Arnold&Porter

July 31, 2025

America’s AI Action Plan: What Full Steam Ahead Means for Your Company

Advisory

By

Peter J. Schildkraut ,

Travis Annatoyn ,

Eun Young Choi ,

Deborah A. Curtis ,

Ronald D. Lee ,

Thomas A. Magnani ,

Soo-Mi Rhee ,

Sandra E. Rizzo ,

Allison B. Rumsey ,

Ethan G. Shenkman ,

Junghyun Baek ,

Daniel M. Elsen-Rooney ,

Emily Orler

Underscoring the centrality of artificial intelligence (AI) to U.S. economic and national security strategy, the White House has released America’s AI Action Plan — a sweeping policy roadmap outlining how the Trump administration intends the United States to win what it calls the “AI race.” Accompanying the release of the Action Plan, President Trump signed three executive orders: Preventing Woke AI in the Federal Government (Preventing Woke AI EO), Accelerating Federal Permitting of Data Center Infrastructure (Permitting EO), and Promoting the Export of the American AI Technology Stack (Export EO).

Three “pillars” underlie the administration’s strategy: accelerating innovation, building American AI infrastructure, and leading in international AI diplomacy and security. They add up to a vision of an unhindered path for AI development and deployment in the United States while driving adoption of American AI systems abroad and ensuring that allies and partners support U.S. efforts to remain ahead of China.

Other actors, of course, have their own visions — sometimes aligned, and sometimes at cross-purposes, with the administration’s. Below, we break down the Action Plan’s implications for businesses.

Key Takeaways for AI Developers, Deployers, Infrastructure Providers, and Their Investors

- The Trump administration will use the tools at its disposal to ensure government does not slow down AI innovation. While the administration recognizes that some limits may be necessary to achieve other objectives, it clearly will tip the scales to push the AI revolution forward. The administration’s toolbox, however, may not be able to relieve developers and deployers from the growing body of state and local AI regulation, let alone regulation from jurisdictions outside the United States.
- AI developers will have to adjust the guardrails they put on large language models (LLMs) they sell to the federal government. These

LLMs will have to display “[t]ruth-seeking” and “[i]deologically [n]eutrality” although, apparently, the federal government will not require guardrails against output disfavored by the administration.

- The National Institute of Standards and Technology (NIST) will remove references to misinformation; diversity, equity, and inclusion (DEI); and climate change from its AI Risk Management Framework (RMF). Deployers that have relied on the RMF for AI governance may have to supplement the RMF if they wish to continue to protect against risks related to these concepts.
- The administration clearly supports the right of AI developers to use copyrighted materials to train their models as “fair use,” but the administration has limited ability to shield AI developers from potential copyright infringement exposure. Although early district court decisions have been favorable to AI developers, until the U.S. Supreme Court or an act of Congress definitively decides the fair use question, AI developers should proceed with caution in using copyrighted materials to train their AI models without permission from rightsholders.
- The administration wants to instill a “dynamic, ‘try-first’ culture for AI across American industry” — even in regulated industries — through regulatory sandboxes and supporting the development of the AI evaluation ecosystem to give deployers greater confidence in AI systems.
- The administration will reduce the burdens of environmental permitting on providers of AI infrastructure (principally, data centers and related energy projects). It also seeks to enhance the electric grid’s ability to deliver reliable power to data centers.
- To address the labor-market changes from AI, the administration will prioritize AI skill development; retraining of displaced workers; and ensuring the availability of workers with the skills to build, operate, and maintain AI infrastructure.
- To secure continued American AI leadership, the administration will invest in AI research and development, including by expanding the availability of federal datasets. The administration also will back research into AI interpretability, control, and robustness against adversarial attacks.
- The administration will harden critical infrastructure against AI-enabled cybersecurity threats while leveraging AI itself for cyber defense. The administration’s strategy centers on standing up a new hub for threat intelligence.
- The administration will promote export of full-stack U.S. AI packages to allies and partners, in large part to ensure allies and partners do not become dependent on AI technology from China.
- The administration will increase its efforts to prevent Chinese use of U.S. technology to close the AI gap between the countries. It plans to expand AI-related U.S. export controls and associated enforcement to deny China and other “foreign adversaries” access to U.S.-developed AI compute, protect U.S. competitive advantages in developing components and products across the semiconductor supply chain, and press other countries to align their export-control regimes with that of the United States.
- The U.S. Department of Commerce (DOC) will continue to work with frontier AI developers to understand the national security risks from AI, particularly as they relate to proliferation of chemical, biological, radiological, nuclear, or explosives weapons capabilities.

Contents

Innovation Without Interference: A Deregulatory Doctrine and Its Limits

The Action Plan contains a clear message: regulation must not stifle innovation. At the outset of his administration, President Trump rescinded prior executive actions that he views as overly cautious. In the Action Plan, the White House again directs agencies to identify and dismantle “unnecessar[y]” regulatory barriers to AI development or deployment.

Focus on the U.S. Federal Trade Commission (FTC)

The White House calls for a review of FTC “investigations commenced under the previous administration to ensure that they do not advance theories of liability that unduly burden AI innovation.” It likewise seeks a review of “all FTC final orders, consent decrees, and injunctions, and, where appropriate, seek[s] to modify or set-aside [sic] any that unduly burden AI innovation.”

Under former Chair Lina Khan, the FTC tried to establish itself as perhaps the leading regulator of the AI industry through, among other things, its never-completed privacy rulemaking; an investigation, according to The Washington Post, into whether inaccurate output from OpenAI’s ChatGPT unfairly or deceptively resulted in reputational harm to consumers (see our July 2023 Blog); a settled investigation into Rite Aid’s use of AI-based facial-recognition surveillance technology; and an enforcement sweep the FTC called “Operation AI Comply.” Last year, the current Chairman Andrew Ferguson decried the FTC’s “aggressive move into AI regulation” under his predecessor as “premature.”

Following the Action Plan, we expect the current FTC to undo one Operation AI Comply consent agreement, involving a generative AI company called Rytr. In that case, the FTC charged — over the strong dissents of now-Chairman Ferguson (who labeled the FTC’s theory “inconsistent with our precedents and common sense”) and Commissioner Melissa Holyoak — that Rytr had unlawfully committed “unfair or deceptive acts or practices” by “furnishing others with the means and instrumentalities to engage in … deceptive practices.” Rytr chose to settle. Until the consent agreement is set aside, the precedent may leave other generative AI companies to fear that they, too, are at risk of investigation and potential liability merely for providing a service that others could put to ill use.

The FTC might also seek to modify the Rite Aid settlement, which requires the pharmacy chain to establish, implement, and maintain an extensively detailed risk-management program before using any automated biometric security or surveillance system. Former Commissioner Alvaro M. Bedoya hailed these requirements as a “baseline for … a comprehensive algorithmic fairness program,” but the current commissioners may see them as premature regulation.

Pushing Back on State AI Regulation

States have gotten ahead of the federal government in regulating AI, with several — California, Colorado, Texas, and Utah — adopting fairly broad statutes and rules while other significant efforts remain pending. Whether fearing a patchwork of different regulations, opposing regulation altogether, or somewhere in between, the technology industry largely has sought federal help to rein in the states. The proposed moratorium on state AI regulation in the One Big Beautiful Bill budget reconciliation legislation could not secure majority support in the Senate and was dropped. A watered-down version has reappeared in the Action Plan, which directs federal agencies to consider state AI regulatory climates when disbursing federal funds related to AI. We will have to see how much this incentive will inhibit state policymakers.

The Action Plan also recommends that the U.S. Federal Communications Commission “evaluate whether state AI regulations interfere with the agency’s ability to carry out its obligations and authorities under the Communications Act of 1934.” On its face, this proposal is rather surprising because state AI regulations have little to do with the FCC’s oversight of the telecommunications, broadcast, and pay-TV industries. Instead, this recommendation may harken back to the first Trump administration’s effort to have the FCC interpret Section 230 of the Communications Decency Act (which is contained within the Communications Act) to limit the statute’s shield against liability for social media platforms’ content-moderation decisions. While the FCC had never claimed authority to interpret Section 230, its general counsel at the time released a long blog post explaining that Congress, indeed, had delegated that authority to the agency. Under this analysis, the FCC could adopt a rule preempting state laws that impose liability on AI providers for wrongdoing by their customers. If it does, however, the Supreme Court’s *Loper Bright* decision undoing *Chevron* deference may undermine the force of such a rule — at least to some extent.

The Limits

In a global economy, other countries get a vote on the degree to which regulation should slow innovation. Companies that want access to their markets will have to comply with their rules and may decide it is not worth the cost to develop or deploy AI models and systems differently depending on local requirements.

Even domestically, the Trump administration can do only so much on its own. Wholesale preemption of state regulation will require congressional action. The debate over the One Big Beautiful Bill Act revealed that there are not 50 senators who support preemption, let alone the 60 required for legislation subject to a filibuster. Indeed, there may not be majority support in the House of Representatives either, as some Republicans who voted for the bill said they did not know it contained the moratorium on state regulation.

Moreover, neither the Trump administration nor congressional Republicans propose to make "the AI did it" a defense to enforcement of generally applicable laws. In other words, a company still cannot use AI to do something that a human cannot do lawfully.

A New Standard for AI: "Truth-seeking" and "Ideological Neutrality"

The Preventing Woke AI EO accompanying the Action Plan introduces novel requirements for federal AI procurement: "[t]ruth-seeking" and "[i]deological [n]eutrality," which the order terms the "Unbiased AI Principles." It instructs federal agencies to procure only LLMs that are demonstrably free from what the administration characterizes as "top-down ideological bias," including concepts associated with DEI; critical race theory; and other frameworks the administration opposes.

Agencies apparently do not have to mandate that LLMs have guardrails that *prevent* such output if those "partisan or ideological … judgments are prompted by or otherwise readily accessible to the end user." Further guidance on the Unbiased AI Principles will come from the Office of Management and Budget by November 20, 2025. AI procurement contracts will have to make vendors liable for decommissioning costs for noncompliance with the Unbiased AI Principles after a reasonable cure period. AI developers may have to create separate models or systems for federal customers if they also sell to customers, or in jurisdictions, expecting models and systems to have guardrails inconsistent with the Unbiased AI Principles.

The Action Plan separately requires NIST to remove references to misinformation, DEI, and climate change from its AI Risk Management Framework. The RMF has been a leading AI governance tool, and compliance with the latest version of the RMF is an element of an affirmative defense to violations of state laws, including the Colorado AI Act and the Texas Responsible Artificial Intelligence Governance Act. Deployers that have relied on the RMF for AI governance may have to supplement the RMF if they have business or legal commitments to protect against risks the RMF no longer will address.

President Trump Weighs in on AI Model Training

Neither the Action Plan nor any of the accompanying executive orders addresses the debate over whether AI model training constitutes copyright infringement or fair use, which is playing out in courtrooms across the country. However, President Trump weighed in on this issue at the "Winning the AI Race" summit during which he signed the executive orders. In his speech, President Trump clearly came down on the side of the AI developers, saying, "You can't be expected to have a successful AI program when every single article, book, or anything else that you've read or studied, you're supposed to pay for … . [I]f you read an article and learn from it, we have to allow AI to use that pool of knowledge without going through the complexity of contract negotiations, of which there would be thousands every time we use AI." The president seemed particularly concerned that worries about liability for copyright infringement would hinder U.S. companies and give Chinese AI developers a competitive advantage.

The president's remarks give credibility to reports that he recently fired the Register of Copyrights because the U.S. Copyright Office, which the Register leads, issued a report on AI model training that expressed significant skepticism that training AI models with copyrighted material qualifies as fair use. It remains to be seen whether the Copyright Office will rescind that report once new leadership is installed (if the former Register's suit to regain her position does not succeed).

Other than influencing the Copyright Office's official position on AI model training as fair use, President Trump has limited ability to address the issue. Only an act of Congress can enshrine a specific fair use exception for AI model training in the U.S. Copyright Act. Until then, the fair use question is in the hands of the judiciary. The only two district court decisions directly addressing the issue have both found that training AI models on copyrighted material is fair use. However, dozens of other cases have yet to be decided, and no appellate court has addressed the issue yet. If Congress does not act, the issue ultimately is likely to be decided by the U.S. Supreme Court in several years. Until then, AI developers must continue to operate in an environment of significant uncertainty regarding their potential copyright infringement exposure.

Open Models, Open Markets

The administration seeks to promote open-source and open-weight AI models as a complement to the proprietary or "closed" models offered by most of the leading developers. Freely available open models enable startups and academic researchers to innovate without dependence on a larger provider.

Likewise, to improve startups' and academics' access to computing resources for model and system development, the federal government will work with industry to develop financial markets for compute access — akin to spot and forward markets for commodities. The administration will also expand the National AI Research Resource to provide researchers with greater access to compute, data, and models.

Promoting AI Adoption in Critical Sectors

The successful deployment of artificial intelligence across the U.S. economy hinges not only on technological readiness, but also on institutional willingness to experiment, adapt, and scale. The administration wants to help businesses overcome their hesitancy to adopt AI solutions due to "a variety of factors, including distrust or lack of understanding of the technology, a complex regulatory landscape, and a lack of clear governance and risk mitigation standards." To lead, instead, to a "dynamic, 'try-first' culture for AI across American industry," the administration will offer regulatory sandboxes ("centers of excellence") and will support national standards development.

Regulatory sandboxes — structured environments where companies can test AI tools under regulatory supervision — offer proving grounds for novel applications without regulatory risk. While the Action Plan specifically mentions the U.S. Food and Drug Administration and the U.S. Securities and Exchange Commission, other agencies may take part, allowing developers to demonstrate efficacy and safety while regulators observe real-world performance. For private-sector actors, early participation in sandboxes may provide a competitive edge and shape future regulatory frameworks — or at least yield greater visibility into evolving compliance expectations. However, participation will likely require transparency, data sharing, and other obligations that companies may find problematic.

To make AI more trustworthy for businesses, especially regulated ones, the administration also plans to promote the nascent AI evaluation ecosystem. The Action Plan charges NIST with developing guidance and resources for regulators evaluating AI systems used by their regulatees. At the same time, it directs federal science agencies to support research into AI metrics and evaluation. Better understanding of how to measure and evaluate AI systems ultimately will give companies greater confidence in managing the risks of deployment.

AI Infrastructure Permitting Reform: "Build, Baby, Build"

The Action Plan and Permitting EO attempt to tackle chronic environmental permitting challenges that are now, according to the White House, an obstacle to building AI infrastructure. Together, they represent the most significant federal initiative to date aimed at streamlining permitting processes for data centers and the energy projects necessary to support AI development.

AI development is the latest front in the long-running battle for permitting reform. Building major infrastructure projects in the United States often requires numerous federal, state, tribal, and local authorizations, many of which are subject to judicial review. But, despite bipartisan consensus on the need for reform, Democrats and Republicans have struggled to find common ground on implementation. (See our explanation of the recent monumental changes to National Environmental Policy Act (NEPA) procedures.) By creating urgent demand for unprecedented numbers of energy infrastructure projects to support new large-scale data center operations, the AI revolution puts further strain on this system.

This latest, AI-focused effort at reform through executive action uses available tools in some novel ways. To encourage development of data centers, the White House uses federal lands and financing, and the U.S. Department of Energy has already announced four sites available for use. The Permitting EO also offers federal financial assistance, by directing all relevant agencies to identify any existing financial support that can be used for certain data centers and associated infrastructure (Qualifying Projects). (For energy infrastructure, such Qualifying Projects notably include "*dispatchable* baseload energy sources" such as natural gas, coal, nuclear, and geothermal, thereby excluding non-dispatchable solar and wind projects.) To ease the environmental review burden typically associated with use of these tools, the White House directs relevant agencies to conduct programmatic consultation under the Endangered Species Act "for Qualifying Projects that will occur over the next 10 years." (This programmatic consultation will obviate or narrow any site-specific analysis for Qualifying Projects falling under its ambit.) The administration also interprets NEPA not to apply to federal financial assistance when that assistance accounts for less than half of total project cost.

More broadly, the Permitting EO also directs the U.S. Environmental Protection Agency to aid in siting data centers and other associated infrastructure on brownfield and superfund sites; the Permitting Council to enable use of the FAST-41 process; and relevant agencies to potentially expand existing fast-track reviews — categorical exclusions under NEPA and nationwide permits to impact waters of the United States under Section 404 of the Clean Water Act. And, in a point of innovation, the Action Plan makes clear that AI will be a part of the solution — by giving support to agencies using AI to expedite permitting processes. (For more detailed analysis of the permitting aspects of the Action Plan and Permitting EO, see our July 2025 Blog.)

These federal moves do not guarantee smooth sailing for siting data centers and associated infrastructure. In a countervailing trend, local municipalities are starting to flex their zoning and land-use authorities to restrict data center development. For example, the City of Atlanta recently tightened zoning restrictions by requiring City Council-approved special use permits for data centers and prohibiting data centers in certain commercial districts. Similarly, Fairfax County, Virginia has imposed size limitations in designated districts and established setback requirements for such facilities.

Powering the AI Revolution

The Action Plan recognizes the impact the electric grid will have on AI innovation. The administration emphatically desires to prevent premature decommissioning of critical power generation resources, stressing the importance of resource adequacy and capacity sufficiency in the Action Plan. It also encourages innovative solutions like new approaches to demand response and promotes grid management technologies and transmission upgrades.

Interconnection queue backlogs have extended — sometimes by many years — the timeline for adding new generation resources to the system. The administration proposes prioritizing interconnection of dispatchable power supply resources, including enhanced geothermal, nuclear fission, and nuclear fusion, instead of the current typical processing of requests in the order of submission.

In the Action Plan, the administration also endorses aligning the financial incentives for power generators with the goals of achieving grid stability and meeting the system's needs. Given the accompanying references to dispatchable resources (those that can respond to a request to produce power), this proposal may refer to providing additional payments to generators that have desirable reliability attributes (e.g., availability, fuel assurance, ramp-up capability, voltage stability, rapid start-up, and long-duration energy at high output) instead of paying the same unit price to each generator. Some regional grid operators, the Midcontinent Independent System Operator in particular, have undertaken to identify essential reliability attributes as a first step in this process. The Action Plan endorsement may lead other regional operators to follow suit.

Workforce Development for the AI Era

The AI Action Plan addresses both the promise and peril of disruptive technological change. To prepare students and workers for tomorrow's economy, the U.S. Department of Labor, the U.S. Department of Education, the U.S. National Science Foundation, and the DOC will prioritize AI skill development in efforts that they fund. The U.S. Department of the Treasury will issue guidance on tax-free employer reimbursement for AI training, helping companies to invest in their workers' futures.

Meanwhile, rapid retraining and proactive upskilling funds will offer a safety net for workers whose roles may be automated or fundamentally transformed by AI. The creation of the AI Workforce Research Hub, among other efforts, will study how AI adoption affects the labor market, enabling policymakers and employers to anticipate shifts in demand and target interventions more effectively.

The success of America's AI strategy also hinges on the availability of a skilled workforce capable of building, operating, and maintaining the infrastructure that powers AI systems. The federal government is launching a national initiative to identify high-priority occupations, define skill frameworks, expand apprenticeships, and align career and technical education programs with industry needs. Employers in the construction, manufacturing, and energy sectors should anticipate new funding opportunities for workforce training and upskilling. Community colleges and technical schools likely will play a central role in these efforts.

Support for New Advances

The administration plans to take a number of steps both to propel AI development and to capitalize on AI capabilities to drive progress in other fields.

- To secure leadership in advanced manufacturing — including autonomous drones, robotics, and self-driving vehicles — the administration will prioritize investment in AI and robotics as core enablers for innovation in both manufacturing and logistics.
- The administration will provide support for automated, cloud-enabled laboratories to conduct research using AI.
- The administration will make high-quality datasets available for AI-enabled research by assembling and facilitating access to federal and federally funded data.
- The upcoming National AI R&D Strategic Plan will emphasize investment in novel theoretical, computational, and experimental approaches to further AI science.
- The administration will launch a cross-agency effort to achieve breakthroughs in AI interpretability, control, and robustness against adversarial attacks — especially for national security applications.

Bolstering Cybersecurity for Critical Infrastructure, Homeland Security, and National Security

As AI systems become embedded in the cyber and critical infrastructure of the U.S. economy, they are exposed to cyber attacks. The Action Plan contains the outline of a multi-pronged strategy to harden infrastructure against AI-enabled threats while leveraging AI itself for cyber defense.

At the center of this strategy is the creation of an AI Information Sharing and Analysis Center (AI-ISAC), led by the U.S. Department of Homeland Security in collaboration with the DOC and the Office of the National Cyber Director. The AI-ISAC will serve as a hub for threat intelligence, enabling infrastructure operators, technology vendors, and federal agencies to share information about vulnerabilities, threats, and mitigation strategies.

The administration also emphasizes the importance of secure-by-design, robust, and resilient AI systems used in critical infrastructure, homeland security, and national security. Developers of such AI systems should focus on safeguards against adversarial inputs, data poisoning, and other attacks, and on the detection of malicious activities. Federal agencies will issue guidance and standards to support this effort while enhancing their incident-response capacity related to AI and AI systems.

Exporting the Complete AI Stack to Allies and Partners

In the Action Plan, the Trump administration calls for the export of full-stack U.S. AI packages to allies and partners. It asserts that the diffusion of U.S. AI technology both encourages countries to "join America's AI alliance" and prevents its allies from becoming "dependent on foreign adversary technology." Therefore, the administration requires the DOC to gather AI trade proposals from industry consortia and to coordinate with other agencies to "facilitate deals that meet U.S.-approved security requirements and standards." The accompanying Export EO gives the DOC, U.S. Department of State (DOS), and Office of Science and Technology Policy (OSTP) 90 days to "establish and implement the American AI Exports Program … to support the development and deployment of United States full-stack AI export packages."

Enforcing, Expanding, and Aligning Export Controls

The administration also plans to expand AI-related U.S. export controls and associated enforcement to deny China and other "foreign adversaries" access to U.S.-developed AI compute, protect U.S. competitive advantages in developing components and products across the semiconductor supply chain, and induce other countries to align their export-control regimes with that of the United States.

First, the Action Plan includes a call for "creative approaches" to enforcement to prevent U.S. foreign adversaries from obtaining chips that enable advanced AI compute. Under the Action Plan, the DOC, OSTP, and National Security Council will partner with industry to exploit existing location-verification features on chips to prevent diversion to countries of concern. (Relatedly, in May 2025, twin bills were introduced before both the U.S. House of Representatives and U.S. Senate — both titled the Chip Security Act — calling for AI chips to be outfitted with tracking technology to implement location verification. The bills have bipartisan support but have not been put to a vote.)

In addition, the Action Plan recommends that the DOC collaborate with intelligence community officials on enforcement. The collaboration is to include enhanced monitoring of emerging technology developments in AI compute to ensure that possible diversion destinations are fully covered. This enhanced monitoring may entail end-use monitoring in countries that present a high risk of diversion. This focus on diversion continues from guidance the DOC's Bureau of Industry and Security (BIS) announced in May 2025, which identified transactional and behavioral red flags that signal a risk of illicit diversion of advanced computing items. (For a discussion of BIS' May 2025 guidance, please refer to our May 2025 Advisory). We expect to see increased enforcement against such diversion.

The Action Plan also calls for "new measures to address gaps in semiconductor manufacturing export controls." BIS has continuously broadened export control measures targeting semiconductor manufacturing equipment, including through new rules announced in October 2022, October 2023, and September 2024. These rules, however, focused on "major systems necessary for semiconductor manufacturing." The new measures will target semiconductor manufacturing sub-systems. However, the Action Plan provides scarce detail on these expected new controls, and the administration has yet to offer further guidance.

The Action Plan also envisages greater U.S. global leadership in shaping AI-related export controls, including through pressure on U.S. partners and allies to align with the U.S. export-control regime. The United States has previously incentivized its allies and partner countries to adopt certain export controls aligning with U.S. diplomatic and national security objectives. For instance, in 2024, the Biden administration engaged with the Netherlands and Japan to align Dutch and Japanese export controls on semiconductor manufacturing items with U.S. export controls. When allies and partner countries adopted — or committed to adopting — export controls similar to the United States', the United States excluded certain exports to those countries from its export controls because those countries' export controls provide comparable protections. While the administration will continue with this approach, the Action Plan contains threats to use existing tools, such as the Foreign Direct Product Rule or secondary tariffs, against allies and partners that resist alignment.

In addition, the Action Plan calls for promoting plurilateral controls for the AI tech stack and avoiding relying solely on multilateral treaty bodies.

Promoting Global Leadership on AI Policy in International Governance Bodies

The administration seeks a greater U.S. role in shaping global AI policy and countering Chinese influence in international governance bodies. The Action Plan directs the DOC and DOS to leverage the U.S. position in the United Nations, Organization for Economic Co-operation and Development, G7, G20, and other international bodies to encourage standard-setting that "promote[s] innovation, reflect[s] American values, and counter[s] authoritarian influence" (and, in particular, influence from China). China, meanwhile, proposed a new Shanghai-based global AI cooperation organization and global AI governance on July 26, 2025.

Addressing National Security Risks Presented by AI

The Action Plan addresses national security threats posed by advanced AI systems, including the use of AI to carry out cyberattacks, or to develop chemical, biological, radiological, nuclear, or explosives weapons. Among other requirements, the DOC is responsible for working with industry to understand and evaluate national security risks emerging from AI, collaborating with national security agencies to assess vulnerabilities in the U.S. infrastructure and economy, and identifying and mitigating the methods that foreign adversaries or other entities can use to exploit those vulnerabilities. Additionally, federally funded institutions engaged in DNA synthesis will have to implement enhanced screening of customers to prevent malicious actors from using AI-driven methods to synthesize harmful pathogens and other biomolecules.

Final Thoughts

The Trump administration's AI priorities are clear. From its first days in office, the administration has sought to accelerate AI development and to ensure that the United States remains at its forefront. The Action Plan and the Permitting and Export EOs continue those themes. Preventing Woke AI continues a different theme that the administration also established at its outset.

AI developers, deployers, infrastructure providers, and their investors will have to determine how the administration's priorities match their own and how the administration's changes affect their global compliance obligations. For further analysis, ongoing updates, or strategic guidance on your company's specific opportunities and challenges, please contact the authors or other members of our multidisciplinary AI team.

This Advisory is intended to be a general summary of the law and does not constitute legal advice. You should consult with counsel to determine applicable legal requirements in a specific fact situation.