PALADIN

GLOBAL INSTITUTE

APPLYING THE DISCIPLINE OF LAW TO FRONTIER AI: PRINCIPLES, AUTHORITIES, AND OPTIONS



Acknowledgements

Many thanks to our contributors including:

- Leonard Bailey
- Steve Bunnell, Raj Ferber PLLC
- Adam Isles, The Chertoff Group
- Robert Kang, Adjunct Faculty, Loyola Law School, Los Angeles & USC Viterbi School of Engineering
- Devin Lynch, Senior Director, Paladin Global Institute
- Katie Moussouris, Founder and CEO, Luta Security
- Chris Painter, The Cyber Policy Group and Adjunct Professor Johns Hopkins University SAIS, Bologna
- Elliott Phaup
- Kiran Raj, Raj Ferber PLLC
- Haley Ring, Director, Paladin Global Institute
- Matthew Shabat, Chief Risk Officer, NAIS LLC and Adjunct Faculty, Georgetown University
- Daniel Sutherland, Senior Advisor, The Chertoff Group and Adjunct Professor, George Washington University
- Kemba Walden, President, Paladin Global Institute

About The Paladin Global Institute

The Paladin Global Institute (PGI) serves as an alliance building platform connecting civil society, global leaders, innovators, investors, and industry to share knowledge and insights, develop a shared comprehension of challenges and implement lasting solutions, foster and sustain partnerships and alliances, and provide stability amid an ever-changing landscape.

Under the leadership of Paladin's Kemba Walden, Devin Lynch, and Haley Ring, this paper was developed as part of PGI's Paladin Connects program. The mission of Paladin Connects is to engage thought leaders, industry, government, civil society, and relevant stakeholders in policy conversations that advance the security and safety of critical infrastructure.

Applying The Discipline Of Law To Frontier AI: Principles, Authorities, and Options

The emergence of artificial intelligence, and the safety and security risks it presents, has become one of the defining policy challenges of our time. The results of the November 3 midterm elections highlighted even more clearly the growing public concern about the societal impacts of AI. Despite intense public debate and growing concern among tech CEOs, AI professionals, and political leaders, the United States has yet to develop a coherent framework for governing the development and deployment of frontier AI systems. One reason for this lack of progress is that much of the current debate has not been sufficiently informed by legal analysis and the accumulated experience of legal institutions. Discussions of AI safety and security often overlook:

- how legal systems function in practice,
- the lessons derived from decades of experience designing, implementing, and evaluating regulatory regimes,
- the wide range of tools available to lawmakers and regulators, and
- the body of law that already bears on AI development and deployment.

This paper builds on discussions held at the Paladin Global Institute through its Paladin Connects convening series. It is organized into two parts. Part I examines how law can inform the governance of frontier AI. It identifies several foundational principles about law and regulation that can help guide the development of effective AI policy. It also reviews the current legal landscape. Part II turns to policy options, evaluating a range of legal, regulatory, and standards-based approaches for governing the development and deployment of frontier AI systems.

Introduction

Since making artificial intelligence (AI)[1] accessible, several of the companies building frontier AI consistently maintain that they cannot reliably control or fully explain the systems they are building, testing, deploying, and selling to the public and to the U.S. government every day. They published research announcements, essays, manifestos, and unprecedented cybersecurity incident reports with these concerning revelations, and generally called for regulation.

In May 2023, leaders of those companies affixed their names to a simple statement: "Mitigating the risk of extinction from AI should be a global priority alongside other societal-scale risks such as pandemics and nuclear war."[2] That same July, OpenAI acknowledged it did not have a solution "for steering or controlling a potentially superintelligent AI."[3] In 2025, Anthropic CEO Dario Amodei wrote that the field does not understand how its own AI creations work.[4]

Over the past year, these companies' warnings have grown in urgency, specificity, and consequence—and so have the cybersecurity incidents behind them. In November 2025, Anthropic reported that a group it assessed to be Chinese state-sponsored had manipulated its Claude Code tool into carrying out an estimated 80 to 90 percent of an espionage campaign against roughly thirty organizations.[5] In February 2026, Anthropic and Google each disclosed industrial-scale campaigns to extract their models' capabilities: more than 16 million exchanges across some 24,000 fraudulent accounts at Anthropic, and a single campaign of more than 100,000 prompts against Google's Gemini.[6] In July 2026, OpenAI models under cybersecurity evaluation got around the controls meant to isolate them from the internet and compromised Hugging Face's systems.[7] On July 28, after OpenAI disclosed its role in the Hugging Face incident, employees of frontier AI companies asked the U.S. government to support an international effort to build the "technical and governance tools" needed to pace frontier development.[8]

On September 12, 2026, Amodei wrote that "we must slow the pace at which we improve the capabilities of AI models."[9] Within hours, OpenAI CEO Sam Altman agreed and pledged that OpenAI would also give independent evaluators employee-like access.[10] Eight days later, an OpenAI agent bypassed its network restrictions to query an outside chatbot, and the company paused training of its most capable models for the second time in under three months.[11]

By contrast, other leading tech executives and political leaders dispute the assertions that AI promises significant safety and security impacts. Moreover, they dispute that new laws and regulations are needed in this space. Instead, they argue that current laws and regulations provide strong consequences that should motivate companies to deploy only those models that have been sufficiently vetted for safety and security. They argue that if companies release models that create safety and security consequences, those companies can be held liable under current tort laws or can be investigated by federal and state regulators.

On September 29, 2026, leaders of six of the most capable AI technology companies in the world joined the President of the United States at the White House and signed a one-page pledge.[12] The accord asks every company that trains frontier models to stand up four layers of control: (1) internal controls that track model capabilities and alignment, (2) an internal team to check those controls, (3) an independent external auditor, and (4) an independent board committee to review the results. The Paladin Global Institute's June 2026 Governance Layer primer calls this work internal governance: the policies, processes, and practices that turn rules into organizational practice.[13]

The accord sets no enforcement mechanism nor requires public disclosure of audit results. The accord's operative verb is "should," not "must."

It is likely that this accord sets the stage for more regulation in the future and concedes as much: it says it "may make sense" to codify its steps in law over time.[14]

The November 3 midterm elections demonstrated growing public concern about the societal impacts of artificial intelligence and increased public support for stronger oversight of AI development and deployment. Exit polling found that 56% of voters considered AI and data centers a major factor in their congressional vote, while pre-election surveys found that approximately three-quarters of Americans believed government was doing too little to regulate AI. The success of candidates emphasizing AI accountability and the widespread approval of measures restricting data-center development suggest growing public support for more rigorous oversight of AI development and deployment.[15]

So, what could those future laws and regulations realistically look like? Tech industry leaders and professionals have identified their concerns—malevolent safety and security outcomes, all while the technology still needs freedom to innovate. It now falls to those with expertise in technology, law, and policy to (a) inform the discussions by laying out some core principles of law and regulation; and (b) identify specific options for moving forward that address these concerns. We know that when these disciplines—law, policy, and technology—collaborate, innovations are possible; when they operate in their own siloes, poor results are certain.

Part One: Principles Of Law To Inform The Debate

Several core principles about law and regulations should help shed light as the country considers possible policy on the development of artificial intelligence. Lawyers and lawmakers have tools in their toolbox, including the following:

A. The structure of law

Not all requirements are cut from the same cloth. Some requirements come from *within a company*—internal governance documents, policies, and standards that management establishes to guide employee behavior and operational practice. These carry only the authority of the employer, and they are the most flexible since the organization can revise them as circumstances change. Other requirements emerge from *industry-wide actions.* For example, industry standards bodies—technical communities and professional associations—establish frameworks and benchmarks that reflect the best collective practices. These are rarely legally binding. However, such standards are at times incorporated into contract provisions and thereby bind the parties. Such standards can also carry weight in regulatory investigations and civil litigation, where conformance to an industry standard can be the difference between a defensible and an indefensible position.

Statutes are requirements passed by the legislative bodies of states and the Congress—enacted through democratic processes and enforceable through the courts. *Regulations* are requirements issued by government agencies to implement the legislative requirements consistent with authorities granted by the Congress. Regulations are necessarily more technical and detailed than statutes, and often directly applicable to specific industries or activities.

Thus, in the context of potential new requirements guiding the development of artificial intelligence, we need to keep in mind the following categories:

- Actions that can be taken unilaterally by AI developers
- Actions that can be taken across industries
- Actions that can be taken by the Executive Branch
- Actions that can be taken by the U.S. Congress
- Actions that can be taken by state governments
- Actions that can be taken by international cooperation
- Actions that can be taken by other governments that affect U.S.-based AI companies.

B. Time

When a lawyer or lawmaker considers options for establishing requirements, they understand that each type of requirement brings with it different time horizons. An action that depends on passage of a law by the U.S. Congress will require at least a several month horizon;

in the specific case of AI, the time horizon might be even longer given the President's stated opposition to such legislation. Tasking a federal agency to promulgate a new rule also takes an extended period of time. An action that depends on an international treaty might require years of negotiation, whereas a less formal agreement between countries can be contained in a declaration or communique. Actions undertaken by individual companies, by contrast, could be implemented immediately.

C. Specificity

Requirements can be written in a very detailed, specific and proscriptive manner. Or, they can be broad and process-oriented, permitting each company to apply the rules in a way that makes most sense for their context. Neither approach is inherently right or wrong. There are circumstances where industry pleads for specific rules, which allow them to operate in an environment with complete clarity. In other circumstances, detailed rules can strangle and over-regulate.

For example, the Americans with Disabilities Act required the Executive Branch to publish rules that include design specifications for new buildings. While some decried how prescriptive the regulations were, architects were pleased—they now knew exactly how many inches off the ground toilets need to be, where grab bars should be placed, etc. Specificity provided needed clarity. By contrast, in the cybersecurity field, it is widely understood that Congress should not pass laws that require specific technology practices. If Congress passes laws that, for example, require passwords to be of a certain complexity, before the ink is dry that requirement will be obsolete. Instead, in the cybersecurity area the rules are often process-based (e.g., companies need to perform risk assessments), and cybercrime laws are deliberately technology-neutral, so they need not be revised every time a new technical means of attack is developed by cybercriminals.

An innovative form of regulation was found in the Chemical Facility Anti-Terrorism Standards (CFATS), passed by Congress in the aftermath of 9/11 to try to mitigate the risk that chemicals would not be used by terrorists for mass destruction. Congress found that the production, storage and distribution of chemicals is actually quite complicated, and one-size-fits-all rules would not work. Therefore, Congress created a unique scheme, implemented by regulations developed by the Department of Homeland Security: broad standards were set, and then each company was required to develop its own security plan against those standards. DHS reviewed each company's plan and approved those that met the standards in a meaningful way. Going forward, the companies were judged against their own plan rather than against a set of proscriptive rules imposed on them by the government. This scheme offers regulators another tool in the toolbox—creating broad standards and allowing each company to customize rules to its own context.

D. Who should author regulations?

It is a weighty thing to create rules that bind the conduct of people or companies; we must be careful to give that responsibility to institutions that are well placed to create such rules. Requirements need to be designed by experts in the field, who understand that industry through experience and careful study. It is also critical that requirements be durable, recognizing that industries, and particularly the technology industry, need the flexibility to rapidly evolve.

In some circumstances, Congress can debate, discuss, and then enact laws that are effective in regulating behavior. But in most situations, Congress passes laws that

set broad standards and then empowers Executive Branch agencies to set the specific rules. Agencies have longer time to study; have larger staffs; can hire and train experts in each field; and, can engage with those in the field to ensure that rules are well-crafted. The more technical the field, the more likely it is that Congress will leave the creation of specific rules to the experts in the Executive Branch agencies.

Courts are sometimes compelled to create rules that bind industry action. If Congress or oversight agencies have not acted, courts are often put in the position of resolving disputes which then effectively set the rules of the road. Courts are asked to determine if a company's actions have been reasonable or negligent, and therefore the court either explicitly or implicitly establishes rules for future behavior. Many tort cases are resolved not by a verdict but by a settlement agreement which the parties negotiate and the court approves. Those settlement agreements often contain terms that essentially establish the rules for others. Courts try to avoid these situations. As the Supreme Court has recognized, agencies often act "within [their] area of special expertise, at the frontiers of science," while judges "are not experts in the field." *Baltimore Gas & Electric Co. v. NRDC*, 462 U.S. 87, 103 (1983); *Chevron U.S.A. Inc. v. NRDC*, 467 U.S. 837, 865 (1984) (although *Chevron* was overruled in 2024 regarding judicial deference, its observations about institutional competence remain meaningful). Courts are therefore cautious about creating detailed regulatory regimes in highly technical industries, leaving such policymaking to institutions capable of convening scientific and technical expertise.

E. Innovation

Regulation can strangle or unlock innovation. Poorly designed or overly rigid rules can stifle creativity and penalize novel approaches. When nuanced, sophisticated, and effective rules are put in place catalyzing technological breakthroughs and market growth.

Effective regulation can provide the public, investors, and governments with confidence. For example, regulation of food and drugs gives the public confidence that these products are safe, which catalyzes sales, promotes investment, and thereby results in research into new products. Regulation can establish consumer trust and market certainty.

Effective regulation can also catalyze innovation by establishing clear lanes for an industry to develop within. Predictability and stability are often petri dishes for innovation. When the lines are drawn clearly, it can open up pathways for the future as opposed to attempting to move forward in the grey.

F. Auditing

One of the tools often employed in legal and regulatory schemes is auditing—that is, we make sure that those who are carrying out an activity are checked by those who are a step away from the day-to-day success of that project. These checks can come from both internal or external sources. Internally, companies improve their work sometimes by having a formal internal auditing body. Internal audit—often referred to as the "third line" of defense—can be valuable because these auditors are within the company and can know the company's operations at a granular level. They talk to their colleagues at the water cooler every day, so to speak. But they report to different people—they do not report to the

heads of engineering or product development. Internal auditors often report to the Board of Directors and thereby gain authority and leverage. Internal audit, however, is limited by the fact that they are compensated by the same company and receive performance evaluations by the leadership of that company. In a government agency, this role is often played by the agency's Inspector General. Another model for internal checks comes when companies establish groups who work alongside engineering or product development teams while offering a different perspective or orientation. These colleagues are motivated by the company's success (they are employees), but they do not report to the same leadership. Lawyers, privacy professionals, or others can play this role. As an example, Congress created an Office for Civil Rights and Civil Liberties in the U.S. Department of Homeland Security to provide an internal perspective on the development of new security policies; these professionals were embedded in the Department but approached problems from the perspective of a civil rights and civil liberties specialist. The Office could not stop new security initiatives but could offer perspectives on how to improve policies and practices.[16] Meta created such a unit within its company for similar purposes. These were internal teams reviewing policies and projects from a different perspective.

In terms of external influences, external auditing is also a tool frequently employed in legal and regulatory schemes. External auditors are less effective in the sense that they are not as embedded in the company's systems. But they benefit from impartiality. They can be quite influential, especially if they are compensated by a neutral source and their funding does not depend on the company.

Auditors operate within a well-defined framework of professional standards, and they can face significant professional and legal consequences if they fail to comply with those standards. For example, for private companies audits are generally conducted under Generally Accepted Auditing Standards (GAAS), which are promulgated by the American Institute of Certified Public Accountants (AICPA)'s auditing standard-setting bodies. For public companies, audits are conducted under standards issued by the Public Company Accounting Oversight Board (PCAOB).[17] Traditionally, auditing standards require:

- auditor competence and training,
- independence,
- due professional care,
- adequate planning and supervision,
- obtaining sufficient evidence,
- understanding internal controls,
- proper reporting of findings.[18]

In short, society grants external auditors a specialized role because: they possess specialized expertise; they operate under recognized standards; they owe duties of independence and professional care; and, they can be held accountable when they fail to meet those standards.

Auditing is emerging as a major issue in AI. There is confusion here, though, as terms are conflated—many of the steps AI companies are taking at this time are not focused on creating true external auditors. They are seeking to create teams who can work within the company to: red-team models; test models against established safety and security safeguards; and conduct alignment assessments. It will take substantial effort to create a truly professional cadre of external auditors such as accounting firms have developed over the past many decades. California has passed a statute requiring "Independent Verification Organizations," or,

"IVOs," and it is thought that this may catalyze the development of such an industry.

For external auditing to be meaningful in an AI context, those who are reviewing the work of engineering and product development teams must: be independently funded; be permitted to publish results publicly; and, be permitted to report concerns immediately to government and others in industry.

G. Standards

Another frequent tool of legal and regulatory schemes is the setting of standards. Many times, companies in similar fields decide to lift the entire industry by setting standards for themselves. For example, credit card companies and financial institutions recognized that there needed to be established standards for how electronic payments will be processed—or else the system will collapse. Therefore, the Payment Card Industry Data Security Standard (PCI DSS) was developed. It is a set of security requirements for any organization that stores, processes, or transmits payment card data. It is not a law. It is a private-sector regime created in 2006 by the major card brands (Visa, Mastercard, American Express, Discover, and JCB) and administered by the PCI Security Standards Council (PCI SSC), which they jointly founded. The standard's purpose is to reduce payment card fraud by protecting cardholder data (such as the primary account number) and sensitive authentication data (such as CVV codes and PINs).When standards widely agreed upon have been established, they can be leveraged in several ways. One is for Congress to pass laws that incorporate those standards. Another is for agencies to incorporate those standards into regulatory frameworks. Private parties can also integrate those standards into contract clauses, thereby binding the parties. Finally, standards can become a "standard of care," which can lead to liability being imposed by courts in litigation. It is important to note that in the cybersecurity field, standards and particularly a duty of care leading to liability have been at best elusive. The PCI standards are a good example. The standards are incorporated through contracts, and therefore enforcement runs through the card brands and acquiring banks rather than through a government regulator. Noncompliant entities can face penalties passed down through their acquirer, higher transaction fees, liability for fraud losses after a breach, or, in the extreme, loss of the ability to accept cards. Moreover, several U.S. states (e.g., Nevada, Washington, and Minnesota) have incorporated PCI DSS into statute or safe-harbor provisions. Also, regulators such as the FTC have at times looked to PCI-type controls in assessing whether an entity's data security practices were reasonable.

H. Incentives

Law is fundamentally an incentive-shaping tool. For example, companies learn that if they produce toys that will sell in great quantities but ultimately will cause harm to a certain number of children, they will alter their financial calculus as they understand the cost of future tort lawsuits. If plants produce valuable products but also spill toxins into water systems, companies are incentivized by the law to find ways to produce those products more safely.

Well-designed rules can change the marginal cost of the harmful behavior versus the marginal cost of the safer alternative, then let the regulated party—who has far better information about its own operations than any regulator—choose the most cost-effective way to comply. We understand that in our economic and legal system, this is often more efficient than dictating the method directly (though prescriptive rules still have their place when harms are catastrophic, hard to price, or urgency outweighs efficiency).

Laws often can establish positive incentives for behavior—for example, creating tax breaks if people or companies take certain actions. Another example is creating "safe harbors," which give companies protection from liability in court if they have taken certain steps.

I. Laboratories for Democracy

In our Constitutional system, state and local governments are given great flexibility for action. This is a core benefit of American federalism. Justice Brandeis wrote in a dissenting opinion in 1932: "It is one of the happy incidents of the federal system that a single courageous State may, if its citizens choose, serve as a laboratory; and try novel social and economic experiments without risk to the rest of the country."[19] Women's suffrage started in Wyoming in 1869, five decades before the 19th Amendment federalized the right of women to vote. California set vehicle emissions standards in the 1960s, a forerunner to more comprehensive measures to combat pollution. California enacted a data breach law that was mimicked in dozens of other states, while other states created laws that adopted different mechanisms and tools. Several states, including Texas, Illinois and California, have passed legislation to govern AI. Challenges arise when the states act in ways that are not compatible with one another, making it difficult for companies to determine the appropriate rules to follow. But a biproduct of varying state legislation is that it creates demand to harmonize laws.

J. Current Law

(1) Antitrust

Antitrust laws are passed by Congress to ensure that there is robust competition in each industry. In cybersecurity, however, the country realized more than a decade ago that a great deal depends on collaboration between companies from the same sectors. We need companies to collaborate to: understand the particular threats they face; identify cyber threat indicators that describe "malicious reconnaissance" by bad actors, or the methods bad actors are attempting to use to defeat security controls; share "defensive measures" that companies take to defeat malicious activity; and help one another understand incidents that they confront.[20] Therefore, Congress passed the Cybersecurity Information Sharing Act of 2015 (CISA 2015), which seeks to stimulate a robust information-sharing environment. For example, the law specifically provides that sharing cyber threat indicators or defensive measures does not violate antitrust laws.[21] The tech industry makes daily use of this legislation. Current law permits AI developers to extend their information sharing practices into at least certain areas regarding the security of AI models. It would seem, however, that the CISA 2015 is not crafted to enable sharing on the full range of AI threats, including particularly those that are characterized as "safety" risks (i.e., threats to the physical world, such as through the development of chemical weapons or the production of exploitative materials). Organizations can still seek to share information outside of CISA 2015. However, without the Act's antitrust protections, competition law concerns may discourage or complicate some forms of collaboration. Those concerns become more pronounced when collaboration extends beyond information sharing. Some suggest that the largest companies collaborate on restricting development - agreeing to delay the release of new models until certain safety and security measures are met. This might be struck down by courts as an inherently anti-competitive market allocation. Indeed, four claimants sued four frontier companies arguing that Amodei's open letter calling for collaboration and others agreeing with Amodei's approach is unlawful.[22] This type of coordination might need statutory action by the Congress to create

explicit statutory protections for such actions. Another tool in the toolbox of federal regulators is for the Department of Justice and Federal Trade Commission to issue a letter to industry advising that in certain circumstances they will never seek to prosecute—essentially issuing guidance that permits certain behaviors.

(2) Computer Fraud and Abuse Act (CFAA)

The Computer Fraud and Abuse Act (18 USC §1030) is a bedrock law that governs activity on the internet. The CFAA provides that if someone intentionally accesses a computer without the authorization to do so or accesses a computer in such a way that exceeds their authorized access, a crime has taken place. Even if the government is not willing to prosecute, the statute provides a private right of action—that is, a person or company can sue. The challenge is that the CFAA may not assign criminal liability to a company when an AI agent acts on its own to penetrate the computer system of another company. The law provides, that "Whoever … intentionally accesses a computer without authorization …"[23]Intentionally is a potential bar to charging a company—or an individual at a company—who unleashes an AI agent in a sandbox that ends up escaping the sandbox and attacking other systems. Establishing in those instances that the company intended such an outcome would be difficult; to date it appears that the companies did not in fact intend for the AI agent to engage in such conduct. The Ninth Circuit Court of Appeals recently confronted such a situation. It ruled that as a technical matter an AI agent developed by Perplexity did not access Amazon.com and therefore did not reach the standard at the core of the CFAA whether access by the Agent was authorized.

As the panel explained,

> The Supreme Court has instructed that, "in the computing context, 'access' references the act of entering a computer system itself or a particular part of a computer system, such as files, folders, or databases." *Van Buren*, 593 U.S. at 388 (internal quotation marks omitted). Our focus is thus to ask whether Perplexity uses a tool (the Assistant) to "access" Amazon's computers. On the facts before us, we answer no. It is the user who "accesses" Amazon's computers, with the help of the Assistant to carry out specific acts on Amazon.com. To be sure, Perplexity may receive screenshots of the user's browser and may communicate instructions to the Assistant. But those activities, by themselves, do not mean that Perplexity has "accessed" (gained entry) to Amazon's servers. As part of its analysis, the Court recognized that the CFAA is not adequate to address an emerging technology such as AI Assistants.[24] The Ninth Circuit implied that if there were facts that demonstrated that Perplexity controlled the agent such that Perplexity accessed Amazon.com as contemplated by the CFAA, then there might be a different outcome. The Court also made clear that its decision "do[es] not address whether in other contexts, including tort claims, Perplexity can avoid liability for the Assistant's actions." The court remanded the case to the District Court. The complainants, Amazon, filed an amended complaint alleging Perplexity controlled the AI agent thereby gaining access to Amazon.com as contemplated by the CFAA and added a third claim: "Tortious Interference of Contract."[25]

This area of law bears watching beyond the developments in just one case. As it becomes clearer whether AI developers control the actions of the agents, the legal standard will evolve. It is also worth noting that Senators Josh Hawley (R-MO) and Chris Murphy (D-CT) have introduced the AI Agent Accountability Act, which would amend the CFAA by:

- Imposing criminal and civil liability on AI agent operators who knowingly deploy agents that recklessly cause hacking-related damage or loss.
- Imposing liability on AI developers who fail to implement reasonable safeguards against hacking when they knew or should have known that their systems possessed hacking capabilities.
- Authorizing the U.S. Attorney General and state attorneys general to seek injunctions against AI operators and developers involved in CFAA violations.

It is also worth watching activity under California law. California has a state law parallel to the CFAA, called the California Comprehensive Computer Data Access and Fraud Act (CDAFA). New California AI law seeks to clarify the CDAFA by providing, "it shall not be a defense … that the artificial intelligence autonomously caused the harm to the plaintiff." CA Civil Code Section 1714.46(b).

(3) Tort law

Tort law is the area of civil law that allows a person who has been harmed by another person's wrongful conduct, whether intentional, negligent, or in some cases subject to strict liability, to seek compensation or other legal remedies. Its primary purpose is to compensate injured parties, allocate responsibility for harm, and deter conduct that unreasonably causes injury to people or property (in some cases, even punitive damages can be imposed). In the AI context, many argue that no new laws or regulations are needed because if an AI developer's model is used in such a way as to cause harm, it can be sued and thereby held liable. Tort law has provided a strong incentive to companies in fields like drug production to ensure that their medicines produce the intended effects.There are complications to relying upon tort law. First, in a fast-developing field like AI, a "standard of care" that is the basis for tort liability has not yet been defined and may be slow in coming. In the far more mature field of cybersecurity, at least in the US, a standard of care has yet to be agreed upon and tort liability has been lacking. Second, litigation takes time—a plaintiff must file suit, discovery must take place, trials must be scheduled and held.[26] Therefore, relying on litigation will mean that no new meaningful requirements will be imposed for the foreseeable future. It may be months or years before final judicial resolution that companies can rely upon in understanding whether certain actions involving AI agents withstand legal scrutiny or require modification to assign liability. Finally, it is debatable whether society wishes for the rules governing the development and deployment of AI are set by courts or in settlement agreements approved by courts. Courtrooms are not designed for settling such complex questions.

Part Two: What To Do About It

The following options are organized by who can do what, when, and how. Most of these options require no statute and could begin tomorrow. Others require law, appropriations, and political coalitions to build durable outcomes. Each is intended to build trustworthiness—the demonstrated properties that warrant trust—so that confidence in the American AI tech stack rests on evidence rather than assertion.

Actions AI Developers Can Take Unilaterally

The fastest path to AI governance is the governance a company installs itself. No vote is required, no rule must clear review, and no agency must stand up a new office. What unilateral action cannot supply on its own is proof—and proof is what turns an accord into a durable standard.

It is worth noting that the September 29, 2026 White House Accord on Super Intelligence applies only to the six frontier AI developers who signed it; it does not apply to all of the other companies who are developing AI models nor those who are deploying models in their own environments. The wide range of companies involved in driving the development or adoption of AI could build on the White House accord by taking steps such as:

- Incident reporting: Adopting incident-reporting policies that define, in part, what gets escalated and to whom.
- Transparency: Maintaining model cards and data inventories that document what each model is and what it was trained on.
- Risk Assessments: Completing risk assessments before any new model ships.
- Disclosure: Companies could commit to disclose its work on all three of these points, so outsiders can verify them.
- Sandbox models before they reach power grids, hospitals, financial markets, or defense systems, and fund that testing as a cost of doing business.
- Employ all three layers of review that are described above—establish internal teams of reviewers who will bring a new set of eyes and perspective, particularly focused on safety & security; establish a third line of defense internal audit function; and, hire external auditors with no stake in the company's success. With these layers of review, companies will be assured to have credible independent challenges to their work and a meaningful QA function.

Actions Industry Can Take Together

Some gaps are too large for any single firm to close and too fast for any legislature. Industry can close them by agreeing—on frameworks, on formats, and on what gets shared. Industry could now:

- Set standards: Converge on one governance framework mapped to each layer of the AI tech stack: data, model, infrastructure, application, and cybersecurity and accountability across all of them. The framework may have the incidental effect of setting a *de facto* standard of care that can increase accountability through legal actions. Such a framework could then be tied to procurement rules; be advanced through a NIST profile; or be incorporated into future statutes or regulations. The framework could be built with input from states, civil society, investors, and the international community.
- Share information: Share incident reports, threat feeds, and machine-readable vulnerability data in real time, extending channels such as the Joint Cyber Defense Collaborative's AI playbook.[27.]

- Establish a credible auditing profession: As described above, external auditors need to be independent of the company financially; have the ability to report concerns to law enforcement authorities; and publish their findings regularly. Industry needs to act together to ensure that this is possible. Therefore, industry could create a non-governmental organization (NGO) to fund and oversee independent auditors assigned to AI labs. Independent funding would strengthen auditor credibility, while the NGO could regularly publish findings and report concerns to the government and the public.

Actions the Executive Branch Can Take

The executive branch holds four tools that require no new statute: the power to convene, the power to buy, the power to negotiate, and the power to direct the agencies it already runs. With these attendant tools it could:

- Convene a Federal–State AI Governance Council (CAISSI, Commerce, ONCD, CISA, DOJ, and the states) to document real conflicts between federal and state requirements, the evidence base for any later preemption decision.
- Use its buying power to turn the accord's voluntary commitments into contract terms.[28] In fact, the GSA has already published proposed contract clauses that will govern the purchase of AI tools, which federal agencies will eventually incorporate into new contracts. The Executive Branch could direct OMB and GSA to make demonstrated trustworthiness a condition of award for frontier AI: independent evaluation informed by CAISSI, incident reporting to the buying agency, and access for good faith outside researchers.
- Direct NIST to publish a reference corpus and require every vendor selling AI to the government to disclose how many of its tokens that corpus costs; GSA now prices federal AI use per token.[29]
- Direct State, Commerce, and the United States Trade Representative (USTR) to align taxonomies, testing and documentation practices, and incident-reporting formats with the EU AI Act, which already reaches American firms selling into Europe.[30]

Separately, federal regulators can use authorities already granted to them by Congress to address AI-related conduct within their jurisdictions. Some actions may include:

- Initiate investigations of questionable practices of AI developers. Both the SEC and the Federal Trade Commission, among others, have authorities already in place to undertake these investigations.
- The FTC has, in fact, opened an investigation into OpenAI, Anthropic, and other AI developers to determine whether their AI systems posed consumer risks and whether the companies' conduct violated the FTC Act. The investigation reportedly focused on safety incidents involving autonomous AI agents, representations concerning AI safety, and the adequacy of corporate safeguards. The SEC regulates disclosures and internal controls related to publicly traded companies and securities issuers, among others. Based on its broad authorities, the SEC could:
  - Flag to companies that it will make it a priority in future examinations to scrutinize AI policies and practices. The SEC has, in fact, done this. In its Fiscal Year 2027 Examination Priorities, the Securities and Exchange Commission identified advanced artificial-intelligence technologies as a focus area for examinations, indicating

that registrants should expect scrutiny of their AI governance, cybersecurity controls, disclosures, and compliance practices under existing securities-law obligations rather than under a separate AI-specific regulatory framework.

Broaden its definition of a "cybersecurity incident." That term is currently defined in the SEC's published rules as: "Cybersecurity incident means an unauthorized occurrence, or a series of related unauthorized occurrences, on or conducted through a registrant's information systems that jeopardizes the confidentiality, integrity, or availability of a registrant's information systems or any information residing therein."[31] The SEC could broaden the definition so that the risk to confidentiality, integrity and availability (CIA) is not limited to the registrant's system, but could apply to actions taken by the registrants that limit the CIA of other information systems. Of course, these rules would not yet bind some of the largest Frontier AI companies, who are not yet publicly traded.

- Enforce internal controls and communications channels designed to ensure senior leadership at publicly traded frontier labs are aware of a model's critical technical risks and can assess a related event's materiality before public statements are made.[32] Even limiting enforcement to internal control failures that lead to material financial impacts and quantifiable investor harm (as opposed to purely speculative risks), per the direction of the current Administration, could drive publicly traded labs and those contemplating an initial public offering towards robust model monitoring, risk assessment, and risk communication.
- Enforce broker-dealer AI use violations based on "conflict of interest" and "care" obligations under Regulation BI (Best Interest), "AI-washing" under Section 17(a) of the 1933 Act and Section 10(b) of the 1934 Act, failures in supervisory risk management controls over AI models used for algorithmic and high-frequency trading under Rule 15c3-5 under the 1934 Act, and recordkeeping deficiencies related to the use of AI chatbots with retail clients under Rules 17a-3 and 17a-4 under the 1934 Act.
- Review Form 19b-4 submissions by exchanges, clearinghouses, and other SROs that seek to introduce AI models into institutional market structures.
- Ensure exchanges and alternative trading systems comply with Regulation SCI's (Systems Compliance and Integrity) requirements for rigorous AI model governance and testing, clear escalation procedures in case of model drift or unexpected behaviors, and established termination protocols (e.g., a "kill switch") for models exhibiting drift or erratic operations.

Much of the federal architecture for governing AI already sits inside the Department of Commerce. Commerce funds domestic chip manufacturing through the CHIPS Program Office, controls exports through the Bureau of Industry and Security, sets the technical standards and measurement methods that define cybersecurity and AI risk through the National Institute of Standards and Technology (NIST), and evaluates frontier models through CAISSI.[33] Executive Order 14320 then described a full-stack AI technology package—hardware and data center infrastructure, data pipelines, models, cyber security measures, and applications—in terms that map almost exactly onto those offices.[34]

The Secretary of Commerce could act at every layer of the stack without waiting for new law, including by:

- Directing CAISSI to publish its evaluation methods, the risk categories it tests, and a summary of findings for each of its more than 40 frontier evaluations, withholding classified detail.[35]
- Using the AI Action Plan's directed revision of the AI Risk Management Framework to build a profile with distinct control objectives for data, models, infrastructure, and applications.[36]
- Requiring every American AI Exports Program package to carry CAISSI-cleared models and meet published NIST security baselines.[37]
- Conditioning federal infrastructure support on security baselines: physical and cyber protection for data centers and model weights, hardware supply-chain assurance, and incident reporting.

Actions Congress Can Take

Executive action can start the work. Only Congress can make it durable. While there are a wide range of bills introduced in Congress, several legislative options would give AI governance the authority, evidence, and accountability that voluntary commitments cannot, including:

- Renewing the Cybersecurity Information Sharing Act of 2015, now extended only to December 11, 2026; companies will not share at machine speed under a law that expires every few months.[38][39][40]Congress could expand the CISA 2015 to explicitly permit sharing on AI safety risks, as well as the security risks that are currently within the remit of the law.
- Codify, fund, and empower a federal agency (perhaps the Department of Commerce's CAISSI) as the federal AI systems authority, and thereby provide it authority to create standards and manage the vetting of models before release. That agency, or another, could have an NTSB-style mandate[41] to investigate significant AI incidents and publish what it finds.
- Require pre-deployment certification for high-risk models, including provenance tracing and adversarial testing, paired with post-deployment monitoring and serious-incident reporting, as the EU AI Act already requires for high-risk systems.[42]
  - Pass a revision to the Computer Fraud and Abuse Act, or enact parallel legislation, to clarify when a person or organization is responsible for unauthorized access carried out through an AI agent. Current law does not always clearly assign this responsibility, thus creating legal ambiguity. The clarification should cover cases where the agent acts beyond its operator's specific instructions. Clear rules would strengthen incentives to design, test, and deploy agents more safely.

Actions States Can Take

States hold the police powers that govern employment, housing, education, health care, and law enforcement within their borders, and as laboratories of policy they generate the enforcement signals, observed harms, and feasibility feedback that federal rulemaking otherwise lacks.[43]State legislatures could continue to pass legislation in the absence of Federal action. Already the states are beginning to act in this area. California has enacted the nation's most comprehensive AI regulatory regimes, including frontier-model governance, training-data transparency requirements, AI-content disclosure rules,

employment-related protections, and sector-specific safeguards for health and consumer applications. Illinois initially focused on employment discrimination and government use of AI but has recently joined California in imposing safety, transparency, and incident-reporting obligations on developers of advanced frontier models. Texas law now prohibits harmful or discriminatory uses of AI. However, the Texas law also creates an innovation-oriented governance structure, and empowering the Attorney General to enforce consumer-protection safeguards, while largely avoiding direct regulation of frontier-model development.States could also play a significant role by working to harmonize legislation so that the industry is not stifled by conflicting rules.

States can also play a role in AI safety and security by initiating investigations of company practices and/or filing lawsuits pressuring the companies. Recent actions by state attorneys general illustrate the emergence of a state-based AI enforcement built largely on existing consumer-protection authorities rather than AI-specific legislation. In 2026, the Florida Attorney General filed suit against OpenAI alleging that the company concealed known safety risks, exposed minors to harmful interactions, and misrepresented the safety of ChatGPT, while a bipartisan coalition of attorneys general from eight states brought parallel actions asserting that OpenAI misled consumers about the safety and reliability of its products and failed to implement adequate safeguards for children and vulnerable users. At approximately the same time, the Texas Attorney General filed a similar consumer-protection suit seeking both monetary and injunctive relief. A multistate coalition led by the New York Attorney General launched a broad investigation into OpenAI's advertising practices, child-safety measures, handling of consumer data, and model safety controls.

Collectively, these actions reflect a growing willingness by state attorneys general to use general unfair- and deceptive-practices statutes, privacy laws, and child-protection authorities to regulate the development and deployment of frontier AI systems, effectively positioning state enforcement agencies as some of the most active AI regulators in the United States.

Actions the United States Can Take with Its Allies

Standards are where technological leadership is won or lost. The United States should lead multilateral coalitions to build interoperable global standards for data provenance, model resilience, digital sovereignty, and trusted compute. In 2019, Trump administration officials and allies from more than 30 countries convened in Prague and produced the Prague Proposals, a shared statement of security principles for 5G communications networks. Without that kind of coalition for AI, domestic standards will be undercut by firms shopping for the weakest rules, by foreign influence operations, and by competing claims of AI sovereignty. Moreover, if the United States does not lead such an effort, American companies could be hampered as laws across Europe and Asia come to be applied to their operations. Together, we can:

- Lead multilateral coalitions on interoperable standards for data provenance, model resilience, trusted compute, and digital sovereignty.
- Use the 2019 Prague Proposals, in which officials from more than 30 countries agreed on 5G security principles, as the model.[44]

Appendix A. Options for AI Governance, by Actor, and AI Tech Stack Layer

Option	Core move	Authority	DATA	MODEL	INFRASTRUCTURE	APPLICATION	GOVERNANCE
Actions AI Developers Can Take Unilaterally							
1. Build and disclose internal governance mechanisms	Incident-reporting policies, model cards and data inventories, and risk assessments before any new model ships, disclosed so outsiders can verify them	Voluntary	✓	✓			✓
2. Sandbox before high-consequence deployment	Sandbox testing before a model reaches power grids, hospitals, financial markets, or defense systems, funded as a cost of doing business	Voluntary		✓		✓	
3. Employ three layers of review	Internal safety and security review teams, a third-line internal audit function, and external auditors with no stake in the company's success	Voluntary		✓			✓
Actions Industry Can Take Together							
4. Adopt a unified framework mapped to each stack layer	One governance framework spanning data, model, infrastructure, and application, which could set a de facto standard of care and feed procurement rules, a NIST profile, or future law	Voluntary	✓	✓	✓	✓	✓
5. Share information in real time	Incident reports, threat feeds, and machine-readable vulnerability data, extending channels such as the JCDC AI playbook	Voluntary; but relies on CISA 2015 renewal			✓	✓	
6. Establish a credible AI auditing profession	An industry-created NGO funds and oversees independent auditors assigned to AI labs, publishes findings, and reports concerns to government and the public	Voluntary		✓			✓
Actions the Executive Branch Can Take							
7. Convene a Federal–State AI Governance Council	CAISI, Commerce, ONCD, CISA, DOJ, and the states document real federal–state conflicts as the evidence base for any preemption decision	Executive					✓

8. Set the terms of federal AI procurement	Build on GSA's proposed AI contract clauses; make independent evaluation, incident reporting, and good-faith researcher access conditions of award	Executive	✓		✓	
9. Direct NIST to define the token	NIST reference corpus; vendors selling AI to government disclose its token cost	Executive	✓			
10. Align with the European Union	State, Commerce, and USTR align taxonomies, testing and documentation practices, and incident-reporting formats with the EU AI Act	Executive				✓
11. Investigate AI developer practices	FTC and SEC investigations of questionable practices, such as the FTC inquiry into AI developers' safety representations and safeguards	Existing FTC/ SEC authority			✓	✓
12. Prioritize AI in SEC examinations	Scrutinize registrants' AI governance, cybersecurity controls, disclosures, and compliance, as the SEC's FY2027 priorities signal	Existing SEC authority				✓
13. Broaden the SEC definition of "cybersecurity incident"	Reach registrant actions that compromise the confidentiality, integrity, or availability of other parties' information systems	Existing SEC authority	✓			✓
14. Enforce internal controls on model risk	Controls and communication channels so leaders of publicly traded labs learn of critical model risks and assess materiality before public statements	Existing authority	✓			✓
15. Enforce existing rules on AI use by broker-dealers	Regulation Best Interest, AI-washing, Rule 15c3-5 risk-management controls, and recordkeeping for AI chatbots used with retail clients	Existing authority			✓	
16. Oversee AI in market infrastructure	Review SRO Form 19b-4 filings that introduce AI models; hold exchanges and ATSs to Regulation SCI testing, escalation, and kill-switch requirements	Existing authority	✓		✓	
17. Publish what CAISSI learns	Release evaluation methods, risk categories, and a summary of findings for each frontier evaluation, withholding classified detail	Secretary of Commerce	✓			
18. Rebuild the AI RMF layer by layer	Use the AI Action Plan's directed revision to build a profile with distinct control objectives for data, models, infrastructure, and applications	Secretary of Commerce	✓	✓	✓	✓

19. Attach trust conditions to the AI Exports Program	CAISSI-cleared models and published NIST security baselines in every export package	Secretary of Commerce			✓	✓	✓	
20. Condition infrastructure support on security baselines	Physical and cyber protection for data centers and model weights, hardware supply-chain assurance, and incident reporting	Secretary of Commerce				✓		
Actions Congress Can Take								
21. Renew and expand CISA 2015	Renew the Act beyond its December 11, 2026, expiration and extend its protections to AI safety risks as well as security risks	Statute				✓	✓	✓
22. Codify, fund, and empower a federal AI systems authority	An agency (perhaps CAISI) sets standards and vets models before release; it or another agency holds an NTSB-style mandate to investigate significant AI incidents	Statute and funding			✓			✓
23. Institutionalize model lifecycle governance	Pre-deployment certification for high-risk models, with provenance tracing and adversarial testing, plus post-deployment monitoring and serious-incident reporting	Statute		✓	✓			
24. Clarify CFAA liability for AI agents	Revise the CFAA, or enact parallel law, to assign responsibility for unauthorized access carried out through an AI agent, including beyond its operator's instructions	Statute					✓	✓
Actions States Can Take								
25. Legislate in the absence of federal action	Use state police powers and continue enacting AI laws, as California, Illinois, and Texas have on frontier models, transparency, and harmful uses	State law		✓	✓		✓	
26. Harmonize state legislation	Coordinate across states so industry does not face conflicting rules	State law						✓
27. Enforce existing consumer-protection law	Attorney general investigations and suits under unfair- and deceptive-practices, privacy, and child-protection statutes	Existing state authority					✓	✓
Actions the United States Can Take with Its Allies								
28. Lead interoperable global AI standards	Multilateral coalitions on data provenance, model resilience, trusted compute, and digital sovereignty, modeled on the 2019 Prague Proposals	Diplomatic		✓	✓	✓		✓

Shaded cells mark the layer or layers of the AI Tech Stack each option acts on.

"Authority" notes whether an option needs new law: voluntary, executive, secretarial, and existing-authority options can begin without it; statutory options cannot.

Endnotes

1 15 USC § 9401(3) -- artificial intelligence is defined as: (3) Artificial intelligence The term "artificial intelligence" means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to— (A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action.

2 Center for AI Safety, "Statement on AI Risk," May 30, 2023. https://www.safe.ai/work/statement-on-ai-risk

3 OpenAI, "Introducing Superalignment," July 5, 2023. https://openai.com/index/introducing-superalignment/

4 Dario Amodei, "The Urgency of Interpretability," April 2025. https://www.darioamodei.com/post/the-urgency-of-interpretability

5 Anthropic, "Disrupting the first reported AI-orchestrated cyber espionage campaign," November 13, 2025. https://www.anthropic.com/news/disrupting-AI-espionage

6 Anthropic, "Detecting and preventing distillation attacks," February 23, 2026, https://www.anthropic.com/news/detecting-and-preventing-distillation-attacks; Google Threat Intelligence Group, "GTIG AI Threat Tracker: Distillation, Experimentation, and (Continued) Integration of AI for Adversarial Use," February 12, 2026, https://cloud.google.com/blog/topics/threat-intelligence/distillation-experimentation-integration-ai-adversarial-use

7 OpenAI, "The Hugging Face incident and the road ahead," August 2026, https://openai.com/index/hugging-face-incident-and-the-road-ahead/; Hugging Face, security incident disclosure, July 16, 2026, https://huggingface.co/blog/security-incident-july-2026

8 "Pacing the Frontier" open letter, July 28, 2026. https://pacingthefrontier.com

9 Dario Amodei, "We Must Pace the Frontier," September 12, 2026. https://darioamodei.com/post/we-must-pace-the-frontier

10 Sam Altman, post on X, September 12, 2026. https://x.com/sama/status/2098811563415150910?lang=en

11 OpenAI, "An agent used DNS to reach an external chatbot," Misalignment Report, September 2026, https://alignment.openai.com/misalignment-reports/an-agent-used-dns-to-reach-an-external-chatbot/; Fortune, "OpenAI says its AI agents escaped a secure 'sandbox' again last weekend and it is pausing training for a second time," September 26, 2026, https://fortune.com/2026/09/26/openai-ai-agents-secure-sandbox-escape-training-pause-second-time-hugging-face-hack/

12 Paladin Global Institute, The AI Tech Stack: A Primer for Tech and Cyber Policy, 2025. https://www.paladincapgroup.com/wp-content/uploads/2026/08/AI-Tech-Stack-Report.pdf; White House Accord on Super Intelligence: Joint Commitment on Frontier Responsibilities, September 29, 2026, https://www.forbes.com/sites/saradorn/2026/09/30/white-house-releases-accord-between-billionaire-ai-execs-heres-what-it-says/.

13 Paladin Global Institute, The Governance Layer: A Primer for Tech and Cyber Policy, 2026. https://www.paladincapgroup.com/wp-content/uploads/2026/08/Paladin_AI-Tech-Stack_Governance-Primer.pdf

14 White House Accord on Super Intelligence: Joint Commitment on Frontier Responsibilities, September 29, 2026, https://www.forbes.com/sites/saradorn/2026/09/30/white-house-releases-accord-between-billionaire-ai-execs-heres-what-it-says/.

15 See, e.g., "How AI Backlash Shaped the 2026 Midterms," Time, Nov. 4, 2026, available at https://time.com/article/2026/11/04/midterm-elections-results-ai-industry-data-centers/16 See a collection of oral histories and documents that demonstrate the function of the Office for Civil Rights and Civil Liberties, https://mercyhurst-archives.starter1ua.preservica.com/portal/en-US.

17 The PCAOB's authority comes from the Sarbanes-Oxley Act of 2002, a federal statute. Thus, the standards are not purely a matter of industry self-regulation. For public-company auditors, Congress created a statutory oversight regime, and the Securities and Exchange Commission oversees it (the PCAOB).

18 See pcaob.org.

19 U.S. Const. amend. X; New State Ice Co. v. Liebmann, 285 U.S. 262, 311 (1932) (Brandeis, J., dissenting).

20 Cybersecurity Information Sharing Act, 6 U.S.C. 1501(6) and (7).

21 6 U.S.C. 1504(e).

22 https://storage.courtlistener.com/recap/gov.uscourts.cand.479357/gov.uscourts.cand.479357.1.0.pdf

23 18 U.S.C. 1030(a)(2).

24 https://cdn.ca9.uscourts.gov/datastore/opinions/2026/08/04/26-1444.pdf

25 https://storage.courtlistener.com/recap/gov.uscourts.cand.459191/gov.uscourts.cand.459191.122.0.pdf

26 For example, as referenced above, the Northern District of California received a complaint that Perplexity caused tortious interference in contract between Amazon and its customers because customers agreed to Terms of Use that included provisions relevant to the use of AI Agents. Amazon is alleging that Perplexity intentionally interfered with customer compliance with those terms. An excerpt of the term in dispute is as follows: "No Agent may access, use, or interact with Amazon Services unless, at all times, it identifies itself and operates in strict accordance with the requirements of Section 3 of these Agent Terms."

27 Cybersecurity and Infrastructure Security Agency, Joint Cyber Defense Collaborative, AI Cybersecurity Collaboration Playbook (January 14, 2025). https://www.cisa.gov/sites/default/files/2025-01/JCDC%20AI%20Playbook.pdf

28 U.S. General Services Administration, "OneGov," https://www.gsa.gov/buy-through-us/purchasing-programs/multiple-award-schedule/onegov.

29 "GSA Expands OneGov Partnership with OpenAI," MeriTalk, September 2026 (reporting GSA's September 10, 2026, news release on token-based pricing beginning October 1), https://www.meritalk.com/?p=209515.

30 Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonized rules on artificial intelligence (Artificial Intelligence Act), art. 2(1), OJ L, 2024/1689, https://eur-lex.europa.eu/eli/reg/2024/1689/oj.

31 Item 106(a) of Regulation S-K (17 C.F.R. § 229.106(a)).

32 This follows precedent established under Rule 13a-15(a) enforcement actions by the SEC against First American Financial Corporation (2021), Unisys Corp. (2024), Avaya Holdings Corp. (2024), Check Point Software Technologies (2024), and Mimecast (2024). First American Financial, in particular, involved an internal controls and communications failure related to a critical product technology risk that was not communicated in a timely manner to corporate leadership. This resulted in delayed mitigation and inaccurate public statements. See In the Matter of First American Financial Corp., Securities Exchange Act of 1934 Release No. 92176, Administrative Proceeding File No. 3-20367 (June 14, 2021), https://www.sec.gov/files/litigation/admin/2021/34-92176.pdf.

33 The White House, Winning the Race: America's AI Action Plan (July 23, 2025), https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf.

34 Executive Order 14320, "Promoting the Export of the American AI Technology Stack," July 23, 2025, https://www.govinfo.gov/content/pkg/DCPD-202500790/pdf/DCPD-202500790.pdf.

35 National Institute of Standards and Technology, "CAISI Signs Agreements Regarding Frontier AI National Security Testing," May 2026, https://www.nist.gov/news-events/news/2026/05/caisi-signs-agreements-regarding-frontier-ai-national-security-testing.

36 The White House, Winning the Race: America's AI Action Plan (July 23, 2025), https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf.

37 Executive Order 14320, "Promoting the Export of the American AI Technology Stack," July 23, 2025, https://www.govinfo.gov/content/pkg/DCPD-202500790/pdf/DCPD-202500790.pdf.

38 Congressional Research Service, The Cybersecurity Information Sharing Act of 2015: Expiring Provisions, IF12959 (updated September 2026), https://www.everycrsreport.com/files/2026-09-03_IF12959_e88b818ccb3aca205dcf1ba233b57cd55207205b.html; extension to December 11, 2026, under the Continuing Appropriations and Extensions Act, 2027, as noted in CRS's revision history

39, 40,41 49 U.S.C. § 1131 (general authority of the National Transportation Safety Board); 49 U.S.C. § 1135 (safety recommendations).

42 Regulation (EU) 2024/1689 (Artificial Intelligence Act), arts. 51(2), 72, and 73.

43 U.S. Const. amend. X; New State Ice Co. v. Liebmann, 285 U.S. 262, 311 (1932) (Brandeis, J., dissenting).

44 Prague 5G Security Conference, "The Prague Proposals: Chairman Statement on Cyber Security of Communication Networks in a Globally Digitalized World," May 3, 2019, Government of the Czech Republic.