

Global AI Risk & Regulation

Researched by Hey Lefty

Automated research briefings on topics you choose — heyleftty.com

TL;DR

AI liability and risk management are fracturing along a stark federal-state divide. While the US federal government has pivoted toward voluntary, security-focused partnerships that reject mandatory licensing for advanced systems White House Executive Order, states like Connecticut are enacting highly prescriptive laws that strip employers of legal defenses for automated hiring bias Connecticut SB 5. This leaves enterprise risk teams navigating a landscape where national compliance remains collaborative, but local operational deployment is increasingly high-stakes.

The Federal Shift to Voluntary Cyber-Defense Partnerships

Federal AI policy is pivoting away from comprehensive mandatory regulations in favor of voluntary, security-focused collaboration between developers and national security agencies.

"Notably, the EO explicitly states that it does not authorize any new mandatory government licensing, pre-clearance, or permitting requirements for the development, release, or distribution of AI models." — White House Executive Order

"One immediate takeaway is that AI and cybersecurity are being treated as a combined governance priority." — White House Executive Order

By prioritizing voluntary 30-day pre-release access and benchmarking over hard federal bans, the administration is shifting the burden of risk management onto corporate security teams. As detailed by DLA Piper and Sidley Austin, enterprise counsel must now treat AI deployments not as a check-the-box regulatory hurdle, but as a core cybersecurity and national security vulnerability.

What to watch: How the National Security Agency and CISA define the technical thresholds for "covered frontier models" under their new classified benchmarking process Sidley Austin.

The State-Level Hardening of Workplace AI Liability

While federal oversight softens into voluntary frameworks, state legislatures are moving aggressively to impose strict, non-negotiable liability and disclosure mandates on employers using automated hiring tools.

"...any technology that processes personal data and uses computation to generate any output, including, but not limited to, any prediction, recommendation, classification, ranking, score or other information, that is a substantial factor used to make or materially influence an employment-related decision" — Connecticut SB 5

"SB 5 amends Connecticut's employment discrimination law to specify that the use of covered automated employment-related decision technology to make an employment decision is 'not a defense against a complaint alleging a discriminatory practice.'" — Connecticut SB 5

This statutory shift, highlighted in analyses by Ogletree Deakins and Holland & Knight, forces enterprise risk teams to take direct ownership of third-party algorithms, as they can no longer shift liability back to software vendors when discrimination complaints arise. Organizations must establish rigorous internal anti-bias testing protocols to serve as mitigating factors in state-level enforcement actions.

What to watch: How employers adapt their recruitment workflows before the law's strict disclosure and notice requirements take effect in October 2027 Holland & Knight.

What surprised us

- **The AI-Caused Layoff Disclosure.** Connecticut's new legislation quietly introduces a highly unusual requirement starting October 2026 Connecticut SB 5. Employers filing WARN Act notices for mass layoffs must explicitly state whether the job cuts are related to their use of AI or other technological changes Connecticut SB 5. This creates immediate reputational and PR risks for companies restructuring their workforces.
- **A Classified Process for "Covered" Systems.** The White House's new policy introduces a classified benchmarking process to evaluate AI systems White House Executive Order. This creates a strange scenario where developers must voluntarily submit systems for 30-day pre-release reviews without a publicly transparent framework for how the "covered frontier model" threshold is determined White House Executive Order.

Open threads worth a vote

- **Eightfold AI Motion to Dismiss Hearing on FCRA/ICRAA Claims** — Cast your vote to prioritize coverage of the upcoming federal court hearing on whether automated candidate-scoring tools constitute "consumer reporting" under credit protection laws.
- **EU Product Liability Directive (PLD) Transposition Deadline** — Vote to track how member states are translating software and AI liability rules into national laws as the transposition deadline approaches.

Appendix: Findings

White House Executive Order "Promoting Advanced Artificial Intelligence Innovation and Security" Establishes Voluntary Cybersecurity Framework

White House Executive Order "Promoting Advanced Artificial Intelligence Innovation and Security" Establishes Voluntary Cybersecurity Framework

On June 2, 2026, President Donald Trump signed an Executive Order (EO) titled "Promoting Advanced Artificial Intelligence Innovation and Security," signaling a major shift in federal AI policy. Moving away from mandatory comprehensive regulatory frameworks, the order prioritizes voluntary public-private collaboration, national security, and AI-enabled cyber defense. It explicitly rejects mandatory government licensing, pre-clearance, or permitting requirements for frontier models.

The EO focuses on three core pillars: upgrading federal and critical infrastructure cyber defenses, establishing a secure (but voluntary) frontier model deployment process, and prioritizing criminal enforcement against malicious actors using AI.

Voluntary Frontier Model Benchmarking and Pre-Release Access

The order directs the Secretaries of the Treasury, War, and Homeland Security, through the National Security Agency (NSA) and Cybersecurity and Infrastructure Security Agency (CISA), to develop:

1. A **classified benchmarking process** to evaluate the advanced cyber capabilities of AI models and determine the threshold for designation as a "covered frontier model."
2. A **voluntary framework** under which developers can collaborate with the government to assess whether their models meet the threshold, grant the government 30-day pre-release access to covered frontier models for security testing, and select trusted partners for early access.

Crucially, the order emphasizes its non-regulatory, incentive-aligned posture:

"Notably, the EO explicitly states that it does not authorize any new mandatory government licensing, pre-clearance, or permitting requirements for the development, release, or distribution of AI models." — DLA Piper

AI and Cybersecurity Convergence

The EO establishes an **AI cybersecurity clearinghouse** to coordinate and deconflict scanning for software vulnerabilities, validate vulnerabilities, and prioritize remediation and patch distribution. This clearinghouse will operate as a "voluntary collaboration with the AI industry and operators of

critical infrastructure." It also directs the Office of Management and Budget (OMB) to assess federal grant funding that can be directed toward developers of advanced AI vulnerability detection.

For in-house counsel, these provisions mean that AI governance and cybersecurity are now unified priorities:

"One immediate takeaway is that AI and cybersecurity are being treated as a combined governance priority. Boards, regulators, and counterparties are increasingly likely to ask how a company evaluates AI-enabled cyber risk, how it secures and monitors advanced AI deployments, how it manages its government and critical infrastructure relationships, and how its leadership oversees the intersection of AI innovation, cybersecurity, and national security." — Sidley Austin

Prioritized Criminal Enforcement

Finally, the EO instructs the Attorney General to prioritize the enforcement of federal criminal laws (such as the Computer Fraud and Abuse Act, identity theft, and wire fraud) against those who utilize AI to illegally access or damage computers, or use "AI agents to unlawfully access data or information that is subsequently used for a criminal or unlawful purpose." This places a higher premium on enterprise incident response plans that can detect and preserve evidence of automated or agentic AI intrusions.

Sources

- Promoting Advanced AI Innovation and Security Executive Order: Top points
- Cyber Strategy at the AI Frontier: President Trump Releases Executive Order to Promote Advanced Artificial Intelligence Innovation and Security

Connecticut Enacts SB 5: Comprehensive Workplace AI Regulation and Novel AI-Caused RIF Disclosures

Connecticut Enacts SB 5: Comprehensive Workplace AI Regulation and Novel AI-Caused RIF Disclosures

On May 27, 2026, Connecticut Governor Ned Lamont signed Senate Bill (SB) 5 (designated as Public Act No. 26-15) into law, establishing a landmark bipartisan framework for artificial intelligence (AI) regulation. While the 74-page bill covers subscription-based AI, frontier model whistleblower protections, and AI companion safety, its most immediate and heavy-hitting impact on enterprise risk teams lies in its sweeping regulation of "automated employment-related decision technology" (AERDP).

Starting on a staggered basis (with key employment provisions taking effect on October 1, 2027), the law represents one of the nation's most stringent and prescriptive workplace AI frameworks,

alongside California's CCPA regulations and Colorado's rewritten AI Act ([[colorado-sb26-189-revised-ai-act-2026]]).

Scope and Definition of Workplace AI (AERDP)

The law applies to "automated employment-related decision technology" (AERDP), defined as:

"any technology that processes personal data and uses computation to generate any output, including, but not limited to, any prediction, recommendation, classification, ranking, score or other information, that is a substantial factor used to make or materially influence an employment-related decision" — Ogletree Deakins

An "employment-related decision" includes hiring, promotion, discipline, discharge, tenure, terms, privileges, or conditions of employment. Common software tools (spreadsheets, word processors, spellcheckers, map navigation) are excluded, as are scheduling and productivity monitoring tools, provided they do not make material employment decisions.

Strict Employer Notice and Disclosure Mandates

Starting October 1, 2027, employers in Connecticut must provide plain-language disclosures to employees and applicants interacting with AERDP. Additionally, before making any employment decision where the automated technology is a "substantial factor" (defined as assisting in making and capable of altering the outcome), the employer must provide a written notice disclosing:

- That the technology has been deployed and its trade name.
- The purpose and nature of the employment decision.
- The categories of personal data processed and how they are assessed.
- The sources of that personal data.
- Contact information for the employer.

If an adverse employment-related decision is made, the employer must provide the affected individual with a high-level explanation of the principal reasons, the specific role of the AERDP, the data used, and an opportunity to correct their personal data.

No Defense Against Discrimination Claims

Crucially, SB 5 amends Connecticut's employment discrimination law to explicitly strip employers of an easy defense:

"SB 5 amends Connecticut's employment discrimination law to specify that the use of covered automated employment-related decision technology to make an employment decision is 'not a defense against a complaint alleging a discriminatory practice.'" — Ogletree Deakins

However, proactive "anti-bias testing" or similar efforts can be considered by courts and the Connecticut Commission on Human Rights and Opportunities (CHRO) as a mitigating factor.

Novel AI-Related RIF Disclosures (Effective October 1, 2026)

In a first-of-its-kind statutory requirement, starting October 1, 2026, employers serving written notice to the Labor Department of a plant closing or mass layoff under the federal Worker Adjustment and Retraining Notification (WARN) Act must explicitly disclose:

"whether the reductions covered by the notice 'are related to the employer's use of artificial intelligence or another technological change.'" — Ogletree Deakins

Developer-Deployer Division of Labor

To support employer compliance, the law mandates that AI developers provide deployers (employers) with all information necessary to fulfill their disclosure obligations.

Whistleblower Protections for Frontier Developers

Starting January 1, 2027, "frontier developers" (entities developing "foundation models" capable of influencing physical or virtual environments) must establish anonymous internal reporting channels for public safety or catastrophic risks. Retaliation against whistleblowers is strictly prohibited, with civil penalties of up to \$1,000 per violation.

Enforcement

There is no private right of action for violations of the notice and disclosure provisions (which are enforced exclusively by the state attorney general as unfair or deceptive trade practices, subject to a 60-day cure period for violations occurring before December 31, 2027). However, traditional discrimination claims arising from AI hiring bias can still be pursued privately under existing employment laws.

Sources

- Connecticut Enacts Comprehensive AI Legislation: Key Obligations for Developers and Deployers
- New Connecticut Law Restricts Employer AI Use, Mandates Notice for AI-Caused RIFs