic initiatives that accelerate integration and use, AI can drive innovation and fuel economic growth.

Pilot Projects and Demonstrations

Government-championed and industry-led pilot projects have the potential to play a crucial role in showcasing AI's transformative potential and understanding business implications. While currently underutilized, these projects can demonstrate practical AI applications in areas lacking clear market incentives, such as AI-driven healthcare diagnostics to reduce costs, smart energy grids for efficient cities, and autonomous transportation for logistics breakthroughs. By leveraging federal assets and resources, these pilots can generate insights and interest, proving AI's real-world benefits. Such initiatives can inspire confidence and broader adoption by highlighting AI's power to enhance efficiency, productivity, and decision making across sectors.

Prototyping to Bridge Mission Gaps and Accelerate Industry Adoption

Beyond showcasing potential, prototyping can serve as a powerful mechanism to demonstrate the feasibility of closing critical operational gaps while driving the creation of new programs and requirements. By developing and testing AI prototypes within federal contexts, the government and its partners can validate solutions that address specific agency needs, motivating stakeholders to establish formal initiatives to scale these capabilities. Once proven, these prototypes can be transitioned to industry on behalf of the government, serving as reference implementations or accelerators that reduce development timelines and costs, enabling rapid commercialization and widespread adoption of mission-aligned AI technologies.

Architectures for Competition and Interoperability

To maximize industry participation and innovation, the United States should prioritize developing AI architectures that enable competition and plug-and-play capabilities at the appropriate level of granularity. These architectures—designed with modular, interoperable interfaces and components—allow a wide range of industry players to integrate their solutions without being locked into a single provider's ecosystem. By collaborating with industry and standards bodies, the government can ensure that these frameworks balance openness for competition with the cohesion needed for seamless functionality, such as enabling small businesses to contribute specialized AI modules alongside larger firms. This approach prevents monopolistic lock-in, fosters a vibrant market of interoperable AI solutions, and accelerates the deployment of transformative technologies across sectors.

AI Governance, Assurance, and Risk Mitigation

AI governance, assurance,[10] and risk mitigation are vital to unlocking AI's full potential, driving rapid adoption, and securing U.S. leadership. Agile governance—co-designed with industry—streamlines decision making and ensures AI delivers results, not rework, across its life cycle. Assurance builds confidence by engineering reliability and validity into AI across the AI lifecycle, paving the way for effective applications in healthcare, finance, and transportation. Risk mitigation keeps projects on track, cutting costs and setbacks that could slow progress. By championing these lean, practical processes with minimal federal footprint, the government can turbocharge industry innovation, boost economic growth, and prove AI's transformative power—all while keeping America ahead in the global tech race.

Policy and Regulatory Environment

Creating a supportive policy and regulatory environment is essential for fostering innovation while maintaining accountability. It is important to reduce policy barriers that stifle innovation and establish robust policies where none exist to prevent gridlock. For example, policies should streamline outdated procurement rules or data-sharing restrictions that delay AI deployment, replacing them with agile, innovation-friendly guidelines. Additionally, addressing risk aversion is crucial; if we do not mitigate this, we risk being left behind in the global AI landscape. By fostering an environment that encourages informed risk-taking and supports early AI adoption, we can ensure that innovation thrives and keeps pace with international advancements..

System Auditability and Transparency

Promoting system auditability and transparency is key to maintaining trust and accountability in AI applications. Developing standards and practices for auditing AI systems ensures that their operations are transparent and accountable. By providing visibility into how AI works and makes decisions, we can reinforce confidence in AI technologies. Transparency initiatives can also facilitate collaboration and knowledge sharing, enabling stakeholders to work together to improve AI-enabled systems and address emerging challenges. This openness can lead to more innovative applications of AI, as stakeholders feel more secure in exploring new possibilities.

Secure American AI

As the United States continues to lead in the development and deployment of AI, it is imperative to secure the innovation ecosystem against a wide range of threats. 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 secure military and economic advantages, often employing

[10] 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.

aggressive tactics that threaten global technological leadership and security.[11,12,13] By implementing comprehensive security measures, we can protect AI advancements and ensure they align with national security interests.

Securing AI Systems

To protect AI innovation, robust cybersecurity is essential for safeguarding AI systems, data, and intellectual property from cyber attacks, breaches, and theft. This means deploying advanced threat detection, incident sharing, rapid incident response, and continuous monitoring to keep AI models and applications secure. Looking ahead, investing in post-quantum cryptography will shield AI from emerging quantum threats, ensuring long-term resilience. These measures—driven by industry best practices—build trust in AI's reliability, enabling its bold use in critical areas like healthcare and finance while keeping U.S. tech ahead of global rivals.

Securing AI Infrastructure

The backbone of AI advancement rests on securing its physical infrastructure—like data centers, communication lines, and power and water sources. This requires strong physical security to protect facilities from unauthorized access or disruption, paired with cybersecurity protocols to block digital threats. Data centers, in particular, need secure access controls, regular audits, and operational resilience to safeguard the computational resources powering AI. By prioritizing this lean, practical security—backed by federal assets and industry standards—the United States can ensure its AI infrastructure drives economic growth and national strength.

Integrating AI Assurance with Existing Guidelines

Integrating AI assurance with existing security risk management frameworks, like National Institute of Standards and Technology (NIST) 800-53, is key to a cohesive security strategy that gets AI-enabled systems operational quickly. Pairing this with threat-based approaches, such as MITRE's Adversarial Threat Landscape for Artificial-Intelligence Systems (ATLAS™),[14] strengthens the mix—using proven controls to spot and mitigate AI-specific risks while tapping real-world threat intelligence to counter attacks like model tampering or data breaches. This dual approach, aligning with industry standards and adversary-focused insights, ensures AI systems are reliable, resilient, and ready to drive U.S. innovation, keeping us ahead of global competitors.

Supply Chain Security and Trusted Suppliers

Securing the supply chain is vital to protecting the components that underpin AI innovation. Partnering with trusted suppliers for critical components, such as chips, software, and communication networks, is essential. Federal guidelines and implementation support for supply

---

[11] 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.

[12] A Vision for Competitiveness: Mid-Decade Opportunities for Strategic Victory. 2023. Special Competitive Studies Project, https://www.scsp.ai/wp-content/uploads/2023/04/Vision-for-Competitiveness-1-1.pdf.

[13] 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/.

[14] See https://atlas.mitre.org/.

chains backing frontier AI programs, including supply chain vetting, should be established, similar to DOD facilities clearances but tailored for commercial entities. Industry partners should co-develop supply chain security standards with government agencies, leveraging their expertise to ensure trusted supplier networks.

Commercial Security Clearances

Expanding commercial security clearances is necessary to develop a vetted workforce that extends beyond federal employees and contractors. Federal guidelines and implementation support for suitability and counterintelligence vetting programs should be designed for staff working in critical industrial programs as well as FFLs. These programs, similar to security clearances in the national security space, but specifically geared for work outside the federal government and its contractors, will ensure that individuals involved in AI development and deployment are thoroughly vetted and trusted. We recommend that federal guidelines for commercial security clearances and supply chain security should be finalized in 2025, with full implementation across government AI innovation partners by 2026. This initiative can be seen as an extension to National Security Presidential Memorandum-33.

Information Sharing and Threat Intelligence

Timely information sharing on potential AI vulnerabilities and adversarial attacks is crucial for promoting safer and more reliable AI innovation. MITRE's ATLAS framework plays a critical role in the overall AI landscape by providing a comprehensive understanding of adversarial threats, enabling developers to anticipate and mitigate potential vulnerabilities in AI systems. By offering insights into real-world attack tactics and techniques, ATLAS helps industry develop AI technologies that are inherently secure and resilient. Fostering transparent collaboration between government and industry further facilitates the exchange of threat intelligence, enhancing the ability to proactively address emerging security challenges.

Counterintelligence and Adversary Monitoring

A comprehensive approach to counterintelligence is necessary to prevent adversaries from exploiting AI advancements. Counterintelligence efforts should target specific risks, such as China's state-sponsored intellectual property theft, to protect U.S. innovations. Monitoring adversarial AI tradecraft and understanding how adversaries may use AI to gain an advantage are critical components of this strategy. By staying vigilant and informed about adversarial activities, we can protect AI innovations from being compromised and ensure that they are used to advance national interests.

Threat-Informed Approach to Export Controls

To keep AI technologies out of adversaries' hands, a threat-informed approach to export controls sharpens our edge over the current system. Unlike static lists, this leverages real-time intelligence and threat assessments to pinpoint risks—think AI model theft or weaponization—while balancing national security with U.S. economic strength. Controls would adapt as threats evolve, but with clear, predictable updates to avoid bogging down industry. Partnering with tech leaders, the United States can craft lean policies that block critical tech from rivals like China without choking American innovation, ensuring our AI stays ahead and our economy thrives.

Resilience and Recovery from AI Incidents

A national AI incident response strategy—built with industry input—must tackle two threats: attacks on U.S. AI capabilities and AI-driven assaults on the nation. For attacks targeting AI systems, like hacks on data centers or model theft, the focus is fast recovery—restoring integrity and protecting IP to keep our tech edge sharp. For AI-enabled attacks, like adversaries crashing infrastructure with smart malware, the priority is rapid containment and public confidence, shielding critical systems like power grids. This lean, dual-track approach ensures U.S. AI stays resilient, drives economic strength, and outpaces rivals without bogging down industry.

Preventing Malicious Use

Refining policies is critical to stop AI misuse and crack down on bad actors, keeping U.S. tech safe and strong. Clear guidelines—built with industry input—must separate legit research from malicious use, like AI-generated deepfakes spreading false information, AI-powered cyber attacks crashing our grids, or AI stealing American tech secrets. These lean rules set tough legal lines, ensuring accountability and deterrence without slowing innovation. By targeting these threats, we protect public safety, national security, and our economic edge, keeping AI a force for U.S. leadership.

International Collaboration

The United States should collaborate with allies to share threat intelligence and align supply chain security standards, amplifying our collective resilience. International cooperation can enhance America's ability to secure AI innovation and provide a global advantage in the competition with China.

Build the American Workforce to Drive and Harness AI Innovation Opportunities

As AI continues to transform industries and redefine the future of work, it is crucial to build a workforce capable of driving and harnessing these innovation opportunities. By focusing on strategic initiatives in education and training, we can ensure that the American workforce is prepared to lead in the AI economy and maintain a competitive edge on the global stage.

Enhancing STEM Education

The foundation of a capable AI workforce begins with enhancing science, technology, engineering, and mathematics (STEM) education. The federal government can play a supportive role by encouraging states and private industry to integrate AI concepts into STEM curricula from an early age. This involves curriculum changes, teacher training, and studies on how to approach AI education at various ages. By incorporating AI into STEM education, we can ensure that students are equipped with the knowledge and skills needed to thrive in a technology-driven world. A skilled AI workforce will help ensure the United States outpaces China's state-driven talent pipeline.

AI Literacy and Skills Development

In addition to foundational STEM education, it is important to build AI literacy and skills across the workforce and general population. This involves teaching individuals how to manage, supervise, and interact with AI technologies, ensuring they are equipped to work and live

alongside AI. By promoting AI literacy, we can empower all citizens to leverage AI tools effectively and enhance their productivity. Initiatives that focus on developing these skills will be crucial for ensuring that the workforce is ready to embrace AI technologies.

Targeted Training and Public-Private Partnerships

To address emerging AI skill requirements, targeted training programs must be developed through collaborations between government, industry, and academia. Public-private partnerships, including academia, can leverage the strengths of each sector to create comprehensive training and education programs that align with the needs of the AI economy. Industry partners could co-design training curricula, fund apprenticeships, and commit to hiring program graduates. These programs should focus on preparing workers for roles in critical AI areas, such as data science, machine learning, and AI system design, while also anticipating emerging fields like quantum AI. These training programs should include certificate and non-degree programs. By fostering collaboration, the United States can ensure that training initiatives are responsive to industry demands and provide workers with the skills needed to succeed in AI-related careers.

Retraining and Upskilling Initiatives

As AI technologies continue to evolve, retraining and upskilling initiatives will be essential for helping workers transition into AI-related roles. These initiatives should prioritize equipping workers with the skills needed to adapt to changing job requirements and seize new opportunities in the AI economy. Retraining should leverage online platforms and community college partnerships, ensuring accessibility for rural and low-income workers. By providing access to continuous learning and professional development, the United States can cultivate a skilled workforce ready to lead in AI innovation. Retraining programs should be designed to be flexible and accessible, allowing workers to acquire new skills while balancing their existing responsibilities.