March 4, 2026

Connecticut Joint Committee on General Law
Legislative Office Building, Room 3500
Hartford, CT 06106

Re: SB 5 - "An Act Concerning Online Safety"

Dear Chair Maroney and Chair Lemar, and Members of the Joint Committee on General Law:

On behalf of the Computer & Communications Industry Association (CCIA), I write to respectfully oppose SB 5 in advance of the Committee hearing on March 4, 2026. CCIA is an international, not-for-profit trade association representing a broad cross-section of communications and technology firms.[1] As a trade association whose members provide digital services across state and national borders, CCIA has a strong interest in ensuring that AI policy is risk-based and innovation-enabling.

While we share the Legislature's commitment to ensuring the safety and privacy of Connecticut's users, the bill as currently drafted creates an unworkable regulatory framework that would stifle innovation, jeopardize security, and burden both consumers and businesses.

SB 5's broad and vague scope misunderstands the AI landscape.

The bill relies on broad, imprecise definitions of terms like "foundation model," "frontier developer," "large frontier developer," and "reasonable internal process" that would encompass a vast range of common digital tools, including customer service chat functions, workplace productivity software, educational platforms, accessibility technologies, and automated moderation or safety tools. These technologies differ significantly in function, risk profile, and user interaction, yet the bill treats them as functionally equivalent. Similarly, the open-ended, subjective definition of "artificial intelligence companion" risks scoping in businesses such as customer service chatbots that answer support questions, productivity tools that use conversation interfaces, wellness applications that respond to user prompts about goals or progress, and other products and services without the capabilities this bill contemplates.

The absence of clear, administrable boundaries makes it difficult for companies to determine whether their products fall within the scope of the bill or which obligations apply. Faced with ambiguous standards and significant penalties, providers are likely to respond by limiting features, restricting access, or withdrawing services entirely rather than attempting to interpret and comply with unclear requirements. This approach ultimately reduces consumer choice and limits access to beneficial technologies without meaningfully advancing safety or privacy objectives.

[1] For more than 50 years, CCIA has promoted open markets, open systems, and open networks. CCIA members employ more than 1.6 million workers, invest more than $100 billion in research and development, and contribute trillions of dollars in productivity to the global economy. A list of CCIA members is available at https://www.ccianet.org/members.

AI policy must be balanced, flexible, and understand the ecosystem.

As AI evolves rapidly, it is crucial to find a balance in regulation to ensure the rules are not so rigid that they hinder innovation and U.S. competitiveness. Achieving this kind of balance requires thoughtful and adaptable legislation that is informed by the principles of responsible AI and can be applied across many different contexts. Rather than imposing overly detailed and narrowly tailored rules, the focus must be on establishing frameworks that enable the design of AI systems and allow them to serve society's best interests. In the absence of a single federal framework regulating AI, any single state's efforts to implement broad regulation would likely place a state like Connecticut at a competitive disadvantage by inhibiting the use of new technologies to further growth, while other states may not implement such obstacles.[2]

One example of this is two-factor authentication (2FA) and the transformation that password security has undergone in the last few years. If policymakers had adopted prescriptive regulations for privacy as it relates to 2FA, those rules would have quickly become outdated and potentially blocked secure innovations such as passkeys. The same risk applies to AI policy, especially prescriptive rules that lock in specific technical assumptions, licensing models, or training methods to a rapidly changing landscape.

SB 5 risks creating a fragmented regulatory environment.

The bill would also contribute to a growing panoply of state artificial intelligence laws that impose inconsistent and potentially conflicting obligations on interstate digital services. Artificial intelligence systems are developed, trained, and deployed on a national and global scale. Prescriptive state-level mandates risk becoming outdated quickly, complicating compliance, and discouraging investment in jurisdictions that adopt rigid or punitive frameworks. Connecticut has long benefited from policies that promote innovation and technological growth. A fragmented regulatory approach threatens that position by making it more difficult for companies to deploy new services and features in the state.

Existing laws already address many aspects of AI, including workforce and employment concerns.

Despite the ongoing trend of AI-specific legislation, it is important to recognize that many of the risks commonly associated with AI are already addressed through existing federal and state frameworks. AI does not operate in a legal vacuum but rather, it is a tool used within regulated markets that are already governed by long-standing consumer protection, civil rights, privacy, and other product liability laws. Ahead of proposing such laws, policymakers must consider what laws AI systems are already covered by. It is important to build upon existing legal protections and focus narrowly on clearly defined gaps where demonstrable harms are not yet addressed. A balanced approach that does not layer expansive new liability regimes on AI developers will better protect consumers and preserve the innovation ecosystem.

[2] CCIA, *Understanding AI: A Guide To Sensible Governance* (June 2023), https://ccianet.org/library/understanding-ai-guide-to-sensible-governance/.

Furthermore, the enormous increases in productivity, real wages, and GDP per capita from the pre-industrial era to today are primarily attributable to the development of new automation-enhancing technologies and their deployment across the economy. Technological change resulting in increased automation has reshaped society through the creation of new roles, higher productivity gains, and rising living standards across the economy.[3] The precedent of more bank tellers being hired decades after ATMs were introduced and more radiologists getting hired after AI tools became widespread in radiology practices shows that automation is often a driver of productivity-enhancing transformation and task switching rather than worker replacement.[4]

The bill's subjective definitions would create compliance uncertainty.

It is difficult to objectively determine when a given output meets the listed criteria, including "prioritizing validation of the user's beliefs, preferences or desires over factual accuracy or the user's safety". These vague terms do not allow businesses to determine in advance whether their products and services comply with the law. Precise narrowing is required to focus any regulation solely on the intended targets.

The obligation for "synthetic digital content" to be "marked and detectable" does not define either of those terms. Without clear definitions of what qualifies, businesses are left to guess at the standards required for compliance, increasing the likelihood of inconsistent implementation or arbitrary enforcement. It is unclear if "marking" requires visible disclosure to users, embedded metadata, or a combination of the two, and what satisfies a detectable standard. Such vagueness not only chills competition by requiring services to divert resources to costly compliance burdens but also exposes companies to liability despite good-faith efforts to meet said obligations.

The requirement to include a "protocol to take reasonable efforts to detect and address" particular "ideation" or an "indicator" may require costly, burdensome, or invasive monitoring. Absent clear definitions, companies may feel compelled to deploy expansive monitoring systems to mitigate liability risk. That could necessitate continuous review of user interactions or data retention practices. Such measures are not only expensive to design, implement, and audit but may raise significant privacy and data protection concerns. Smaller firms in particular may lack the technical or financial capacity to build and maintain these systems, effectively creating barriers to entry.

The bill does not define "any activity that poses a specific and substantial danger to the public health or safety" or what it means to "immediately take action to eliminate such danger." This lack of clarity leaves regulated entities without meaningful guidance as to the scope of their obligations or the threshold that triggers them. "Specific and substantial" is inherently subjective, and without definitions or limiting principles, companies may struggle to distinguish

[3] Trevor Wagener, *AI Is an Economic Engine Propelling Productivity and Wages Up, and Prices Down*, CCIA (Oct. 5, 2025), https://ccianet.org/articles/ai-is-an-economic-engine-propelling-productivity-and-wages-up-and-prices-down/.
[4] Trevor Wagener, *What Bank Tellers and Radiologists Can Tell Us about Our Job Security in the AI Era*, CCIA (Feb. 3, 2026), https://ccianet.org/articles/what-bank-tellers-and-radiologists-can-tell-us-about-our-job-security-in-the-ai-era/

between imminent threats or more speculative risks. Additionally, the mandate to "immediately take action" does not indicate the timeframe required, the types of measures contemplated, or the degree of certainty necessary before acting.

Some of the bill's limits may aim to provide some flexibility, such as "as far as technically feasible and in a manner that is consistent with any nationally or internationally recognized technical standards, ensure that such developer's technical solutions are effective, interoperable, robust and reliable, considering the specificities and limitations of different types of synthetic digital content, the implementation costs, and the generally acknowledged state of the art." As read, this language appears to acknowledge that technical capabilities evolve and that not all solutions will be equally effective across different forms of synthetic digital content. However, these qualifiers may ultimately provide limited practical certainty. "State of the art" implementation costs are inherently dynamic and subjective, potentially exposing organizations to second-guessing as technologies and best practices evolve.

Similarly, there are some exceptions for content that "is published to inform the public on any matter of public interest" or "is unlikely to mislead a reasonable person consuming such synthetic digital content," or if it "forms part of an evidently artistic, creative, satirical, fictional analogous work or program," the required disclosure should "not hinder the display or enjoyment of such work or program." While these carveouts are appreciated, they introduce additional ambiguity that may be difficult to operationalize in practice. Together, these exceptions may create as much uncertainty as they resolve, leaving digital services and developers to make difficult, subjective determinations under the threat of potential liability.

Finally, the proposed safe harbor program's utility depends on implementation. A well-designed safe harbor should include predictable timelines for approval, clear documentation requirements, and protection against retroactive second-guessing where a company has adhered to approved standards. Without detailed guardrails and durable protections, a nominal safe harbor risks becoming an additional compliance layer rather than a meaningful shield under uncertainty and liability.

The bill would lead to a chilling effect on the flow of information online and undermine user experiences.

SB 5 requires that "the operator shall provide such notice to the user (1) at the beginning of each artificial intelligence companion interaction, except the operator shall not be required to provide such notice to the user more frequently than once per day, and (2) at least once hourly during any continuous artificial intelligence companion interaction." These kinds of mandatory notices undermine the user experience on services, and are not only ineffective but can even backfire — users may instead just keep an app or site open even more, or just ignore it due to the phenomenon known as 'notice fatigue,' as seen with frequent cookie notices from Europe or California.

Additionally, compliance with the bill is likely to result in a significant operational and technical burden through the mandatory implementation of time spent and the display of the notices themselves, especially for small and medium-sized digital services that fall within the scope of the bill's broad definition. The costs required to redesign interfaces and conduct testing favor

larger companies with the necessary resources, potentially harming smaller competitors or decentralized platforms.

Age verification raises significant privacy concerns.

Since the bill provides no objective way for covered operators to demonstrate that they have "reasonably" determined that a user is over 18, they will be effectively forced to institute age verification to ensure compliance. This approach creates significant problems. Every approach to age determination presents trade-offs between accuracy and privacy[5]—in addition to significant costs, especially for startups[6]—and there is no one-size-fits-all approach. Different services consider various factors, including but not limited to their user base, the service offered, risk calculation, privacy expectations, and economic feasibility. A recent Digital Trust & Safety Partnership (DTSP) report, *Age Assurance: Guiding Principles and Best Practices*, contains guiding principles for age assurance and discusses how digital services have used such principles to develop best practices.[7]

Determining a user's age inherently requires collecting additional sensitive data from those users, and any document capable of verifying a user's age will likely contain sensitive information. The Commission Nationale de l'Informatique et des Libertés (CNIL) analyzed several existing online age verification solutions but found that none of these options could satisfactorily meet three key standards: 1) providing sufficiently reliable verification; 2) allowing for complete coverage of the population; and 3) respecting the protection of individuals' data, privacy, and security.[8]

Best practices must be determined by practitioners.

SB 5 calls for a working group to make recommendations on best practices "to avoid the negative impacts, and to maximize the positive impacts, on services and state employees in connection with the implementation of new digital technologies, including...artificial intelligence." Determining "best practices," in AI governance and technology policy more broadly is not solely a matter of government preference, and should not be dictated by prescriptive regulations. Meaningful best practices require consultation with industry stakeholders who design, deploy, and operationalize these systems, along with the relevant stakeholders who understand how AI functions in real-world settings.[9] The safeguards, tools, and governance approaches that best mitigate risk are those that are principle-based and readily adaptable to the constantly-evolving nature of AI and like technologies.

---

[5] Kate Ruane, *CDT Files Brief in* NetChoice v. Bonta *Highlighting Age Verification Technology Risks* (Feb. 10, 2025), https://cdt.org/insights/cdt-files-brief-in-netchoice-v-bonta-highlighting-age-verification-technology-risks/.
[6] Engine, *More than just a number: How determining user age impacts startups* (Feb. 2024), https://static1.squarespace.com/static/571681753c44d835a440c8b5/t/65d51f0b0d4f007b71fe2ba6/1708465932202/Engine+Report+-+More+Than+Just+A+Number.pdf.
[7] *Age Assurance: Guiding Principles and Best Practices*, Digital Trust & Safety Partnership (Sept. 2023), https://dtspartnership.org/wp-content/uploads/2023/09/DTSP_Age-Assurance-Best-Practices.pdf.
[8] *Online Age Verification: Balancing Privacy and the Protection of Minors,* CNIL (Sept. 22, 2022), https://www.cnil.fr/en/online-age-verification-balancing-privacy-and-protection-minors.
[9] *See, e.g.*, Digital Trust & Safety Partnership, *Best Practices for AI and Automation in Trust and Safety* (Sept. 2024), https://dtspartnership.org/best-practices-for-ai-and-automation-in-trust-and-safety/.

Consumer-facing digital services have already built considerable consensus around mitigating risk to users and other parties. A recently published international standard, ISO/IEC 25389, reflects the evolving industry consensus on mitigating online risk.[10] This reflects the first horizontal standard for safety on online consumer services. Insofar as there is work occurring in this space, it should incorporate the existing consensus around best practice: safety by design, appropriate governance, application, iteration and improvement, and transparency.

* * * * *

CCIA supports sensible AI governance, including appropriate transparency disclosures and robust data protections for minors. However, SB 5 goes beyond these goals, introducing mandates that would make Connecticut a difficult environment for AI innovation. We encourage Committee members to pause advancing legislation that is not adequately tailored to this objective.

We appreciate the Committee's consideration of these comments and stand ready to provide additional information as Connecticut considers proposals related to technology policy.

Respectfully submitted,

Kyle J. Sepe
State Policy Manager, Northeast Region
Computer & Communications Industry Association

[10] ISO/IEC 25389:2025, Information technology — The safe framework (Edition 1, June 2025), https://www.iso.org/standard/90106.html.