TECHNET
THE VOICE OF AMERICAN INNOVATION

1420 New York Avenue NW, Suite 825
Washington, D.C. 20005
www.technet.org | @TechNetUpdate

October 27, 2025

Michael Kratsios
Director
Office of Science and Technology Policy
Eisenhower Executive Office Building
1650 Pennsylvania Avenue, NW
Washington, D.C. 20504

Re: Regulatory Reform on Artificial Intelligence (Docket No. OSTP-TECH-2025-0067)

Dear Director Kratsios:

TechNet appreciates the opportunity to respond to the Office of Science and Technology Policy's (OSTP) Request for Information on Regulatory Reform on Artificial Intelligence (AI). TechNet is the national, bipartisan network of technology CEOs and senior executives that promotes the growth of the innovation economy by advocating a targeted policy agenda at the federal and 50-state level. TechNet's diverse membership includes over 100 dynamic American businesses ranging from startups to the most iconic companies on the planet and represents over five million employees and countless customers in the fields of information technology, artificial intelligence, e-commerce, the sharing and gig economies, advanced energy, transportation, cybersecurity, venture capital, and finance.

AI is a foundational element of our modern economy and a critical driver of future growth and security. From pioneering new medical treatments and creating more efficient supply chains to enhancing cybersecurity and accelerating scientific discovery, AI is poised to deliver transformative benefits across every sector of American life. The global race to lead in AI is underway, and the nation that successfully cultivates an ecosystem of innovation while earning public trust will define the technological landscape for decades to come. The United States has long been the world's innovation engine, thanks to a policy environment that encourages investment, research, and entrepreneurial risk-taking. To maintain this leadership in the age of AI, we must ensure our regulatory frameworks are as innovative and forward-looking as the technologies they govern. TechNet commends OSTP's goal of identifying and modernizing outdated federal statutes, regulations, and administrative practices that unnecessarily hinder the responsible development, deployment, and adoption of AI technologies. America's global competitiveness depends on maintaining leadership not only in AI research and innovation, but also in the regulatory systems that govern their use.

TechNet strongly supports technology-neutral, risk-based regulation as the cornerstone of responsible AI governance. New or AI-specific regulation should be considered only where genuine gaps exist in current, technology-neutral legal frameworks**.** The United States already maintains a comprehensive foundation of generally applicable laws that govern product safety, consumer protection, privacy, competition, and civil rights. Many of these laws are deliberately designed to be adaptable to new technologies, including AI. Policymakers should recognize that most AI risks can be effectively managed under existing frameworks, and that new, AI-specific statutes or mandates should be pursued only when demonstrable gaps cannot be addressed through clarification, enforcement guidance, or interagency coordination. Creating overlapping or duplicative AI-specific rules risks fragmenting compliance requirements, increasing costs, and chilling innovation — particularly for small and mid-sized developers. Instead, federal agencies should focus on interpreting and applying existing authorities consistently to AI use cases, issuing clear guidance where needed rather than crafting new, siloed regulatory regimes.

TechNet believes that modernizing, not multiplying, regulation is essential. Effective regulatory reform should emphasize clarity, flexibility, and experimentation, enabling agencies to fulfill their missions while allowing industry to innovate and deploy AI responsibly. Many existing rules inadvertently suppress innovation by mandating human oversight, rigid certification criteria, or data management approaches that no longer match technical reality. A concerted effort to modernize our regulatory approach will unlock the full potential of AI, allowing the United States to lead the world in both innovation and the establishment of thoughtful, effective governance.

The current federal regulatory environment for AI suffers from a number of recurring deficiencies that include regulatory mismatch where rules misalign with adaptive, algorithmic systems; structural incompatibility with statutes written before AI that often require human decision-makers and therefore foreclose automation; and lack of clarity with each uncoordinated guidance across agencies creating compliance uncertainty. These shortcomings — combined with fragmented state laws — slow AI adoption and weaken U.S. leadership. In sectors from healthcare to transportation to telecommunications, outdated frameworks block technologies that could improve safety, efficiency, and service delivery. However, TechNet believes that a focused and collaborative effort between industry and government can create a modern, pro-innovation regulatory ecosystem that is fit for the AI era.

Our core recommendations include:

- **Modernize Legacy Regulations:** Updating existing statutes and rules to be technology-neutral and performance-based, rather than prescriptive.
- **Enhance Regulatory Clarity:** Calling on agencies to issue clear, consistent, and timely guidance on the application of existing rules to AI systems.
- **Champion Regulatory Sandboxes:** Establishing safe harbors and experimental authorities to allow for innovation under regulatory supervision.

- **Invest in Federal Capacity:** Upskilling the federal workforce and embedding technical expertise within agencies to ensure informed and agile governance.
- **Preempt State Regulation:** Establishing a unified federal approach to AI regulation would help address compliance burdens with varying state regulations.
- **Combat International Overregulation:** Driving global consensus in support of a U.S.-led framework for international AI standards and definitions that enables regulatory coherence and global adoption.

Modernize Legacy Regulations

Regulatory barriers often occur when rules and regulations are based on human-centered assumptions that do not align with how AI systems operate in the world today. The underlying regulatory goal is often achievable, but the prescribed process creates an unnecessary obstacle by requiring human action that can otherwise be easily automated. OSTP, working with OMB and the Office of Information and Regulatory Affairs (OIRA), should establish a formal process to identify and review any statutes and rules that were written for static, human-operated systems. This review should prioritize high-impact areas such as transportation safety standards, medical device regulation, privacy and data governance, export control regimes, and federal procurement rules where human-centric requirements continue to impede the use of adaptive, data-driven technologies.

Modernizing legacy regulation requires shifting from prescriptive, process-oriented rules to performance-based frameworks that focus on measurable outcomes. Prescriptive regulation — still common across health, transportation, and financial sectors — dictates the methods or designs entities must use to comply with the regulation. This rigidity may have suited static, human-operated systems but is fundamentally mismatched for adaptive technologies like AI, which evolve continuously and can achieve compliance through multiple technical pathways. Wherever possible, agencies should replace prescriptive, design-based requirements including mandates for manual oversight, paper documentation, or fixed product configurations with performance-based standards that measure outcomes like safety, reliability, effectiveness, and transparency.

For example, in transportation, performance-based safety standards could focus on collision avoidance rates or system reliability benchmarks such as requiring that an autonomous vehicle system demonstrate a statistically equivalent or superior safety record to human drivers, rather than mandating the presence of manual steering controls or specific sensor configurations. Similarly, in healthcare, regulations should focus on patient safety and outcome accuracy rather than specifying the algorithmic techniques permissible for diagnostics or monitoring.

Agencies should also develop adaptive compliance mechanisms**,** such as voluntary certification programs that can accommodate iterative software updates and machine learning model retraining. OSTP and OIRA can facilitate this transition by developing interagency guidance on how to design and evaluate performance-based standards,

including templates for outcome metrics, methods for verifying compliance, and procedures for periodic reassessment as technologies evolve. Embedding performance-based principles into federal regulation will ensure that the government's oversight remains rigorous, relevant, and resilient to technological change. This approach fosters innovation by creating space for new solutions while maintaining clear, enforceable accountability standards and ensuring that regulation remains effective even as technology evolves.

Enhance Regulatory Clarity

In many sectors, the primary barrier to AI adoption is not a specific rule but a pervasive lack of clarity on how a web of existing rules applies. When companies cannot predict how existing laws will be applied to new AI technologies, the legal risk and compliance costs can become prohibitively high, discouraging investment and delaying market entry. This uncertainty forces businesses to rely on expensive and conservative legal interpretations, slowing innovation. The federal government can significantly reduce this friction by providing clear and authoritative guidance on the application of existing rules to AI.

Agencies are best positioned to interpret their own regulations, and each agency should be directed to issue clear and timely interpretive guidance. However, rather than inventing new regulatory guidance and compliance regimes for each agency from scratch, the adoption of common, interagency frameworks should be encouraged. Without harmonization, sector-specific regulations can unintentionally create duplicative or conflicting compliance obligations**.** For example, health data governed under the *Health Insurance Portability and Accountability* Act (HIPAA), and financial data governed under the *Gramm–Leach–Bliley Act* often overlap in real-world applications such as telehealth platforms or fintech products that handle both medical and financial information. Inconsistent or redundant requirements for consent, data retention, and breach notification can force organizations to run parallel compliance systems for the same underlying activity, increasing costs without improving consumer protection. OSTP should encourage agencies to harmonize cross-sectoral regulations by identifying shared policy objectives — like privacy, security, and accountability — and aligning compliance standards and reporting processes wherever the same action is governed by multiple regimes.

As part of this effort, consensus-based technical standards developed by bodies like the National Institute of Standards and Technology (NIST) should be promoted across agencies regardless of sector. For example, the NIST AI Risk Management Framework (AI RMF) provides a robust, voluntary framework for managing risks associated with AI systems. OSTP should encourage all federal agencies to formally recognize the AI RMF as an acceptable methodology for demonstrating due care in AI development and deployment. As part of this process, NIST should publish authoritative crosswalks from the updated AI Risk Management Framework to existing laws, regulations, and safety standards once the regulatory review has been completed. This would help organizations understand how following the NIST framework can meet legal requirements. NIST should also provide a model "presumption of adequacy" so that

regulators can recognize companies that follow the framework as meeting key compliance obligations, creating a more predictable and streamlined compliance path. This would reduce regulatory fragmentation, lower compliance burdens for U.S. innovators by providing clear target for their governance programs, and create a clear and consistent compliance safe harbor across the government.

Additionally, differing definitions of "AI system" and "automated decision" throughout agency regulatory regimes also create redundant compliance efforts. For example, the definitions adopted by NIST, the Federal Trade Commission (FTC), and the Department of Defense (DOD) diverge in scope, thresholds, and covered use cases — forcing companies to interpret and document the same technology differently for each regulator. This inconsistency results in redundant reporting, duplicative risk assessments, and conflicting audit requirements that do little to enhance public trust or accountability. The lack of definitional harmonization also complicates procurement and interagency coordination, as the same system may be categorized as "high-risk" under one framework and "low-risk" under another. OSTP should prioritize the establishment of a unified, technology-neutral federal lexicon for AI-related terms, developed in consultation with NIST and OMB, to ensure consistent interpretation and reduce unnecessary compliance friction while maintaining robust oversight. Again, TechNet recommends the use of the NIST AI RMF's definition of an AI system for two reasons: 1) the RMF was developed through close coordination with the experts from the AI community, and 2) it was adapted from existing AI industry definitions. Adopting the NIST AI RMF definition across the government will help provide greater clarity for the public's understanding of AI systems.

Agencies rely on nonbinding statements rather than formal interpretive rules, leaving innovators uncertain about enforcement expectations. Ambiguity over how existing rules apply to emerging technologies often leads companies — especially startups and small innovators — to delay deployment or over-comply out of caution, diverting resources from research and development toward legal risk management. Agencies can foster innovation by clearly communicating their enforcement priorities and delineating regulatory enforcement actions, including by regularly publishing enforcement priority statements, interpretive guidance, and advisory opinions that clarify how existing authorities will be applied to AI systems. This should also include outlining enforcement triggers and clear processes for voluntary disclosure and corrective action. Predictable, transparent enforcement not only protects the public interest but also builds the confidence necessary for responsible AI investment and deployment.

For example, a clear AI enforcement policy would be particularly valuable at the FTC, which has increasingly applied its broad authority under Section 5 of the Federal Trade Commission Act to algorithmic and AI-driven practices without issuing detailed interpretive guidance. The FTC has brought enforcement actions related to "algorithmic bias," "data misuse," and "automated decision-making," yet companies often lack clarity on how these concepts are defined, what evidentiary thresholds apply, or what constitutes a violation. For instance, the Commission's consent decrees in cases involving AI-enabled consumer scoring and facial recognition systems

established important precedents but offered little prospective guidance for compliant behavior. An AI enforcement policy statement that outlined enforcement priorities and examples of acceptable risk mitigation practices would provide much-needed predictability. It would help innovators distinguish between legitimate experimentation and prohibited conduct, encourage early engagement with the agency, and promote the development of standardized risk-management practices aligned with FTC expectations. Such transparency would strengthen consumer protection while reducing the chilling effect that uncertainty currently imposes on responsible AI development. A well-defined AI enforcement policy would also be highly beneficial at the Food and Drug Administration (FDA)**,** particularly in the regulation of software as a medical device (SaMD) and AI-enabled clinical decision-support tools. While the FDA has issued draft guidance on the use of machine learning in medical software, enforcement expectations remain opaque, especially regarding when iterative updates to AI models trigger new regulatory submissions or enforcement actions. Innovators face uncertainty about how adaptive algorithms will be evaluated under existing premarket clearance pathways, and whether post-market model retraining could be construed as operating outside an approved authorization. An FDA AI enforcement framework could clarify these boundaries by specifying enforcement priorities, outlining conditions under which real-time learning and model updates are permissible. Such transparency would allow companies to innovate continuously within clear regulatory parameters and improve patient outcomes through faster model improvements and focusing resources on safety and quality rather than procedural compliance. Clear enforcement guidance would thus advance regulatory oversight and medical innovation by aligning incentives around measurable performance and public health impact.

In parallel with clearer enforcement communication, agencies should implement safe harbor programs that encourage proactive compliance and responsible innovation. Safe harbors provide regulated entities with defined protections or reduced enforcement exposure when they operate transparently within approved experimental or supervisory frameworks. For AI, this could include participation in regulatory sandboxes, voluntary reporting programs, or structured pilot initiatives overseen by agencies such as the FDA, DOT, or FTC. By allowing companies to test new models, deployment strategies, or governance tools under real-world conditions and with agency oversight, safe harbors promote early identification of risks while preserving the flexibility needed for innovation. OSTP should work with OMB and OIRA to develop model safe harbor provisions that agencies can adapt to their respective missions — linking participation to demonstrable commitments to safety, fairness, and transparency. Well-designed safe harbors strike a critical balance: they maintain accountability while providing innovators the confidence to experiment, learn, and deploy beneficial AI technologies without fear of inadvertent regulatory penalty.

Champion Regulatory Sandboxes

TechNet supports the establishment of a federal AI regulatory sandbox" framework, including an OSTP-led program such as that authorized in U.S. Senator Ted Cruz's *SANDBOX Act*. Unlike permanent regulatory exemptions, sandboxes provide a controlled space for experimentation while maintaining public accountability and data

transparency. These sandboxes would provide a safe harbor for innovators to test new AI applications under regulatory supervision, allowing for rapid learning and the co-development of smart, effective rules.

To complement this, OSTP should direct agencies to fully deploy and establish streamlined, transparent, and time-bound processes for granting waivers for AI-related pilot programs. Agencies should be empowered to issue time-bound waivers, collect outcome data, and publicly report on best practices and regulatory insights generated through sandbox participation. Additionally, these pilot programs should be coordinated across agencies to prevent fragmentation and allow for the cross-jurisdictional testing of technologies that operate across multiple regulatory domains such as transportation, healthcare, and communications.

To maximize their value, regulatory sandboxes should not operate in isolation but as structured learning mechanisms that feed directly into long-term regulatory modernization. Agencies should be required to evaluate sandbox outcomes systematically, assessing performance metrics, safety data, and compliance strategies, and use those findings to update existing rules, guidance documents, and standards. OSTP, OMB, and OIRA can coordinate this process by establishing a centralized repository where agencies share anonymized results, case studies, and evidence on what regulatory approaches enable innovation without compromising safety or accountability. This evidence-driven feedback loop would transform sandboxes from one-off experiments into engines of continuous regulatory improvement, helping policymakers identify where prescriptive rules can be replaced by performance-based standards or where outdated provisions can be safely retired. By institutionalizing this learning process, the federal government can ensure that the insights gained through experimentation directly inform more adaptive, risk-based, and innovation-friendly regulation over time.

Invest in Federal Capacity

Many federal agencies lack sufficient in-house AI expertise. This skills gap makes it difficult for regulators to assess complex AI systems, distinguish genuine risks from hype, and develop agile, forward-looking rules. Regulators who do not understand the technology are more likely to default to prohibition or inaction, both of which stifle innovation.

Congress and the Administration should authorize and fund a major initiative to recruit and train AI experts for careers in public service. This could include creating something akin to a "U.S. Digital Service" for AI, establishing competitive pay scales and fellowships for technical talent and creating clear career paths for AI specialists within the civil service. Establishing interagency AI training programs that support interagency AI coordination would also help build regulatory literacy and coherence. Additionally, the government should create programs that facilitate the temporary exchange of AI talent between the private and public sectors through secondments or fellowships that allow private-sector experts to serve for six to twelve months in federal agencies without sacrificing their career progression. TechNet has also been a longtime supporter of the creation of a National Digital Reserve Corps. A National

Digital Reserve Corps aims to bridge federal government needs and private sector capabilities by establishing a federal program to manage a reserve of individuals with the credentials to address the digital and cybersecurity needs of Executive Agencies across the federal enterprise. The creation of a National Digital Reserve Corps and expanded AI talent exchange programs could be closely coordinated with the Presidential AI Challenge and the AI workforce recommendations in the President's AI Action Plan to build a unified national strategy for developing and deploying AI expertise. For example, the Presidential AI Challenge could serve as the competitive entry point or pilot framework for selecting fellows who then transition into the Reserve Corps for continued federal service. By integrating these initiatives, the government can create a scalable model for bringing private-sector AI talent into public service, accelerating the adoption of safe and effective AI solutions across agencies while ensuring that workforce development, training, and deployment remain consistent with national AI strategy objectives.

TechNet believes this kind of creative thinking and public-private partnership can buttress the U.S. Government's workforce needs and address our ongoing modernizing efforts. Allowing individuals with technological expertise from industry to serve short-term positions in government, and for civil servants to spend time in industry, would build critical cross-sector understanding and capacity. By investing in its own people and building an internal pipeline of AI talent, the federal government can become a more effective partner for the innovation economy and a more capable steward of the public trust in the age of AI.

Sector-Specific Examples of Regulatory Barriers

Healthcare (FDA and HHS)

AI algorithms can now analyze medical images (e.g., MRIs, CT scans) to detect diseases like cancer with a level of accuracy that meets or exceeds human radiologists. These tools can reduce diagnostic errors, shorten wait times for patients, and allow clinicians to focus on treatment and care. However, healthcare regulations, such as certain provisions within the *HIPAA Privacy Rule* (45 C.F.R. Part 160), were written with human actors in mind. For example, documentation and audit trail requirements often presume a specific clinician is accessing a record at a specific time. This can be difficult to map onto a federated learning system where an algorithm is trained across decentralized data sets without moving the data itself, creating compliance uncertainty for hospitals wishing to adopt this cutting-edge technology. These rules should also be updated to ensure secure and appropriate access by AI processes to health data. The Department of Health and Human Services (HHS) should issue guidance or engage in rulemaking to clarify how privacy-preserving AI techniques like federated learning comply with HIPAA. The rules should be updated to focus on the security of the data and the outcome of the process, not the specific architecture of the system.

Additionally, the Food and Drug Administration (FDA) regulates artificial intelligence tools under its *Software as a Medical Device (SaMD)* framework (anchored in 21 CFR §

820), which covers stand-alone software intended to diagnose, cure, mitigate, or treat a disease. While the *21st Century Cures Act* does exempt a class of software functions — including many administrative functions — from being treated as medical devices under FDA law, the exemption is not absolute, and the FDA retains authority and may use enforcement discretion where safety concerns are involved. As a result, the absence of an approved pathway for adaptive learning models has delayed the adoption of AI-based diagnostic and predictive tools. Instead, TechNet recommends that the FDA issue guidance to explicitly state that only clinical AI that directly informs diagnosis and treatment decisions are under FDA's regulatory authority for SaMD purposes. Distinguishing between clinical AI with direct implications for diagnosis or treatment and administrative AI used for workflow efficiency is critical to allow AI to augment the work of licensed providers. The FDA should recognize that licensed providers are ultimately responsible for clinical treatment and care decisions for patients, even when those decisions involve the use of AI. This will ensure that the standard of care does not change, while still allowing for healthcare technologies to innovate and improve.

Similarly, the Office of the National Coordinator for Health IT (ONC) has implemented the *HTI-1 Final Rule* (effective February 2024), which requires certified electronic health record systems to provide transparency into the operation and limitations of decision support interventions. While important for patient safety, these requirements could be interpreted too rigidly for adaptive or generative AI tools. TechNet encourages OSTP to recommend a principles-based transparency approach, with safe harbors for companies that conduct and document bias testing, including internal use of data for training models, and real-world performance evaluations, rather than prescriptive disclosures that risk exposing proprietary intellectual property.

Transportation (DOT and FMCSA)

Autonomous trucking promises to revolutionize logistics by making supply chains more efficient, reducing fuel consumption, and improving safety on our nation's highways. The National Highway and Traffic Safety Administration (NHTSA) has taken welcome first steps to modernize Federal Motor Vehicle Safety Standards (FMVSS) to account for autonomous vehicles. NHTSA should build on its continued work in this space and further clarify that manually operated controls and equipment intended only to support a human driver are not necessary for Society of Automotive Engineers (SAE) level 4 and level 5 AVs. Removing these outdated requirements will support U.S. innovation and leadership on AVs, enhance safety, and encourage the safe deployment of AVs.

On top of this, existing Federal Motor Carrier Safety Administration (FMCSA) regulations still require a level of involvement by human drivers, precluding full automation. For example, under current regulations, if a commercial motor vehicle (CMV) is stopped on the highway or shoulder for any reason other than a necessary traffic stop, then warning devices (e.g., warning triangles) must be placed within 10 minutes in three locations along the roadway. Given the absence of a human driver in an autonomous CMV, transportation stakeholders and safety advocates have

supported an industry developed solution to support a new safety solution that meets the needs of autonomous trucks while ensuring the safety of all road users.

In 2023, Aurora and Waymo, both leading American AV technology companies, filed an exemption petition with the Federal Motor Carrier Safety Administration (FMCSA) that would allow for the use of cab-mounted beacons that have been shown to achieve a level of safety that is equivalent to, or greater than, the level of safety by the current requirements. If broadly adopted by the trucking industry, the cab-mounted beacon solution would also benefit conventional CMVs by providing added protection to human drivers and other road users by preventing the need for a driver to step out of the truck and walk alongside the roadway to place warning triangles as required under current regulations. Recently, motor carriers operating with a Level 4 Automated Driving System were granted a waiver to use the warning beacons for a limited period under certain conditions and requirements. This is an important step toward a long-term solution. Modernizing the regulations to allow for greater adoption of this system would provide regulatory certainty for AV trucks and provide additional safety measures for drivers and other motorists.

Another example of FMCSA regulation hindering AI adoption is FMCSA's "Hours of Service" regulations (49 C.F.R. Part 395) that are designed to prevent accidents caused by fatigued human drivers. These rules mandate specific off-duty and sleeper berth periods. These kind of regulations and requirements are structurally nonsensical when applied to a fully autonomous vehicle that does not experience fatigue. While waivers are available, the process is slow and the lack of a clear, updated regulatory framework for autonomous operations creates significant uncertainty for a multi-trillion-dollar industry, delaying large-scale investment and deployment. There is a clear need to transition FMCSA and related transportation regulations toward outcome-based metrics — collision avoidance, occupant protection, and system reliability — and expand exemption authority to allow for greater automation.

Financial Services (Federal Reserve, CFPB, SEC)

Rules such as ECOA/Reg B (12 CFR § 1002) and SEC Rule 15c3-5 rely on deterministic model explainability and manual risk checks, which are incompatible with machine learning. This constrains fair-lending innovation and next-gen compliance tools. Joint interpretive guidance that clarifies acceptable algorithmic explainability standards should be adopted, alongside supervised AI testing environments under existing prudential oversight.

Procurement, Accessibility, and Workforce Policy (OMB, GSA)

Government procurement processes must be modernized to ensure easier and more effective AI adoption. For example, the FedRAMP certification process presents a significant barrier when AI features are added to existing platforms. In many cases, it triggers a full recertification, which is time-consuming and resource intensive. This procedural rigidity discourages iterative innovation and slows the deployment of AI-enhanced solutions. Ongoing efforts to modernize the FedRAMP certification process,

including through FedRAMP 20X and AI prioritization effort, have shown an improvement over the previous documentation-focused effort and should be expanded. In particular, the FedRAMP 20X Key Security Indicator (KSI) process aligns with greater automation, rather than static, point-in-time audits that are outdated as soon as they are complete, and significantly reduces the time to get authorized, ensuring government has access to cutting edge technologies. The AI Prioritization path also provides direct, hands-on support and real-time feedback during the audit, which allow for a faster, more successful review upon audit completion. These advancements have significantly improved the FedRAMP program and TechNet recommends continuing and expanding these enhancements.

On top of this, federal acquisition rules (FAR Part 39) require fixed technical specifications, incompatible with iterative AI procurement. OMB's 2023 *M-23-18* guidance broadly restricting AI tool use has deterred experimentation across agencies. Instead, OMB should create standardized *Responsible AI Use Authorities* that allow employees to leverage commercial AI tools under approved data-handling protocols, coupled with agency "AI Centers of Excellence."

Finally, requirements outlined in Section 508 of the *Rehabilitation Act of 1973* related to digital accessibility in the federal government must be modernized to reflect technological advances. Section 508 rightly ensures that federal technologies are accessible to all users, but its rigid application to rapidly evolving AI interfaces can delay pilots and discourage iterative improvement. Because AI interfaces evolve rapidly, requiring full accessibility recertification for every minor update or model iteration imposes significant operational burdens without meaningfully improving accessibility outcomes. A more flexible, tiered compliance framework would better balance these objectives by streamlining requirements for low-risk or pilot deployments and focusing resources on developing durable, long-term accessibility solutions in collaboration with stakeholders. This approach preserves the intent of Section 508 — universal access — while enabling innovation and continuous improvement in AI-enabled tools.

Spectrum Policy and Next-Generation Networks (NTIA, FCC)

AI's success depends on leadership in spectrum and network modernization, and thus policy coordination across these digital infrastructure domains will be critical. However, fragmented rulemaking across these areas currently amplifies compliance friction and slows progress. NTIA and FCC need to align their AI regulatory reform efforts with a spectrum pipeline for AI and next-gen network rollout, recognizing connectivity as foundational to AI competitiveness. AI workloads will exponentially increase network demand, especially for uplink traffic as edge devices generate and transmit data for real-time inference. Without a robust pipeline of flexible-use spectrum, AI applications that rely on low-latency connectivity — autonomous vehicles, telehealth, precision agriculture, and industrial automation — will be throttled by capacity constraints. Regulatory inertia in identifying and auctioning mid-band spectrum has already delayed additional 5G spectrum resources and could now delay AI-driven network intelligence. The lack of clear timelines for releasing

additional spectrum bands undermines investment and slows deployment of AI-optimized networks.

OSTP should work with NTIA and the FCC to treat spectrum availability as an AI-enabling infrastructure priority and establish a spectrum pipeline for AI to ensure predictable release of flexible-use spectrum, coupled with streamlined permitting for network upgrades.

Additionally, next-generation networks — 5G Advanced, 6G, and AI-native architectures — are not only conduits for AI applications but also platforms that use AI to manage themselves. However, legacy regulatory frameworks delay the transition to software-defined and virtualized infrastructure by maintaining outdated certification and equipment authorization processes. Network regulations need to be modernized to support dynamic spectrum sharing, open-RAN deployment, and AI-enabled network management.

Critical Infrastructure (DHS, DOE)

Critical infrastructure sectors — energy, communications, and transportation — depend increasingly on AI for predictive maintenance, grid optimization, and autonomous incident response. Yet some statutory frameworks treat AI-driven control systems as potential vulnerabilities rather than protective assets. For example, DHS and DOE cybersecurity standards often restrict autonomous control actions absent human authorization. Such restrictions may delay response to fast-moving network or grid events and prevent deployment of AI agents that could contain or remediate outages autonomously. Instead, critical-infrastructure protection laws and guidance should be revised to explicitly permit AI-based network optimization, including to allow AI for autonomous monitoring, fault detection, and network optimization, while maintaining safety and cybersecurity standards.

Additionally, regulators should be encouraged to speed up interconnection approvals for energy infrastructure and AI data centers. Lengthy and complex permitting and interconnection processes delay the deployment of critical clean energy projects and the expansion of data center capacity needed to support AI innovation and economic growth. Streamlining these approvals through clearer timelines, standardized application procedures, and coordinated review across agencies would enable faster integration of new renewable generation and transmission assets, while also ensuring that AI data centers can access reliable, sustainable power. By modernizing interconnection policies, regulators can help align infrastructure development with national energy and innovation goals, supporting both decarbonization and technological leadership.

Export Controls (Commerce/BIS)

The Bureau of Industry and Security's (BIS) Interim Final Rule (15 CFR § 744) defining "frontier model" exports represents an important national security safeguard, but its current structure is overly broad and risks capturing legitimate, low-risk

research activity. The rule's inclusion of open-source models, academic collaborations, and commercial partnerships that do not involve sensitive end-uses or foreign adversaries introduces unnecessary friction into the U.S. innovation ecosystem. By relying on expansive definitions of "frontier model" based primarily on model size, parameter count, or training compute, the rule inadvertently encompasses a wide range of general-purpose or precompetitive research projects that pose minimal security concerns. This uncertainty discourages collaboration between U.S. researchers and trusted foreign partners, complicates data-sharing agreements, and can drive cutting-edge research offshore to countries with clearer, more proportionate regimes.

To maintain both national security and U.S. leadership in AI research, export controls must be precise, risk-based, and administratively predictable. BIS should refine the rule to focus on factors that directly correlate to misuse risk — such as compute capacity, model capability thresholds, and end-use or end-user intent — rather than broad technical characteristics like size alone. Clear definitions of what constitutes a "controlled frontier model," coupled with transparent criteria for exemptions and license exceptions, would provide innovators with certainty while preserving flexibility to address genuine threats. OSTP should work closely with the Department of Commerce to issue joint guidance distinguishing legitimate AI research and open collaboration from activities that raise national security concerns, ensuring that routine research, benchmarking, and model evaluation are not inadvertently restricted.

Finally, Commerce should promote international coordination on AI export control norms to prevent regulatory divergence that disadvantages U.S. innovators. Working through forums such as the OECD, the Global Partnership on AI (GPAI), and the U.S.–EU Trade and Technology Council (TTC), the United States can advocate for shared, evidence-based standards for AI model classification, compute risk thresholds, and end-use monitoring. Such collaboration would reinforce national security goals while preserving the openness and interoperability essential to the global AI research community. By modernizing export controls in this balanced, risk-based manner, the federal government can ensure that U.S. policy continues to protect security interests without undermining the innovation capacity that drives America's AI leadership.

Copyright and Fair Use (U.S. Copyright Office, USPTO, USTR)

Like other transformative technologies, the rapid advancement of AI is facing a wave of litigation under copyright law that could substantially affect the development and deployment of generative models, content moderation tools, and automated data analysis systems. Inconsistent interpretations and litigation risks over how traditional concepts like fair use, authorship, and derivative works apply to modern AI training and outputs threaten to discourage investment in both AI innovation and creative industries.

Federal agencies should reaffirm that existing copyright principles — particularly the fair use doctrine — remain technology-neutral and continue to apply when AI systems

use data for transformative purposes such as model training, research, or model evaluation. Protecting the right and ability to perform computational analysis on data is necessary to create effective AI models and ensure the U.S. remains the leader in AI development. Any changes to existing copyright law or efforts to narrowly tailor fair use exemptions that would restrict access to AI training data will hamstring America's ability to innovate and compete globally. Rather than implementing new statutory requirements or restrictions, the U.S. Copyright Office should ensure that any further guidance focus on protecting the core provisions of copyright law and fair use doctrine that offer technology-neutral protections to legitimate rightsholders and innovators. This will encourage a balance where creators' rights remain honored while ensuring continued progress in developing groundbreaking AI technologies. Federal agencies should advocate for copyright policies, both in the U.S. and abroad, that support and reinforce our national leadership in innovative AI technologies.

Preempt State Regulation

This year alone, over 1,000 AI bills were introduced in state legislatures. These bills are not uniform, contain different definitions of AI and related terminology, and require different disclosures for engineering content. This developing patchwork makes compliance burdensome for businesses and confusing for consumers, and it serves as a significant barrier to America's AI leadership. To this end, we believe it is important for the administration to develop federal regulations and responsible safety practices and to harmonize national standards around AI testing and evaluations. A unified federal approach to AI regulation would help address compliance burdens with varying state regulations while still making room for states to address concerns related to high-risk consumer-facing applications where clear gaps have been identified and no existing regulation is applicable.

For example, states are adopting conflicting definitions of "automated decision systems," "algorithmic accountability," and "high-risk AI," forcing companies to design different compliance programs for each jurisdiction even when operating a single national product. This duplicative system increases costs, deters small innovators from entering the market, and slows the deployment of beneficial technologies. In some cases, state rules impose obligations that contradict or exceed federal frameworks, particularly in privacy and bias auditing, creating uncertainty about which standard prevails.

In addition, any federal efforts to regulate AI innovation should recognize the importance of publicly available data in providing U.S. companies with a competitive advantage. Any state regulation that would diminish access to and use of publicly available data should be preempted by federal regulation. As set forth above, publicly available data is necessary to create effective AI models and ensure the U.S. remains the leader in AI development. Any efforts to restrict access to publicly available data for AI training purposes will significantly undermine development efforts, and differing state treatment will create substantial compliance risk and burden particularly for small and medium-sized businesses.

Given the rate at which we are seeing overregulation at the state level, the federal government should look to impose a moratorium on state legislation related specifically to the development of frontier AI models until a federal regulatory framework and national standards are adopted.

Combat International Overregulation

International regulation is increasingly creating fragmented and burdensome compliance regimes that stifle American AI innovation and disadvantage U.S. companies competing globally. Divergent requirements across jurisdictions — ranging from the *EU AI Act*'s rigid classification system to overlapping data, privacy, and algorithmic transparency mandates — impose duplicative costs and uncertainty on developers seeking to deploy AI products across borders. These inconsistencies discourage cross-border research collaboration, limit market access for startups, and divert resources away from responsible innovation toward legal compliance. Without greater regulatory alignment and interoperability, the global AI landscape risks becoming balkanized, slowing the pace of technological progress and undermining the United States' ability to lead in developing trustworthy, human-centered AI systems.

The United States must drive global consensus in support of a U.S.-led framework for international AI standards and definitions that enables regulatory coherence and global adoption. This includes working closely with trusted partners and allies to harmonize AI standards and regulations to ensure that misaligned regulatory frameworks do not create unnecessary barriers to AI adoption, increase compliance costs, or slow innovation. The administration should empower the Center for AI Standards and Innovation (CAISI) to lead this effort and guard against burdensome international regulations by coordinating with the NIST Information Technology Laboratory (ITL), the State Department, and other agencies to ensure consistent U.S. representation in key international standards bodies, while issuing clear guidance to U.S. companies on how to engage effectively. OSTP should also work with the Department of Commerce and the United States Trade Representative (USTR) to leverage trade negotiations in collaboration with allied nations to ensure aligned AI regulatory practices globally. This international engagement strategy should protect U.S. market access and promote an innovation-oriented approach, including advocating for adherence to international consensus-based technical standards, the use of existing regulatory frameworks where possible, and AI-specific rules only where gaps exist.

Additionally, as the world demands more and more technology, the administration should move assertively to accelerate its efforts to export American AI and technology solutions, leveraging the directives laid out in Executive Order 14320, *Promoting the Export of the American AI Technology Stack*, and the newly launched American AI Exports Program at the Commerce Department to amplify American competitiveness on the world stage. By accelerating efforts aimed at reducing barriers to export — including streamlining approval, aligning export controls, offering diplomatic and financial tools, and promoting U.S. standards abroad — the administration can better support

innovation and the development and deployment of American AI across the globe.

Conclusion

The United States stands at a pivotal moment. The decisions we make today regarding the governance of AI will determine our economic competitiveness, national security, and global leadership for the remainder of the 21st century. To seize the immense opportunities presented by AI, we must urgently address the outdated and misaligned regulatory frameworks that are currently acting as a brake on innovation.

Regulations should evolve with technological advancement. OSTP's leadership in this regulatory reform effort is essential to harmonize agency actions, promote experimental flexibility, and integrate AI policy with broader infrastructure initiatives. TechNet believes that a focused and collaborative effort between industry and government can create a modern, pro-innovation regulatory ecosystem that will ensure America wins the global AI race. We remain eager to collaborate with the administration in developing balanced AI policies that safeguard public interests while ensuring the United States maintains its global leadership and continues to foster AI innovation.

Sincerely,

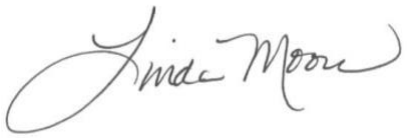

Linda Moore
President and CEO