Q1 2026

WILSON SONSINI

AI/ML
ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING

ALL EYES ON AI: REGULATORY, LITIGATION, AND TRANSACTIONAL DEVELOPMENTS

The White House Releases Federal Artificial Intelligence Framework

On December 11, 2025, President Trump issued an Executive Order (EO) to protect American artificial intelligence (AI) innovation from "the most onerous and excessive laws emerging from the States that threaten to stymie innovation." This EO was signed in response to a flurry of AI-related bills being introduced across every state in America.

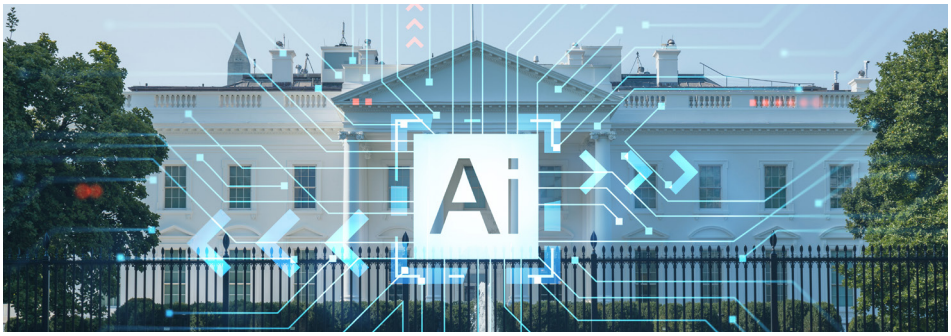

In This Issue

To further the Trump administration's goal of creating a uniform national AI policy framework, the White House announced its National Policy Framework for Artificial Intelligence (the Framework) on March 20, 2026. Through this Framework, the Trump administration aims to secure American leadership in AI to advance economic competitiveness, strengthen national security, and ensure broad societal benefit for Americans. Recognizing public concerns about AI's impact on issues such as children's online safety and household utilities costs, the Framework addresses seven core objectives:

1. **Protecting Children and Empowering Parents:** The Framework calls on Congress to create several protections, including providing parents with tools to manage their children's digital experiences such as privacy controls and safeguards against exploitation and self-harm.

2. **Safeguarding and Strengthening American Communities:** The Trump administration believes that residential ratepayers should not experience increased electricity costs as a result of AI data center construction and operation. As such, the Framework urges Congress to streamline permitting processes so that data centers can generate power on site, thereby enhancing grid reliability. The Framework also asks Congress to enhance the federal government's ability to combat AI-enabled scams and address AI-related national security concerns.

3. **Respecting Intellectual Property Rights and Supporting Creators:** The Framework calls on Congress to balance protections for American innovators, creators, and publishers with the need for AI systems to learn from available information.

4. **Preventing Censorship and Protecting Free Speech:** The Framework reflects the Trump administration's concern with safeguarding freedom of speech by proposing guardrails to ensure that AI "can pursue truth and accuracy without limitation."

Continued on page 2...

The White House Releases Federal AI Framework *(Continued from page 1)*

The Framework aims to prevent AI systems from being used to silence or censor lawful political expression or dissent.

5. **Enabling Innovation and Ensuring American AI Dominance:** The Framework asks Congress to remove unnecessary regulatory barriers, accelerate AI deployment across sectors, and expand access to testing environments needed to build and deploy world-class AI systems.

6. **Educating Americans and Developing an AI-Ready Workforce:** The Framework also encourages Congress to expand workforce training and education to ensure Americans can participate in and benefit from AI-driven economic growth.

7. **Creating a Uniform Federal Framework to Preempt State AI Laws:** Lastly, the Framework calls upon Congress to create one minimally burdensome national standard that would respect the key principles of federalism but would enable the administration to advance its national strategy of achieving global AI dominance. The Framework notes that states should not be permitted to regulate AI development "because it is an inherently interstate phenomenon with key foreign policy and national security implications."

The Framework is intended to be applied uniformly by the federal government, and the White House notes that a "patchwork of conflicting state laws would undermine American innovation and [the United States'] ability to lead in the global AI race." It is the Trump administration's goal to work with Congress to turn this framework into legislation. Just days before the Framework was issued, Senator Marsha Blackburn released a draft of proposed legislation that would replace the existing patchwork of state AI regulations. Similar to the Framework, this proposed bill seeks to promote children's online safety, prevent censorship and promote freedom of speech, and protect the intellectual property of American innovators and creators. While it remains to be seen whether Senator Blackburn's proposed legislation will become law, it is clear that the Trump administration is keen to work with Congress to enact federal AI legislation that would preempt existing state laws in this sector and advance the administration's objectives surrounding AI.

For more on recent AI EO developments, see this client alert.

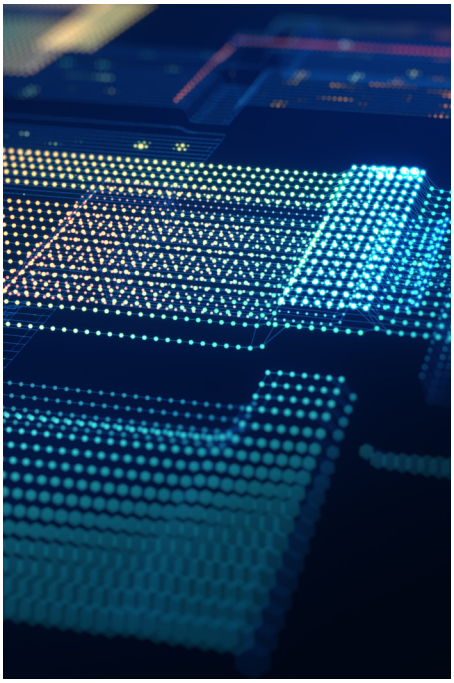

Trend Toward AI Chatbot Safety Laws

State regulation of chatbots is accelerating, driven by concerns that AI systems may contribute to harms to minors, including mental health risks. Nearly 100 chatbot-related bills have been introduced in approximately two dozen states this year. Several states, including California, Idaho, Maine, New Hampshire, New York, Oregon, Utah, and Washington, have already enacted chatbot-related measures. These proposals vary significantly in scope and terminology, including how they define "chatbot" or "artificial intelligence," creating a fragmented landscape in which similar systems may be subject to different requirements across jurisdictions.

Although there is variation in scope, a defining trend is emerging around transparency, content safety, and harm prevention. Many of the enacted laws and proposed bills contain provisions that contain requirements for harm detection and response. Common requirements include disclosures that users are interacting with AI, safeguards for high-risk scenarios (such as self-harm-related interactions), and restrictions on harmful outputs, including sexually explicit content involving minors. For example, California's SB 243 requires operators of "companion" chatbot platforms to implement various safety measures for

Continued on page 3...

Trend Toward AI Chatbot Safety Laws *(Continued from page 2)*

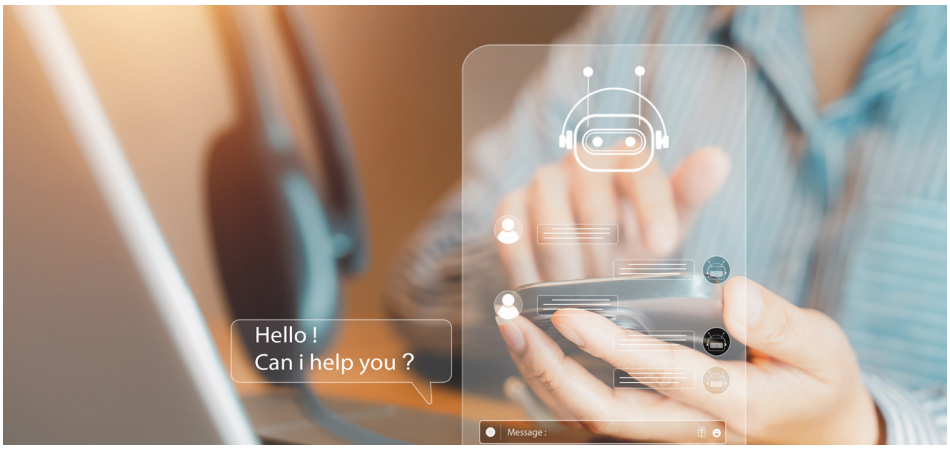

users known to be minors and to develop and publish a protocol to prevent content related to suicide and self-harm, among other requirements. Some measures go further by limiting the use of chatbots in sensitive contexts, such as substitutes for licensed professionals.

Another emerging trend is the focus on protections for minors. Various chatbot proposals include age-verification requirements for chatbot platforms, but they vary in approach. For example, Florida's SB 1344 requires chatbot platforms to use age verification to identify minors and then take specific actions, including requiring the account to be affiliated with a parent account, obtaining parental consent, and blocking access to sexually explicit content.

Other proposals, including South Dakota's SB 168, Oklahoma's HB 4083 and Maine's LD 2162, prohibit minors from communicating with chatbots with "human-like" features, some permitting communication if the "human-like" features are removed.

Many of these laws and proposals authorize enforcement by state attorneys general, while some, such as California's SB 243 and New Hampshire's HB 143, include private rights of action. These developments point to a rapidly evolving and fragmented regulatory landscape in which companies deploying consumer-facing or companion chatbots should expect heightened scrutiny of product design, user interactions, and safety controls across jurisdictions.

For more on recent AI regulatory developments, see this client alert.

Implications of SDNY Ruling for Discovery of AI Communications

On February 10, 2026, Judge Rakoff of the Southern District of New York (SDNY) considered in *U.S. v. Heppner* whether, when a user communicates with a publicly available AI tool in connection with a pending criminal investigation, the user's communications are protected by attorney-client privilege or the work product doctrine. The court concluded that the user's communications were not privileged or protected work product and issued a written memorandum explaining this decision on February 17, 2026. The court's order is summarized below:

In the criminal case, after receiving a grand jury subpoena, the defendant used a publicly-available AI tool to prepare dozens of documents related to his case. According to the order, these documents were not generated at the direction of counsel, but the defendant subsequently shared these documents with his attorneys. The government moved for a determination that the documents were not privileged.

Attorney-Client Privilege: The court concluded that the communications did not meet the three-part test for attorney-client privilege: communications (1) between a client and their attorney (2) that are intended to be, and in fact were, kept confidential (3) for the purpose of obtaining or providing legal advice. The court found that the defendant failed to

Continued on page 4...

Implications of SDNY Ruling for Discovery of AI Communications *(Continued from page 3)*

meet at least the first two elements (if not all three). First, the court found the communications were not between the defendant and their attorney, since the AI tool was not an attorney. Second, the court found the communications were not confidential because the defendant communicated with a third-party AI platform, and the AI platform's privacy policy stated the company may use certain data to train its model and disclose that data to third parties, including "governmental regulatory authorities." The court also concluded that "AI users do not have substantial privacy interests in their 'conversations with [another publicly accessible AI platform] which users voluntarily disclosed' to the platform." Finally, the court indicated that the defendant did not communicate with the AI platform for the purpose of obtaining legal advice. The court acknowledged that this issue "perhaps presents a closer call," but the defendant conceded that he did not do so at the suggestion or direction of counsel.

Work Product Doctrine: The court also concluded that the documents were also not protected by the work product doctrine, which the court noted provides protection for "materials prepared by or at the behest of counsel in anticipation of litigation or for trial." The court noted that even if the documents were prepared in anticipation of litigation, they were not prepared at the direction of counsel and did not reflect defense counsel's strategy.

The case raises novel issues regarding privilege as applied to the use of AI in litigation. Other courts are grappling with these issues as well, and the law is still developing.

The Meta-Manus Ruling Highlights Growing Global AI Deal Risk

On April 27, 2026, China's National Development and Reform Commission (NDRC) ordered Meta to unwind its $2 billion acquisition of Manus, an AI-agent start-up. This decision is a vivid reminder that foreign investment review regimes well beyond the U.S. Committee on Foreign Investment in the United States (CFIUS) are viewing acquisitions of AI-related businesses as a strategic national security matter. Governments around the world are actively scrutinizing AI-related transactions, whether through inbound foreign investment screening regimes or outbound investment restrictions that limit where companies can deploy capital, and those operating in this space should be tracking developments across multiple jurisdictions when preparing for equity fundraising, M&A, or other corporate transactions.

The announcement came after four months of regulatory scrutiny by multiple Chinese authorities and, as the ruling was issued post-closing, Meta had already begun the process of onboarding Manus personnel and integrating the company into the Meta group. One of the most interesting features of this decision was that Manus had relocated its headquarters from China to Singapore in 2025, and yet the NDRC asserted jurisdiction based on the technology's Chinese origin, the nationality of the founders, and the company's historical ties to Chinese data. This extraterritorial reach is not unique to China. Regulators in the EU, UK, Australia, and elsewhere have similarly expanded their foreign investment review frameworks to capture AI-driven deals with only limited connections to those jurisdictions. Companies, funds, investors, and other market participants engaged in the AI space should carefully assess the investment regimes that may be implicated and remain vigilant to evolving regulatory obligations, potential national security risks, and the substantial impact these reviews can have on both deal certainty and closing timelines.

EU Legislators Agree on EU AI Act Amendments with Implications for AI Companies Operating in the EU

On May 7, 2026, EU legislators reached a political agreement on amendments to the EU AI Act. The key changes are:

- **Delayed application of the AI Act requirements for high-risk AI systems (HRAI)** from August 2, 2026, to December 2027. For HRAI systems covered by EU product safety legislation, the deadline is postponed to August 2028.
- **New industrial AI carve out** meaning industrial AI embedded in machinery will be subject to AI-related obligations under EU machinery rules, rather than the AI Act's HRAI regime.
- **Deadline to mark AI-generated content extended** from August 2, 2026, to December 2, 2026. The European Commission is expected to issue a transparency code of practice before the new deadline.
- **New prohibitions on the use of generative AI to create certain harmful content**, including (i) AI systems to generate non-consensual sexualized deepfakes (e.g., nudification apps), and (ii) AI systems to generate child sexual abuse material.

For further details of the AI Act amendments, see here.

The formal legal amendments are expected to take effect before the original August 2, 2026, deadline for the HRAI regime. Meanwhile, EU legislators continue to negotiate revisions to other EU digital regulations, notably the General Data Protection Regulation (GDPR), which may further impact AI companies operating in the EU. For more information on the proposed changes to other EU digital regulations, see here.

European Regulators Are Scrutinizing AI Tools That Can Generate Sexually Explicit Images

Regulators across Europe are scrutinizing X and its AI tool, Grok, for generating sexualized images. On January 26, 2026, the European Commission announced that it launched a formal investigation into X under the EU Digital Services Act in relation to the dissemination of AI-manipulated sexually explicit images through Grok. In addition, on February 17, 2026, the Irish privacy regulator announced it had opened an inquiry into X concerning the creation and publication of non-consensual intimate and/or sexualized images involving processing personal data of EU individuals, including children, using a generative AI functionality associated with Grok within the X platform.

Similarly, on January 12, 2026, the UK online safety regulator, Ofcom, announced it had opened a formal investigation into whether X had done enough under the Online Safety Act (OSA) to assess and mitigate the risk of sexual deepfakes, including of children, spreading on its social media platform, and to take them down quickly once identified. As the OSA does not currently apply to AI chatbots, it was unable to investigate the creation of illegal images by the standalone Grok service. However, on February 3, 2026, the UK privacy regulator also announced that it is investigating X and xAI's processing of personal data of UK individuals by Grok and its potential to produce harmful sexualized image and video content.

All regulatory investigations are ongoing. In parallel, EU legislators have agreed to amend the EU AI Act to prohibit AI systems to (i) generate non-consensual sexualized deepfakes (e.g., nudification apps) and (ii) generate child sexual abuse material, and it is now a criminal offense in the UK to create or facilitate the creation of intimate images of a person, including generating an intimate image using AI.

UK Government Takes Steps Toward More AI Regulation Measures in the UK

To date, the UK has adopted a principles-based approach that relies on regulatory cooperation to regulate AI. It has been reported that the UK government plans to legislate to regulate the highest risk AI models, but a draft bill has been delayed and, to date, no proposal has been brought forward. However, the UK government has recently taken steps towards regulating AI in targeted areas, affecting providers and deployers of AI in the UK.

On April 30, 2026, the UK government was granted a new power to bring generative AI services into the scope of the UK Online Safety Act (OSA). The government could exercise this power by introducing regulations setting out which duties under the OSA AI services must comply with (e.g., to minimize or mitigate the risk of harm caused by illegal AI-generated content or the use of the AI service to carry out a criminal offense). There is no timeline for such regulations to be introduced. Further information on the duties that AI services could become subject to under the OSA is available here.

In addition, on May 12, 2026, regulations to require the UK data protection authority, the Information Commissioner's Office (ICO), to prepare a code of practice in relation to developing and using AI and automated decision-making under UK data protection law (Code) became law. Once published, the Code will not be legally binding. However, it will be admissible in evidence in legal proceedings, and a court, tribunal or the ICO must account for the Code when relevant to assessing the merits of a case or investigation started once the Code is in force. There is no timeline for the Code to enter into force, but the ICO must publish a draft before it is finalized.

Algorithmic Pricing Global Roundup: The UK's Hotel Investigation, a Canadian Consultation, and Concerns About a Chinese Hotel Pricing Tool

Competition authorities globally are showing a sustained interest in algorithmic pricing. However, cases and investigations are not equally developed in all jurisdictions, with European cases at most in the investigation phase.

On January 22, 2026, the Competition Bureau Canada (CBC) published the results of its consultation on algorithmic pricing and competition, following a discussion paper it had released in June 2025. Responses from individuals showed concern about the potential for unfairness and discriminatory practices, while longer responses addressed both potential benefits from economic efficiencies as well as possible anticompetitive behavior. The CBC emphasized that it was merely summarizing responses from the public to its consultation, and not providing conclusions of its own.

On February 24, 2026, the UK's Competition and Markets Authority (CMA) launched an investigation into whether the three hotel chains Hilton, IHG Hotels, and Marriott may have used the hotel data analytics tool STR, provided by CoStar, to share competitively sensitive information.

Continued on page 7...

Algorithmic Pricing Global Roundup . . . *(Continued from page 6)*

STR collects and shares anonymized and aggregated hotel performance data, such as occupancy, pricing, and revenue trends. The CMA is likely to focus on whether such data can be de-aggregated and how algorithms can use such data to anticipate and react to decisions of competitors.

On March 10, 2026, according to public sources, Ctrip, a Chinese subsidiary of online travel platform Trip.com, shut down its automated hotel pricing tool "AI Business Assistant." Some hotel partners had criticized the tool, claiming that it automatically scanned competitor prices and forced price reductions on their own listings.

UK CMA Publishes Analysis on the Impact of Agentic AI on Consumers; France and Turkey Launch AI Sector Inquiries

On March 9, 2026, the CMA published an analysis on the impact of agentic AI on consumers, outlining its concerns and expectations for businesses.

While recognizing that agentic AI can have beneficial outcomes for consumers including reduced friction, personalization of services, and lower prices, the CMA warned that increased autonomy brings new risks. In the CMA's view, it may enable manipulative behaviors, especially when designed to increase engagement or other commercial objectives. Agents can also make mistakes with significant consequences especially for financial or contractual decisions, and reinforce existing biases through opaque processes, preventing consumers from challenging decisions. Agentic AI may also increase the risk of coordinated market outcomes in a manner similar to what may result from anticompetitive algorithmic pricing.

The CMA emphasized that companies remain responsible for agents' behaviors and should ensure compliance with consumer law and competition law. They are also expected to be transparent about AI use, train and test AI systems properly, ensure human oversight, and monitor and correct harmful outcomes efficiently.

On January 9, 2026, the French Competition Authority (FCA) launched a sector inquiry into competition in the "conversational agent" or chatbot sector. This follows a 2024 FCA inquiry into the upstream part of the AI value chain focusing on foundation models, and a 2025 study into competition issues of energy and environmental impacts of AI. The FCA stated it was especially interested in the use of chatbots in the e-commerce sector, advertising, and partnerships, but that the inquiry would exclude the interplay of chatbots and search engines. The FCA's public consultation closed on March 6, 2026.

On April 7, 2026, the Turkish Competition Authority (TCA) announced its own comprehensive inquiry into the AI sector. The TCA stated it wanted to study how the AI ecosystem is being developed, including the main models, relationships between the different layers of the value chain, access to critical inputs, interactions between large technology undertakings and innovators also from a merger control perspective, and effects of data and computing power on competition. It is as of yet unknown on which timeline the TCA's sector inquiry will progress.

Select Client Highlights

- Wilson Sonsini Advises Thinking Machines Lab on Strategic Partnership with NVIDIA
- OpenAI Prevails in Landmark Italian AI and GDPR Enforcement Case
- Wilson Sonsini Advises Replit on $400 Million Series D
- Wilson Sonsini Advises ElevenLabs on $500 Million Series D
- Firm Advises Lead Investor G2 Venture Partners on Waabi's $1 Billion Fundraise

Newsletter Contributors

- Deirdre Carroll
- Laura De Boel
- Tom Evans
- Julius Giesen
- Joshua Gruenspecht
- Angela Guo
- Jordan Jaffe
- Kara Millard
- Maneesha Mithal
- Manja Sachet
- Taylor Stenberg Erb
- Michelle Ullman
- Hattie Watson
- Malcolm Yeary

The following attorneys have editorial oversight of Wilson Sonsini's All Eyes on AI: Regulatory, Litigation, and Transactional Developments.

Laura De Boel
ldeboel@wsgr.com

Jordan R. Jaffe
jjaffe@wsgr.com

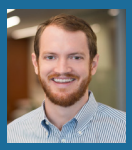

Scott McKinney
scmc@wsgr.com

Maneesha Mithal
mmithal@wsgr.com

Manja Sachet
msachet@wsgr.com

WILSON SONSINI

650 Page Mill Road, Palo Alto, California 94304-1050 | Phone 650-493-9300 | Fax 650-493-6811 | www.wsgr.com

Wilson Sonsini has 17 offices in technology and business hubs worldwide. For more information, visit wsgr.com/offices.