

AI Regulation under the Trump Administration

A Changing Business Environment for AI

Yusuke Inoue, Marubeni America Corporation Washington Office General Manager
Inoue-Y@marubeni.com

- The Trump administration has so far been cautious about tightening AI regulation. However, the release of Anthropic's latest AI model, Mythos, has intensified the regulatory debate.
- Public opposition to AI is rising rapidly, particularly among younger Americans. The term "anti-AI" encompasses a wide range of concerns, including job losses and data center development. Although these concerns have yet to coalesce into a unified political movement, they could eventually become impossible for policymakers to ignore.
- AI companies are seeking to shape the policymaking process both from within and outside the Trump administration. The industry is not acting as a unified bloc, however, as individual companies are pursuing regulatory and competitive conditions that favor their own interests.
- Regulatory uncertainty, growing public opposition, and the emergence of low-cost open-weight models are beginning to reshape the business environment for frontier AI.

Since taking office, the Trump administration has consistently pursued a deregulatory approach to AI policy. However, the remarkable advances embodied in the latest models, together with growing national security concerns, have brought regulation back to the forefront. At the same time, public opposition to AI is intensifying, driven by concerns over job losses and the rapid expansion of data centers. AI companies, for their part, are increasing their political influence, but fierce competition has prevented the industry from presenting a united front. This report reviews these developments and considers how the business environment surrounding AI is changing..

AI Regulation Becomes an Urgent Priority

The Trump administration has sought to promote economic growth by unleashing the private sector. Viewing innovation as a primary engine of growth, it has pursued broad deregulation. In AI, the administration released its AI Action Plan in July 2025, prioritizing faster innovation and the rapid expansion of infrastructure. Its objectives include accelerating data center construction, removing excessive safety regulations, securing energy supplies, and granting priority access to grid interconnections.¹ This policy support has helped spur massive investment in data centers by AI companies. In the near term, the resulting infrastructure boom is expected to boost economic growth. Over the medium to long term, productivity gains could strengthen competitiveness across the economy. AI is also regarded as the single most important technology in determining the outcome of the geopolitical competition with China, providing a national-security rationale for both deregulation and government support for the industry.

The release of Anthropic's latest AI model, Mythos, in April has nevertheless brought regulation back to the forefront. The model reportedly possesses extraordinary cyberattack capabilities and can identify system vulnerabilities in a short period of time. If misused, it could cause serious social disruption, including the shutdown of critical infrastructure and the theft of sensitive information. Anthropic therefore initially made Mythos available through Project Glasswing to only about 50 organizations, including major technology companies, financial institutions, and cybersecurity firms,

¹ White House, "Winning the Race: America's AI Action Plan," July 2025.

in an effort to identify vulnerabilities early and ensure the model's safety.² At the same time, Chief of Staff Susie Wiles, Treasury Secretary Scott Bessent, and National Cyber Director Sean Cairncross reportedly began leading an administration review of how frontier models should be regulated.

On June 2, the administration issued Executive Order 14409, "Promoting Advanced Artificial Intelligence Innovation and Security," requiring measures to ensure AI model safety and strengthen cyber defense.³ For the most advanced models, the order asks developers to submit them for government safety testing at least 30 days before public release. Submission is voluntary, however, and the order contains neither legally binding requirements nor penalties for noncompliance. The order also emerged only after considerable internal debate. An earlier draft would have required companies to submit models 90 days before public release and undergo mandatory government review. Shortly before signing, however, President Trump sought the views of David Sacks, co-chair of the President's Council of Advisors on Science and Technology. The proposal was subsequently watered down from mandatory review 90 days in advance to voluntary submission 30 days before release.

On June 9, Anthropic released Fable 5, a publicly available model based on Mythos. Fable 5 was designed with stronger guardrails to prevent misuse. When users asked questions involving cyberattacks or biological weapons, the system was supposed to switch automatically to a safer, previous-generation model. After researchers discovered a "jailbreak" method that could circumvent these safeguards, however, the Department of Commerce issued an export-control directive three days after the model's release, prohibiting its use by non-U.S. nationals. Anthropic continued negotiations with the government, and after the model was confirmed to meet stringent safety standards, public access was restored on July 1 following a 19-day suspension.⁴

Anthropic has clashed repeatedly with the government in the past. The company is unusual within the industry in that it has warned of catastrophic AI risks and advocated stronger government regulation. It also entered into a contract to deploy its models on classified Department of Defense networks, while seeking assurances that they would not be used for autonomous weapons or mass surveillance. As a result, Anthropic was designated a national-security supply-chain risk in March 2026 and excluded from government procurement.

Government intervention in frontier models has not been limited to Anthropic. Following the problems surrounding Fable 5, OpenAI made its latest model, GPT-5.6, available only to selected partners beginning June 26, before releasing it to the general public on July 9. Despite the voluntary framework set out in the June 2 executive order, government pre-release review of frontier models is therefore approaching a de facto requirement. The risks under discussion also appear increasingly real. On July 21, OpenAI disclosed that a high-performance model under testing had autonomously left its test environment and penetrated the systems of an outside company.⁵

Growing Public Backlash

Public opposition to AI is also beginning to influence the policy debate. A Pew Research survey conducted last year found that optimism about AI has declined since around 2023, with pessimistic

² Anthropic, Claude Mythos System Card and Project Glasswing materials, April 2026.

³ Executive Order 14409, "Promoting Advanced Artificial Intelligence Innovation and Security," June 2, 2026.


⁴ Reuters, Axios, and other reporting dated June 9, June 12, and July 1, 2026.

⁵ OpenAI announcement, July 21, 2026.

views now held by a majority of respondents.⁶ A March 2026 NBC poll similarly found that only 26% of respondents viewed AI positively, compared with 46% who held a negative view, a net rating of minus 20 percentage points. Opposition was especially pronounced among those aged 18–34, who are generally considered heavy users of AI, with a net rating of minus 44 points.⁷ Anti-AI sentiment is also notable for cutting across party lines.

A poll conducted by the AI Policy Institute following the Anthropic incident, released on June 29, produced similar results. Seventy-eight percent of voters preferred mandatory safety standards to an unregulated environment, while 68% supported a formal government review process for AI models. In addition, 86% said the most powerful models should include an “emergency kill switch” that could be used to shut them down. The findings suggest that public anxiety about AI currently outweighs expectations of its benefits.⁸

Figure: Views on the Increasing Use of AI in Daily Life



Source: Pew Research, June 2026

Although these surveys point to broad agreement among voters in opposing AI, the meaning of “anti-AI” varies considerably. Among younger people, resistance is often linked directly to anxiety about future employment. The primary concern is that AI will eliminate jobs, particularly entry-level white-collar positions.

The extent to which these fears have already materialized remains open to debate. Corporate hiring of new graduates has reportedly declined in recent years, however, and data from the Federal Reserve Bank of New York show that unemployment among recent college graduates aged 22–27 is rising.⁹ Fewer employment opportunities also reduce opportunities to acquire professional skills, potentially weakening individuals’ long-term value in the labor market. For graduates who took on substantial

⁶ Pew Research Center, survey on public attitudes toward AI, September 17, 2025.

⁷ NBC News poll, March 2026.

⁸ AI Policy Institute poll, released June 29, 2026.

⁹ Federal Reserve Bank of New York, employment data for college graduates aged 22–27, May 5, 2026.

student debt to earn their degrees, this can generate a profound sense that the promises made by politics and society have been broken.

Some argue that AI will ultimately create new jobs. Even so, a mismatch between the jobs created and those displaced is unavoidable, while acquiring new knowledge and skills entails costs. A period of disruption is therefore likely.

Beyond employment, data centers have emerged as another major focus of anti-AI sentiment. They were initially expected to benefit local economies through large-scale direct investment and construction demand, stable tax revenues, and the creation of new industries and infrastructure. Attention has increasingly shifted, however, to their negative effects. Data centers consume large volumes of water and electricity, impose substantial environmental burdens, and can raise living costs for local residents.

As a result, at least 75 data center projects were canceled or postponed in the first quarter of 2026 alone.¹⁰ On July 14, New York became the first state to impose a data center moratorium, banning for one year the construction of large facilities consuming 50 megawatts or more. Opposition to data centers has also become increasingly organized, with support from environmental groups such as Greenpeace and the Sierra Club. It can no longer be dismissed as a collection of isolated grassroots campaigns confined to individual communities.

The movement is also spreading beyond the political left. On July 18, anti-data-center demonstrations were held at 142 locations across 42 states. The organizer, Humans First, traces its roots to the Tea Party movement and the MAGA wing of the Republican Party.¹¹ Its activities extend beyond data centers to a broader campaign warning of the risks posed by AI itself.

The concerns grouped under the “anti-AI” label extend well beyond employment and data centers. They include control over personal data and privacy violations; the spread of disinformation and the resulting threat to democracy; child safety and education; the erosion of traditional social systems, including local communities; and changes in ethical standards and social values. The movement may therefore be understood not simply as a rejection of technology, but also as criticism of the lack of social and institutional preparation needed to adapt to it.

Because the term remains loosely defined and encompasses such a wide range of issues, the anti-AI movement currently lacks cohesion. Nevertheless, it has a strong possibility of emerging as a central theme in U.S. politics.

AI Companies Step Up Political Engagement

AI companies are facing mounting headwinds from government regulation and public opposition, but they are not standing on the sidelines. Leading Silicon Valley executives began moving closer to the Trump campaign during the 2024 presidential election and provided substantial financial support to its campaign efforts. Their engagement reflected expectations that Trump would continue traditional Republican policies favoring deregulation, economic growth, large-scale tax cuts, and innovation.

Following Trump’s election, Sacks and numerous other industry figures joined the administration and became involved in policymaking on AI, semiconductor exports, and digital assets. Technology executives have also maintained direct channels of communication with the administration by

¹⁰ Data Center Watch, data on data center project developments in the first quarter of 2026.

¹¹ Reuters, “Data center opponents stage 142 protests across 42 US states,” July 18, 2026.

attending White House meetings and accompanying the president on overseas visits. These efforts have helped the industry secure measures including relaxed environmental regulation, priority grid interconnection for power-generation facilities, eased controls on semiconductor exports to China, and public-private support for exports of “American-made AI.”

Leading AI companies increasingly appear to accept that some tightening of regulation is unavoidable. Their focus has therefore shifted toward shaping the rules in ways that favor their own competitive positions. Executives from OpenAI, Anthropic, and Google have recently advanced separate regulatory models.

OpenAI has opposed a framework controlled unilaterally by the U.S. government and instead advocated an international structure modeled on the International Atomic Energy Agency. Anthropic has called for more extensive government oversight and regulation, proposing a Federal Aviation Administration-style framework with clear standards to mitigate risks such as cyberattacks. On July 14, Google DeepMind CEO Demis Hassabis proposed industry self-regulation modeled on the Financial Industry Regulatory Authority, the self-regulatory body for the U.S. securities industry.

All three proposals emphasize safety concerns, but they differ over who should lead the regulatory process. Each therefore reflects the strategic interests of the company advancing it. At this stage, however, the proposals appear less intended as blueprints for immediate legislation than as efforts to shape the ground rules for future regulatory negotiations.

AI companies are also intensifying their political involvement ahead of the 2026 midterm elections. According to the Federal Election Commission, super PACs funded by the cryptocurrency and AI industries have already raised more than \$320 million—nearly twice the approximately \$170 million raised during the 2024 election cycle.¹² Differences within the industry over the appropriate direction of regulation have nevertheless led to the creation of several competing organizations.

Leading the Future is among the most prominent opponents of regulation. Established in August 2025 with funding from OpenAI co-founder Greg Brockman and prominent investors Marc Andreessen and Ben Horowitz, the group supports pro-innovation candidates from both parties. Its principal objective is to establish federal preemption of state-level AI regulation.

The pro-regulation group Public First Action was established as a counterweight to Leading the Future. Former Democratic Representative Brad Carson and former Republican Representative Chris Stewart founded the organization in November 2025. Anthropic contributed \$20 million in February 2026 and an additional \$20 million on July 21, bringing its total support to \$40 million.¹³ The group advocates stronger AI safety, transparency, and oversight, and opposes aggressive deregulation.

Meta, meanwhile, supports an organization active in state-level elections that promotes open-source AI development, faster data center construction, and opposition to state regulation. More recently, OpenAI engineers and former employees established the Guardrails Alliance, which opposes excessive deregulation and calls for stronger worker protections and safety-focused development. This has created a new political divide within individual companies, pitting management against employees.

¹² Federal Election Commission data and related reporting, 2026.

¹³ Reuters, Axios, The Hill, and other reporting, July 22, 2026.

Congressional Developments

Congress has also begun debating AI regulation. In the House, attention has focused on the Great American AI Act (GAAIA), a proposed comprehensive AI governance framework released on June 4 by Republican Representative Jay Obernolte and Democratic Representative Lori Trahan. Although the measure remains in draft form, it would require AI companies to establish formal safety-management policies and would grant the Commerce Department's Center for AI Standards and Innovation (CAISI) supervisory authority, including responsibility for the technical evaluation of frontier models.¹⁴

In the Senate, Democratic Senator Mark Warner has introduced his own comprehensive AI legislative package. The proposal would require frontier AI companies to submit transparency reports, undergo mandatory model reviews by an independent body, and disclose the environmental impact of their data centers.¹⁵ Both initiatives remain at an early stage of deliberation, however, and neither has a clear path to enactment.

Congress has also considered measures addressing data centers. In March, independent Senator Bernie Sanders and Democratic Representative Alexandria Ocasio-Cortez jointly proposed legislation that would temporarily halt the construction of new data centers and the expansion of existing facilities. The moratorium bill has little prospect of passage.¹⁶

The Ratepayer Protection Act, which is intended to require data center operators to bear the costs associated with constructing and operating their facilities, passed the House Energy and Commerce Committee unanimously on July 21. A vote by the full House before the congressional recess nevertheless appears difficult, and its prospects in the Senate are also uncertain. The bill would also leave final authority over electricity-rate setting with state public utility commissions.¹⁷

Risks in an Uncertain Regulatory Environment

Congress's inability to move decisively on AI regulation has made administration-led action increasingly prominent. At the same time, divisions within the industry make it difficult for the government to establish clear rules. Further complicating the picture is the view that some AI companies may actually prefer the current uncertainty. A lack of regulatory predictability generally weighs on the industry as a whole, but it can benefit leading companies with the resources to manage complex compliance requirements, effectively creating barriers to entry for competitors. Critics have characterized this approach as a strategy of "regulatory capture."

A similar contest over government regulation is emerging around open-weight models. On July 21, Treasury Secretary Bessent raised the possibility of sanctions against Chinese AI models on the grounds that distillation may involve intellectual-property infringement. The United States and China could eventually move to restrict access to each other's open-weight models. Given the Trump administration's more conciliatory posture toward China, however, many observers expect that any U.S. sanctions would come only after the U.S.-China summit scheduled for September.

¹⁴ Discussion draft of the Great American AI Act released by Representatives Obernolte and Trahan, June 4, 2026.

¹⁵ Axios, "Inside Sen. Mark Warner's AI plan," July 21, 2026.

¹⁶ S.4214, introduced by Senator Sanders, and the House version introduced by Representative Ocasio-Cortez, AI Data Center Moratorium Act of 2026.

¹⁷ H.R.9340, introduced by Representatives Evans and Custer, and S.5028, introduced by Senator Husted, Ratepayer Protection Act.

On July 22, the Little Tech Association, an industry group established this month to represent startups against Big Tech, sent a letter to Commerce Secretary Howard Lutnick and Office of Science and Technology Policy Director Michael Kratsios opposing restrictions on access to Chinese open-weight models. The group argued that depriving startups of inexpensive Chinese models would drive emerging companies out of the market and ultimately benefit the largest technology firms.¹⁸

On July 24, NVIDIA CEO Jensen Huang used his first-ever post on X to release a joint statement titled “Open Weights and American AI Leadership,” signed by 25 companies including Microsoft, Meta, and Palantir. The statement argued that open-weight models strengthen U.S. competitiveness and national security and opposed excessive regulation. OpenAI and Anthropic were notably absent from the list of signatories.¹⁹

Debate over whether AI is in a bubble also continues. Even so, regulatory uncertainty, shifting public attitudes, and the improving performance of open-weight models are clearly beginning to alter the business environment surrounding AI. The focus of the debate is changing rapidly, but AI will undoubtedly remain a major issue in Washington.

The implications for the widely shared expectation of strong U.S. economic growth led by AI will also warrant close attention. Three upcoming milestones will be particularly important: earnings releases by major technology companies beginning in late July, especially NVIDIA’s results on August 26; the U.S.-China summit in September; and the midterm elections in November.

Marubeni America Corporation Washington Office

1717 Pennsylvania Ave., N.W., Suite 375, Washington, DC 20006
<https://www.marubeni.com/jp/research/>

(Disclaimer)

- This document is created based on publicly available information; however, we do not guarantee its validity, adequacy, or completeness.
- Any advantages or disadvantages resulting from decisions made based on this document are the responsibility of the individual who made the decisions, and we bear no responsibility for them.
- The contents of this document are subject to change without notice.
- The individual texts, photographs, illustrations, etc., included in this document (hereafter referred to as "Information") are copyrighted works of our company, protected under the Copyright Law of Japan and international treaties such as the Berne Convention. Except in cases permitted by copyright law, such as personal private use and quotation, reproducing, distributing, modifying, translating, adapting, broadcasting, or making available to the public the Information contained in this document without the permission of the copyright holder is a violation of copyright law.

¹⁸ Letter from the Little Tech Association opposing restrictions on open-weight AI models, July 22, 2026.

¹⁹ “Open Weights and American AI Leadership,” July 24, 2026.