LUMO INSIGHTS

Thursday, 9 July 2026

Global crackdown and New threats put AI governance in spotlight

Executive Summary

In the past two days, a flurry of developments has underscored the urgent need for responsible AI governance. From the European Union launching enforcement of its landmark AI Act to a potentially game-changing U.S. federal AI bill and a precedent-setting court ruling on AI liability, regulators and courts worldwide are making clear that managing AI risk is now a core business issue. At the same time, the emergence of powerful "agentic" AI systems - and even an AI-driven cyberattack - shows that companies must strengthen oversight and safety measures to remain both compliant and competitive.

European Union: AI Act enforcement kicks off

The European Union's AI Act - the world's first comprehensive law regulating artificial intelligence - is moving from theory to reality. This month marks the start of enforcement for the Act's "high-risk" AI provisions, following a two-year transition period. Regulators have made it clear this is not a soft launch; the law carries teeth, with penalties for non-compliance that can reach up to 7% of a company's global annual revenue.

As of August 2, businesses deploying AI in sensitive sectors like employment, finance, critical infrastructure, law enforcement, healthcare, and education will be subject to strict requirements. They must maintain detailed technical documentation for high-risk AI systems, conduct rigorous risk assessments and audits, ensure meaningful human oversight of AI-driven decisions, and register qualifying applications in a new EU-wide database. For companies operating in or serving the EU market, these obligations are no longer theoretical. Organisations should be inventorying their AI systems now and implementing governance measures to ensure compliance.

Early enforcement actions have already begun to emerge. The European AI Office has reportedly sent information requests to multiple firms to assess their AI compliance. In one instance, several U.S. tech companies were asked to provide documentation for AI-driven hiring tools used in European markets. An EU-based bank's automated lending system is under scrutiny for lacking adequate human oversight, and a national health service's AI diagnostic platform was flagged for missing risk assessment documentation. While no fines have been issued yet, these investigations show regulators' enforcement priorities - transparency, human oversight, and robust documentation - and their willingness to act swiftly on potential violations.

On July 2, the European Commission also published final guidance that broadened the law's scope. Notably, it clarified that any company which customises or fine-tunes a general AI model for a specific use is considered a 'deployer' of an AI system (rather than just an end-user), and thus bears full responsibility for complying with the AI Act's requirements. In practice, this means firms adapting third-

party AI models (for example, tailoring a large language model for a customer service chatbot) must implement the same rigorous oversight, documentation, and risk controls as the original model developers.

Even participation in a government-sponsored "sandbox" offers no immunity from scrutiny. French authorities this week issued formal warnings to six companies operating under France's AI sandbox programme after auditors found they lacked required transparency documentation for their AI systems. The clear takeaway is that regulators expect compliance even from experimental AI projects. The EU's message to enterprises: if you deploy high-risk AI, be prepared to demonstrate accountability and safety measures now - not at some undefined future date.

United States: federal AI Bill vs. state regulations

In the United States, tensions are growing between federal ambitions for AI governance and a patchwork of state rules. In a significant step, the U.S. Senate on July 3 passed the Great American Artificial Intelligence Act by a 67-31 vote. This sweeping federal bill - the closest the country has come to a national AI law - would establish comprehensive AI oversight, including requiring federal agencies to conduct AI impact assessments for high-risk government systems and creating a national AI Safety Board under NIST.

However, the bill's most controversial feature is a provision to preempt (override) state AI laws that conflict with new federal standards. That clause has set up a showdown with states such as California and Texas, which have been pressing ahead with their own AI regulations on issues like automated hiring practices and AI safety disclosures. Several state officials, including attorneys general, argue the federal law's preemption language is too broad and could undermine stronger local protections, setting the stage for potential legal challenges.

As the bill moves to the House of Representatives for debate, its fate remains uncertain. In the meantime, companies must contend with an expanding thicket of state-level AI requirements. Just this week, the Texas legislature approved a new law mandating that employers disclose AI use in hiring and allow candidates a human review of algorithmic decisions. Similarly, New York's financial regulator (NYDFS) on July 1 issued binding model risk-management guidance for insurance companies using AI in underwriting and claims, with compliance required by January 2027. From financial services to HR, many states are implementing their own AI rules in the absence of a single national standard.

If the Great American AI Act becomes law with its preemption clause intact, some of these state measures may be nullified - likely prompting court battles over states' rights. If the federal effort stalls, however, businesses will continue to face a fragmented regulatory landscape, needing to comply with varying AI laws across jurisdictions. Either outcome demands that U.S. executives stay closely engaged with regulatory developments and invest in flexible AI governance programmes that can adapt to fast-changing requirements.

United Kingdom: AI Regulation takes shape

Across the Atlantic, the United Kingdom is accelerating its move towards formal AI oversight. On July 3, the proposed Artificial Intelligence Regulation and Safety Bill passed its Second Reading in the House of Lords - a key step in the UK's legislative process. This progress signals growing political consensus for a dedicated AI governance framework, after previous governments favoured a sector-by-sector approach instead of a single overarching law.

The draft law is expected to set cross-sector principles for "responsible AI" and empower existing regulators (like the Financial Conduct Authority for finance and the Information Commissioner's Office for data protection) to enforce new requirements on AI use within their domains. The UK's approach is envisioned as more "light-touch" and decentralised than the EU's top-down AI Act, but the turn towards a national AI law marks a notable shift in strategy. In recent months, UK regulators have already begun issuing sector-specific AI guidelines - for example, in financial services and data privacy - to bridge the gaps while broader legislation is in progress.

For businesses operating in the UK, this legislative momentum means that clearer compliance obligations are likely on the horizon. Organisations may soon face minimum standards for AI transparency, risk assessment, and oversight across many industries, replacing today's voluntary guidance with enforceable rules. Forward-looking companies are advised to start aligning their AI governance practices with anticipated requirements now; establishing internal AI oversight committees, robust risk management processes, and thorough documentation will help ensure they are prepared once the UK's regulatory framework comes into force.

Legal precedents: AI liability and IP risks

New legal developments are translating AI-related risks into real-world liability for enterprises. In a landmark case last month, a court in Germany ruled that Google is liable for defamatory content produced by its AI-driven search summary feature (www.dw.com [1]). The Munich judges concluded that the AI-generated summary - which had falsely linked two local publishers to scams - wasn't just aggregating search results but was presenting a new, libelous statement that Google must answer for (www.dw.com [2]) (www.technology.org [3]). This precedent-setting judgment represents the first instance of an AI provider being held directly responsible for its algorithm's output (www.technology.org [4]). It has prompted widespread discussion about accountability in AI, and other companies deploying generative AI in customer-facing applications are now on notice that they could face legal claims if their systems produce harmful misinformation (www.technology.org [5]) (www.technology.org [6]).

Beyond defamation, a broader legal battle is intensifying around intellectual property and data privacy in the age of AI. More than 70 AI-related lawsuits are ongoing or recently settled in the U.S. and internationally, targeting tech firms for allegedly infringing copyrights or privacy rights through AI training data and outputs. The financial exposure from these cases has ballooned: industry analysis estimates over $50 /billion in cumulative damages are claimed across active AI litigation as of mid-2026. In one high-profile example, an AI developer agreed to a $1.5 /billion settlement in 2025 after being accused of using hundreds of thousands of pirated books to train its models - the largest such settlement on record. As regulators and courts sharpen their focus on these issues, companies must rigorously vet their AI supply chains (from training data licensing to output monitoring) to mitigate legal risks. The era of unchecked AI experimentation is ending; going forward, businesses will need to demonstrate that their AI innovations respect existing laws and rights or face substantial liability.

References:

[1] www.dw.com — https://www.dw.com/en/german-court-holds-google-liable-for-fake-ai-answers/a-77527661#:~:text=Images%2FZUMA%2Fpicture%20alliance%20Advertisement%20A%20court,It%20had%20%E2%80%8B%E2%80%8Blinked
[2] www.dw.com — https://www.dw.com/en/german-court-holds-google-liable-for-fake-ai-answers/a-77527661#:~:text=treatment%20as%20conventional%20search%20results,party
[3] www.technology.org — https://www.technology.org/2026/06/12/german-court-google-ai-overviews-liable/#:~:text=scams%20and%20dubious%20business%20practices,%E2%80%94%20and%20the%20court%20sided
[4] www.technology.org — https://www.technology.org/2026/06/12/german-court-google-ai-overviews-liable/

#:~:text=Google%20Liable%20for%20False%20Claims,decision%20that%20could%20haunt%20every
[5] www.technology.org — https://www.technology.org/2026/06/12/german-court-google-ai-overviews-liable/#:~:text=Google%20Liable%20for%20False%20Claims,decision%20that%20could%20haunt%20every
[6] www.technology.org — https://www.technology.org/2026/06/12/german-court-google-ai-overviews-liable/#:~:text=cited%20sources,%E2%80%94%20and%20the%20court%20sided

Emerging AI threats and enterprise responses

The past 48 hours have also highlighted the double-edged nature of rapid AI advancement, as new capabilities bring new risks. Elon Musk's AI venture (now part of his newly rebranded "X" enterprise) launched its most advanced model yet, known as Grok 4.5 (aigovernance.com [1]). Billed as a next-generation 'agentic' AI, Grok 4.5 can self-direct long-running tasks such as writing and executing its own software code, using reinforcement learning on extensive programming data to act autonomously for hours at a time (aigovernance.com [2]). While such capabilities promise significant productivity gains, they also present unprecedented governance challenges. An AI that can make and act on decisions with minimal human intervention requires stringent oversight: companies adopting these cutting-edge models will need to implement robust controls to ensure the AI's actions remain safe, legal, and aligned with organisational policies - including mechanisms for human override, thorough testing for unintended behaviours, and assurances that AI-generated code does not violate intellectual property licenses (aigovernance.com [3]).

At the same time, malicious use of AI is becoming a very real enterprise threat. This week, cybersecurity researchers documented what appears to be the first-ever ransomware attack orchestrated largely by an AI system itself. In the incident, an autonomous agent (a self-improving AI programme) dubbed 'JadePuffer' exploited a known software vulnerability, harvested access credentials, and proceeded to lock up a company's database via encryption - all without direct human involvement. The AI even fixed its own errors on the fly and executed over 600 steps in rapid succession during the attack, something no human hacker could match (cyberscoop.com [4]) (cyberscoop.com [5]). Security experts warn that this leap in attack automation drastically lowers the barrier to launching sophisticated cyberattacks, potentially enabling less-skilled actors to carry out breaches that previously required extensive expertise (cyberscoop.com [6]).

In response to these emerging risks, industry groups and companies are hastening to strengthen their AI governance practices. A global IT standards consortium, the GSD Council, this week published an eight-step framework for safely implementing generative AI in IT service management operations (aigovernance.com [7]). The framework emphasises measures like strict access controls for AI systems, proactive "hallucination" detection and correction processes, and alignment with evolving regulatory compliance requirements, all intended to prevent AI failures from disrupting critical business services (aigovernance.com [8]).

Meanwhile, leading firms are sharing their own playbooks for managing advanced AI. For instance, OneTrust - a major compliance and privacy software provider - recently detailed how it established an internal AI governance committee to oversee its use of autonomous AI tools (aigovernance.com [9]). The company's governance framework mandates full traceability for AI decision-making, limits any AI system's access privileges to a "least privilege" level, and requires a formal review before deploying AI that can act without human intervention (aigovernance.com [10]). By publicizing these controls, OneTrust offers other enterprises a concrete example of how to proactively tame the risks of "agentic" AI within an organisation.

The takeaway for businesses is that as AI technology evolves, so must risk management. Whether it's adopting the next wave of powerful AI systems or defending against new AI-enabled threats, companies need to update their governance frameworks in real time. Boards and executives are

increasingly expected to ensure that AI deployments are transparent, secure, and compliant with the law - a responsibility that cannot be delegated or delayed in today's environment.

References:

[1] aigovernance.com — https://aigovernance.com/news#:~:text=2026,5%20xAI%20agentic%20AI%20AI
[2] aigovernance.com — https://aigovernance.com/news#:~:text=8%2C%202026%2C%20positioning%20it%20as,AI%20Act%20software%20engineering%20AI
[3] aigovernance.com — https://aigovernance.com/news#:~:text=Enterprises%20adopting%20Grok%204,AI%20Act%20software%20engineering%20AI
[4] cyberscoop.com — https://cyberscoop.com/sysdig-judepuffer-ai-agentic-ransomware-attack/#:~:text=narrated%20their%20objectives%20in%20plain,switched%20its%20approach%20from%20subprocess
[5] cyberscoop.com — https://cyberscoop.com/sysdig-judepuffer-ai-agentic-ransomware-attack/#:~:text=AI%20agent%20also%20quickly%20diagnosed,were%20used%20in%20the%20attack
[6] cyberscoop.com — https://cyberscoop.com/sysdig-judepuffer-ai-agentic-ransomware-attack/#:~:text=For%20Clark%2C%20there%20is%20a,stops%20at%20Cybersecurity%20Dive%2C%20CIO
[7] aigovernance.com — https://aigovernance.com/news#:~:text=collaboration%20procurement%20governance%20Research%20Global,IT%20operations%20functions%20should%20treat
[8] aigovernance.com — https://aigovernance.com/news#:~:text=collaboration%20procurement%20governance%20Research%20Global,IT%20operations%20functions%20should%20treat
[9] aigovernance.com — https://aigovernance.com/news#:~:text=AI%20security%20Corporate%20Policy%20Global,that%20compliance%20teams%20at%20other
[10] aigovernance.com — https://aigovernance.com/news#:~:text=AI%20security%20Corporate%20Policy%20Global,that%20compliance%20teams%20at%20other

Key Statistics

- 72% of S&P 500 companies disclosed AI-related risks in 2025, up from just 12% in 2023 ([www.thomsonreuters.com](https://www.thomsonreuters.com/en-us/posts/sustainability/ai-governance-gap-esg-risks/#:~:text=implementation%20gap%20is%20alarming%20%E2%80%94,of%201%2C000%20companies%20indicates%20a)).

KEY TAKEAWAY

Global regulators are enforcing new AI rules, courts are setting first liability precedents, and even malicious AI use is now a reality. Boards and executives must treat AI risk governance as a non negotiable priority to stay compliant and competitive.

Sources

AI Governance News (Daily Roundup) – AI Governance Institute, July 7-9 2026
https://aigovernance.com/news

German court holds Google liable for fake AI answers (DW, June 12 2026)
https://www.dw.com/en/german-court-holds-google-liable-for-fake-ai-answers/a-77527661

JADEPUFFER: Agentic ransomware for automated database extortion (Sysdig Threat Research, Jul 1 2026)
https://www.sysdig.com/blog/jadepuffer-agentic-ransomware-for-automated-database-extortion

SpaceXAI launches Grok 4.5 model for coding, agentic tasks (Reuters, Jul 8 2026)
https://money.usnews.com/investing/news/articles/2026-07-08/spacexai-launches-grok-4-5-model-for-coding-agentic-tasks

New data reveals AI governance gap between policy and practice, creating ESG risks (Thomson Reuters, Feb 23 2026)
https://www.thomsonreuters.com/en-us/posts/sustainability/ai-governance-gap-esg-risks/