



## TechNet's Federal AI Policy Framework

Artificial intelligence (AI)<sup>1</sup> is a transformational technology that has the potential to revolutionize how we live and work and help us solve the most significant challenges of our time. AI can enhance productivity, democratize and expand access to important services, and improve product innovation. TechNet members represent many of the leading AI and automated systems developers, researchers, deployers, and users.

In an era of rapid technological advancement, it has become imperative for federal policymakers to navigate the complex landscape of innovation and regulation. This comprehensive policy framework comprises five distinct sections, each addressing critical facets of this evolving ecosystem. From deploying risk-based regulations to fostering responsible AI evaluations, mitigating potential bias, securing advanced systems, and building a resilient innovation workforce, our recommendations are the result of collective expertise and commitment to shaping a forward-looking, prosperous future for our nation.

### Deploy Risk-Based Regulations

- As policymakers consider new regulations for AI, it is important to note that there are already existing laws and regulations that prohibit the use of AI in furtherance of unlawful behavior. Many existing civil rights laws apply to AI models used in education, healthcare, employment, housing, financial services, and accessing goods and services.<sup>2</sup> Such laws and regulations, which benefit from existing well-developed regulatory and enforcement frameworks, focus on preventing and providing recourse against the prohibited conduct rather than the means by which the conduct was accomplished.
- Any new laws or regulations, as well as guidance documents and enforcement statements, should focus on known or rationally-anticipated harms that could be prevented or addressed by filling gaps in existing legal regimes.
  - It is crucial for policymakers to recognize the diverse array of stakeholders involved in AI systems, including developers, researchers, deployers, and users. Careful consideration must be given to designating regulatory responsibility that aligns with the roles and interactions of these entities.
  - The AI startup ecosystem is vital to maintaining America's competitive edge. Potential implications for small and mid-size businesses must be considered.
  - Any new regulations should be subject to existing Regulatory Impact Assessment analyses.

---

<sup>1</sup> The National Institute of Standards and Technology (NIST) [Artificial Intelligence Risk Management Framework](#) (AI RMF 1.0) defines an AI system as an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy.

<sup>2</sup> Several existing enforcement statutes were outlined in the [National AI Advisory Committee's Year One Report](#): *Civil Rights Act of 1964, Equal Educational Opportunities Act, Americans with Disabilities Act, Individuals with Disabilities in Education Act, Genetic Information Nondiscrimination Act, Immigration and Nationality Act's Anti-Discrimination Provision, Fair Housing Act, Equal Credit Opportunity Act, Violent Crime Control and Law Enforcement Act, and the Omnibus Crime Control and Safe Streets Act.*

- Efforts to require approval for commercial AI systems by a federal/state/local government agency should be calibrated to the level of risk that the intended use case poses, consistent with any new AI frameworks applicable to the private sector. Overly broad requirements to gain government approval will likely entrench leading existing players and stifle innovation to the detriment of America's global leadership and American consumers.
- We believe there should be a central coordinator of the federal government's development, deployment, and use of AI systems and ensure that AI policy and regulations are consistent across agencies and industries. This coordinator should ensure that AI policies are risk-based and the regulations that actors are subject to are based on the level of risk the AI use case entails and not just what regulatory body may claim authority over an entity. The agency must coordinate with existing subject matter agencies on particularly complex or technical use cases that may benefit from specialized expertise.
- The National Institutes of Standards and Technology (NIST) AI Risk Management Framework (AI RMF 1.0) should be promoted as a voluntary model for AI lifecycle management, including design, development, deployment, and post-deployment.
- Private rights of action have the potential to undermine innovation and subject small and large businesses to abusive and frivolous litigation tactics, and therefore, must be avoided.
- Existing enforcement mechanisms and protections from intermediary liability should be utilized to address AI enforcement challenges.
- A consistent risk-based federal AI framework is needed to provide clear guidance and to prevent a patchwork of differing state laws that could impede innovation and progress. A consistent and level playing field for all entities developing and deploying AI is essential.
- Establish a national privacy standard to promote consistent regulation of Americans' data. The passage of a federal consumer data privacy law is an essential component of a coherent national AI-focused policy. A comprehensive federal privacy law will help consumers understand their data rights and will assist developers in knowing their liability when managing large datasets. A clear national framework will help build trust in AI systems. To learn more about TechNet's principles on privacy, [please visit here](#).

### **Responsible AI Evaluations**

- Any transparency, explainability, or audit requirements imposed on AI systems must account for protecting the personal information of consumers, as well as proprietary and trade secret information regarding the AI system and the technical feasibility of implementing such requirements. It must also not jeopardize the safety systems of AI-driven services.
- Leading AI developers and academics are continuing to research and improve how we can best explain the output of generative AI systems. We encourage the federal government to support continued research and development into best practices for explainability, transparency, and auditing and discourage "one-size-fits-all" regulations as this technology continues to evolve.
- TechNet believes that any regulations requiring enhanced disclosures for users or regulators should apply only to high-risk applications that lack existing regulatory structure to govern situations where the AI system's compromise, misuse, or destruction would be reasonably likely to result in loss of life, liberty, or significant legal effects.

- **Explainability**
  - When considering explainability requirements, TechNet suggests considering NIST's Four Principles of Explainable AI:
    - **Explanation:** Systems deliver accompanying evidence or reason(s) for all outputs.
    - **Meaningful:** Systems provide explanations that are understandable to individual users.
    - **Explanation Accuracy:** The explanation correctly reflects the system's process for generating the output.
    - **Knowledge Limits:** The system only operates under the conditions for which it was designed or when the system reaches sufficient confidence in its output.
  
- **Transparency**
  - We urge policymakers to avoid one-size-fits-all transparency requirements on AI systems. A better approach is to require transparency or audit reports only for high-risk uses of AI.
  - Support public education efforts on how AI systems operate in order to help demystify AI.
  - TechNet supports the disclosure of generative AI content to users. Industry leaders are still researching how to best indicate content has been AI-generated and when such indications are appropriate. We are supportive of this ongoing discussion and research to best inform the American public about the content they are viewing.
  
- **External Reviews**
  - TechNet believes that it is premature to mandate independent third-party auditing of AI systems. There is not currently a well-established credentialing regime for AI auditing, such as exists for other high-impact sectors such as financial services. In some cases, particularly with sophisticated AI developers that have robust AI systems, their internal auditing programs far surpass third-party options. Mandating an independent audit before the market reaches maturity could open AI systems to national security threats, trade secrets theft, and inaccurate audit reports.
    - TechNet supports the White House's voluntary commitments for third-party discovery and reporting of vulnerabilities for generative models that are overall more powerful than the current industry frontier.
  - We believe that AI auditing standards, ethics, or oversight rules must consider the use-case-specific auditing needs, calibrated to the risk of the specific use case, set to measurable benchmarks, and ensure safe and ethical practices to promote continued innovation while also protecting intellectual property.
  - Reciprocity of AI audit findings across local, state, and federal jurisdictions should also be accepted to limit resource burden and sustain market access for the AI startup ecosystem.

### **Mitigate Potential Bias**

- Throughout its lifecycle, AI development must reflect our society's highest ideals, and its performance must be appropriately monitored and evaluated. Measures to identify, track, and mitigate unintended bias and discrimination should be implemented.

- Owners of AI systems should ensure appropriate oversight and accountability to assess the need for improvements to ensure safety, fairness, and trustworthiness; protect against malicious activity; and address flawed data sets or assumptions.
- Importantly, it is worth underscoring that existing anti-discrimination laws apply to AI models in many important contexts, including housing, employment, and consumer financial services (i.e., the *Fair Housing Act*, Title VII of the *Civil Rights Act of 1964*, and the *Equal Credit Opportunity Act*). Therefore, additional legislative and/or regulatory obligations in these areas at this time would be unnecessarily duplicative, create inconsistent or conflicting standards, and chill innovation in the U.S. Instead, policymakers should leverage existing tools to address concerns of bias.
- Bias in human processes is well documented but can be difficult to spot until it is too late to correct. By contrast, TechNet members are developing AI systems that can detect and avoid or mitigate bias. TechNet members follow legal guidelines at all stages when developing, testing, and monitoring AI assessments, and in many cases, they test for group differences beyond those required by law.

### **Secure Advanced Systems**

- Leverage security by design principles to enhance cybersecurity within AI systems at the start of their lifecycle.
- Empower America’s cyber defenders by funding the use of AI-enhanced cybersecurity services and tools within the federal government.
- Strengthen the adoption of AI cybersecurity awareness training to help minimize risk and prevent loss of intellectual property, data, or money.
- Support bidirectional information sharing and cyber threat programs accounting for threat actors leveraging AI.

### **Build the Innovation Workforce**

- Support public-private partnerships in establishing and maintaining upskilling programs to help Americans best utilize and improve their productivity with automated tools.
  - Some of these programs will be government-funded and designed, but many companies are already providing useful resources to help Americans advance their careers. Governments at all levels should seek to understand and build on what is already working.
  - Promoting upskilling and investing in workforce programs offer a proactive approach to fostering diversity among AI developers, deployers, monitors, and users. This is a valuable strategy to address bias concerns throughout the AI lifecycle.
- Support government funding for AI safety research and infrastructure
  - Congress must authorize and fund the National AI Research Resource (NAIRR). The NAIRR is important to foster the development of the U.S. domestic AI research ecosystem and maintain U.S. leadership in AI on the global stage.
  - Most of the world’s leading AI developers are outside of government institutions; governments need to engage these experts by utilizing public-private partnerships to inform the development of regulation and guidance, build modern government AI systems, and incorporate AI efficiencies into government services.

- Support the creation of a science, technology, engineering, and math (STEM) visa for foreign students who have earned Master's level or higher degree from U.S. colleges and universities. This would promote economic growth and innovation in AI by ensuring that talented innovators educated and trained in the U.S. can become citizens and create jobs here.<sup>3</sup>
- Support the federal government's strategic hiring of AI experts and the filling of vacant technology roles. Bolstering our federal workforce with needed talent will allow key government agencies to enhance their capacity to monitor, utilize, and ensure responsible and impactful AI development and deployment.
- TechNet supports expanded government utilization of AI to improve access to important services, enhanced efficiency, cost savings, data-driven decision-making, and more equitable and inclusive service provision, ultimately benefiting citizens and society as a whole.
  - TechNet supports the government in developing "AI Ready Data." The United States federal government is one of the biggest producers of data in the world, and these important datasets are already fueling innovation in the public and private sectors. As we move to greater deployment of AI systems, ensuring this data is well-organized will allow these modern tools to deliver faster, cost-effective, and more accurate insights.

---

<sup>3</sup> 28 of 43 (65%) of the top AI companies in the U.S. were founded or co-founded by immigrants; 70% of full-time graduate students in fields related to AI are international students; 18 of 43 (42%) of top U.S. based-AI companies have a founder who studied in the U.S. as an international student.