



2026 National AI Policy And Regulatory Framework

CTA supports one nationwide rulebook for AI that will end the chaos of conflicting state laws. With one set of rules, American innovators can focus on what they do best: inventing the future and advancing American AI leadership. This document sets out the elements of a national policy framework that ensures AI systems are safe, trustworthy, and transparent and that promotes industry-led voluntary standards.

I. PRINCIPLES

Leveraging Existing Law

Lawmakers should leverage existing laws and standards that already govern AI deployment, development, and use. For example, existing laws already guard against illegal bias and discrimination, whether or not AI is involved. Any federal law relating to AI should recognize where such laws provide existing remedies and procedures and should avoid duplication wherever possible.

Technology Neutrality

Lawmakers should avoid distorting markets and stifling innovation by favoring specific AI technologies or techniques. Federal AI law should, to the greatest extent possible, focus on principles, outcomes, and risk rather than the specific technologies employed in AI systems. Technology neutrality allows law to evolve alongside technology as AI systems become more complex and independent.

Risk-Based Approach

Many AI use cases have safely integrated into Americans' daily lives and demonstrated very low risk. Given the wide range of possible AI systems and use cases, lawmakers should tailor regulatory frameworks to risk. This includes scoping regulatory frameworks to AI systems that pose risks and focusing regulatory obligations on reasonably managing those risks.

Use of Existing Standards and Frameworks

Federal law should, where appropriate, reference and acknowledge widely recognized standards and frameworks—particularly the National Institute of Standards and Technology AI Risk Management Framework (NIST AI RMF), which was developed through a highly collaborative process with feedback from industry and other stakeholders.

II. NATIONAL POLICY FRAMEWORK

Preemption

Federal law should preempt state and local regulation of AI, except for: (1) states' traditional police powers (such as generally applicable laws related to prohibiting child sexual abuse material, protecting children, and preventing fraud); and (2) state and local laws regulating state and local government use of AI.

Transparency

When a conversational AI system interacts directly with consumers (such as a chatbot), federal law should require clear and conspicuous disclosure that the system is AI, if a reasonable person interacting with the AI system would believe they were interacting with a human. Developers and deployers should work together to enable that disclosure, with each responsible for the elements under their control.

Federal law should require frontier developers to publish, and regularly update, a frontier AI framework that describes the technical and organizational measures the developer uses to manage, assess, and mitigate catastrophic risks related to their frontier AI models. Federal law should also require frontier developers to file corporate disclosure statements with the Federal Trade Commission (FTC) containing basic business name and address information, ownership information, and points of contact.

In addition, all frontier developers should be required to publish a transparency report when deploying a new or substantially modified frontier model. The transparency report should include basic information about how the model can be used, such as intended uses, languages supported, types of outputs, and restrictions or conditions on use. For frontier developers, the transparency report should also summarize the developer's assessment of catastrophic risk and other steps undertaken pursuant to its frontier AI framework.

The frontier AI framework and transparency report requirements should allow frontier developers to redact information as needed to protect public safety and national security, to comply with applicable law, and to protect developers' proprietary information, trade secrets, and cybersecurity.

Reporting

Federal law should require frontier developers and deployers to report critical safety incidents to the Center for AI Standards and Innovation (CAISI) within 15 days after discovery and verification of a reportable incident. The report should be limited to the date of the incident, a short description of the incident, and an explanation of why the incident qualifies as a "critical safety incident." The reporting obligation should fall on the entity that discovers the incident.

Frontier model developers should further be required to regularly submit to CAISI summaries of assessments of catastrophic risk resulting from internal use of frontier models. Reporting should be exempt from public disclosure under the Freedom of Information Act and other applicable laws.

Whistleblower Protections

Lawmakers should prohibit employers from retaliating against employees for, or prohibiting employees from, reporting credible concerns about catastrophic risk, critical safety incidents, or violations of law to employers, regulators, or other appropriate government entities.

Regulatory Sandboxes

Federal law should establish and encourage the use of regulatory sandboxes. Regulatory sandboxes allow regulators, industry, and civil society to collaboratively develop effective frameworks for promoting safe, trustworthy, and transparent development and deployment of AI.

Standards Development

Lawmakers should encourage timely industry-led, consensus-based voluntary standards development, building off the strong foundation of the NIST AI RMF and responding to the evolving nature of the technology. Federal law should formally establish CAISI and codify its role.

Support for Small Businesses

To encourage the adoption of AI, lawmakers should direct NIST to develop or identify AI resources specifically for small businesses. These resources should be technology neutral and based on industry-led, consensus-based voluntary standards, where applicable.

Intellectual Property

Lawmakers should not restrict AI developers from relying on fair use when training their models and should encourage voluntary practices that support a rights-respecting ecosystem. Courts have long recognized that fair use furthers the ultimate aim of copyright—promoting innovation—and they have extensive experience applying the fair use framework to cutting-edge technologies. Legislation on AI-related fair use questions would be premature.

Demonstrating Compliance Through Industry Solutions

Federal law should incentivize entities to self-regulate and demonstrate compliance by adopting accepted industry solutions such as third-party certification, self-certification of compliance, alignment with the NIST AI RMF, or alignment with accredited national or international standards for AI governance.

Sector-Specific Regulations

To avoid an overly complex and conflicting regulatory framework, lawmakers should recognize that some AI uses are already subject to an established, sector-specific safety regime administered by a federal regulator with scientific or technical subject-matter expertise in the relevant field, such as the approval of drugs and medical devices and the safety regulation of aviation and motor vehicles. When more general AI rules and these sector-specific regulations conflict, the more specific rule should control.

To encourage consistent rules across industries and reduce customer confusion, sector-specific regulators should avoid, to the greatest extent possible, regulating AI use in general business activities of regulated entities (for example, hiring and accounting) that are not specialized to the sector.

Oversight and Enforcement

Lawmakers should not create a new rulemaking body to regulate AI. While subject matter experts throughout the federal enterprise should continue to develop standards and to assess and manage AI risks consistent with their existing authorities, the FTC should be exclusively responsible for enforcing compliance with the transparency and reporting requirements, in addition to applying its existing unfairness and deception authorities in the context of AI.

To the extent feasible, the FTC should facilitate recognition of comparable disclosures, certifications, and attestations across jurisdictions. Organizations operating across jurisdictions should not have to repeat substantially equivalent assessments solely because different regulators use different terminology or documentation formats.

Lawmakers should also exempt, as appropriate, the use of AI for military, defense, or national security purposes and for scientific research and development.

III. DEFINITIONS

Definitions must adapt to rapidly changing technology. In addition to adopting the definitions below, lawmakers should create procedures for CAISI, in consultation with NIST and relevant stakeholders, to regularly review and update the definitions' technical, capability, and impact thresholds.

AI System

Federal law should employ the definition of "AI system" used in the NIST AI RMF: "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" and that is "designed to operate with varying levels of autonomy." Federal law should further clarify that neutral security, authentication, and attestation infrastructure—which does not itself generate the outputs described above—is not an "AI system" and that supplying such infrastructure does not make a person a developer or deployer.

Catastrophic Risk

A foreseeable and material risk that a frontier developer's development, storage, use, or deployment of a frontier AI model will materially contribute to the death or serious injury of more than 50 people or more than \$1 billion in damages resulting from a frontier model engaging in: (1) the creation or release of chemical, biological, radiological, or nuclear weapons; (2) cyberattacks without human oversight; (3) certain crimes; or (4) evasion of the developer or end user's control.

Critical Safety Incident

Any incident involving: (1) unauthorized access to, modification of, or theft of model weights of a frontier model that materially increases catastrophic risk; (2) harm resulting from a materialization of a catastrophic risk; (3) a loss of control of a frontier model that materially increases catastrophic risk; or (4) circumvention of safeguards, monitoring, or human oversight that materially increases catastrophic risk.

Deployer

A person or entity that uses an AI system to interact with end users or uses the outputs of an AI system to make decisions affecting end users.

Developer

A person or entity that codes, designs, produces, or substantially modifies an AI system for use by members of the public. Developers are sometimes also deployers of their own systems.

Foundation Model

An AI model that is trained on a broad data set, designed for generality of output, and adaptable to a wide range of distinctive tasks.

Frontier Developer

A person or entity who trains, or initiates the training of, a frontier model using computing power that meets the technical specifications set forth in the definition of frontier model.

Frontier Model

Current definitions of frontier model generally rely on the amount of compute used for training, defining a frontier model as a foundation model trained using more than 10^{26} integer or floating-point operations. Although this may be appropriate now, federal lawmakers should consider more future-proof definitions.