Testimony of Meghan Pensyl
Director, Policy
Business Software Alliance

Subject Matter Hearing

Illinois Senate Executive Subcommittee on AI and Social Media

April 9, 2026

Good afternoon, Chair Cunningham, Vice Chair Castro, and Senator Rezin. My name is Meghan Pensyl, and I am a Policy Director at the Business Software Alliance.

BSA is the leading trade association for the global enterprise software industry.[1] BSA members are at the forefront of developing business-to-business technologies—including AI—and their products are used by companies across every sector of the economy.[2] I commend the subcommittee for convening today's hearing, and I thank you for the opportunity to testify.

AI is changing the way we live and work, and it has real-world benefits. But realizing the potential of AI requires trusting that the technology is developed and deployed responsibly. Crafting AI legislation that promotes the responsible adoption of AI and protects against its misuse is one of the most important technology issues today, and one we already see governments beginning to tackle, including in the European Union and across the states. The most effective way to address this issue is through a single, national law. However, just as states took the lead in adopting consumer privacy laws, we recognize that states are again leading with AI legislation.

It's critical to get this right. As you consider how to regulate AI, I want to underscore the key role state governments should play in encouraging AI adoption. States that most effectively promote secure AI adoption in the private and public sectors will see the greatest economic benefits, stimulate further innovation, and deliver substantial economic gains across every industry sector. Proactive policies that encourage the use of trustworthy enterprise AI—through talent, infrastructure and data, and workable frameworks—can ensure that these benefits are widely distributed and aligned with economic goals.

---

[1] BSA's members include: Adobe, Alteryx, Amadeus, Asana, Atlassian, Autodesk, Avalara, Bentley Systems, Box, Cisco, Cohere, Cohesity, Dassault Systemes, Databricks, Datadog, Docusign, Dropbox, Elastic, EY, Graphisoft, HubSpot, IBM, Kyndryl, MathWorks, Microsoft, Notion, Okta, OpenAI, Oracle, PagerDuty, Palo Alto Networks, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Trend Micro, TriNet, Veeam, Workday, Zendesk, and Zoom Communications Inc.

[2] *See* Business Software Alliance, Artificial Intelligence in Every Sector, *available at* https://www.bsa.org/files/policy-filings/06132022bsaaieverysector.pdf.

To help ensure companies across every sector can benefit from enterprise AI, any AI legislation should seek to create clear obligations for companies and provide strong protections for consumers. To achieve this, we recommend AI legislation:

- Aims to broadly distribute the benefits of AI and encourage trustworthy AI adoption;
- Addresses the uses of AI that have the greatest impact on consumers;
- Reflects the different roles and responsibilities of different actors along the AI value chain;
- Ensures strong enforcement; and
- Promotes interoperability and incorporates stakeholder feedback.

We also offer our views on privacy bills pending before the subcommittee.

I. AI Adoption Creates Benefits Across Industry Sectors

At the outset, it is important to recognize that the economic benefits of AI are not limited to one industry sector or one business model. Instead, the promise that AI may one day impact every industry is quickly turning into a commercial reality. Airlines now use AI systems to more efficiently clean planes between flights; farmers use AI to analyze large amounts of weather information to maximize their harvest; manufacturers use AI to test new prototypes; and construction companies build AI-generated "digital twins" of real-life cities to understand the impacts of a proposed design.

I want to illustrate a few examples of how companies in all industries are adopting AI-powered enterprise software:

- In healthcare, a large pharmacy chain uses an advanced platform to forecast demand and redistribute medications across thousands of store locations and to deliver near real-time insights and recommendations for pharmacists to provide more personalized advice to patients. This helps managers understand the supply chain, store labor and productivity, patient vaccine scheduling, and prescription pickup processes.

- In manufacturing, a car maker used generative AI technology to redesign a seat bracket, which secures seat belt fasteners to seats and seats to floors, that is 40 percent lighter and 20 percent stronger than the previous iteration. Changes like these can help reduce the amount of material needed to build a car and make vehicles more fuel efficient.

- In agriculture, the research division of an enterprise software provider partnered with a climate risk company to develop software capable of providing more accurate long-range weather predictions. Traditional weather forecasting methods can provide accurate predictions for a seven-day window. By leveraging AI, the researchers are developing new forecasting models to provide accurate predictions of weather trends two to six weeks out from a given date. By providing reliable extended forecasts, these tools will help water managers predict snowpack and water availability for irrigation, hydropower, and other critical agricultural and environmental uses.

Because BSA members work with companies across every sector of the economy, we have unique insight into AI's tremendous potential to further spur digital transformation and the policies that can best support the responsible adoption and use of AI. BSA's views are informed by our experience working with member

companies to develop the BSA Framework to Build Trust in AI,[3] a risk management framework for mitigating the potential for unintended bias throughout an AI system's lifecycle. Built on a vast body of research and informed by the experience of leading AI developers, the BSA Framework outlines a lifecycle-based approach for performing impact assessments and highlights corresponding risk mitigation best practices.

For state governments in particular, AI adoption offers efficiencies as states face growing demands on constrained resources. Enterprise AI tools, when deployed with appropriate safeguards, can help state governments automate routine tasks, augment human decision-making, and unlock insights that could otherwise remain buried in data. AI-powered tools can help states cut processing times and improve service delivery for consumers.

As you consider AI policies, we encourage you to recognize the benefits of adopting AI across both the public and private sectors, to ensure its benefits are broadly distributed across the state's economy.

II. Guardrails Are Important for High-Risk Uses of AI

Legal guardrails are warranted for high-risk uses of AI that have the most significant impacts on consumers' lives. Many everyday uses of AI present few risks to individuals and create significant benefits, like helping to organize digital files, auto-populate common forms for later human review, improve a company's ability to forecast supply chain issues, and detect, prevent, and respond to cybersecurity threats.

However, when AI systems decide whether someone receives important benefits and services, like housing, healthcare, and employment opportunities, companies should be accountable for developing and deploying those systems responsibly.

In addition to efforts to protect consumers when AI systems make consequential decisions about them, states have begun to regulate certain AI technologies, including frontier models. Laws that govern frontier model safety implicate national security issues and should be addressed through a single, national law. BSA continues to work with Congress towards that goal.

For high-risk uses of AI that decide whether consumers are granted or denied important life opportunities, new safeguards are important—and legislation can leverage tools that already exist to help companies identify and mitigate potential risks. BSA supports requiring companies that develop or deploy AI for high-risk uses to: (1) adopt risk management programs; and (2) conduct impact assessments. These measures can help companies identify and mitigate risks when AI makes important decisions about consumers—and increase trust that AI is developed and used responsibly.

a. Risk Management Programs

Companies should implement risk management programs that help them identify and mitigate risks associated with high-risk AI systems.

---

[3] *See* Business Software Alliance, Confronting Bias: BSA's Framework to Build Trust in AI, *available at* https://www.bsa.org/reports/confronting-bias-bsas-framework-to-build-trust-in-ai.

Risk management is particularly important in contexts like AI, privacy, and cybersecurity, where the combination of quickly evolving technologies and highly dynamic threat landscapes can render traditional approaches to compliance ineffective. Risk management programs establish repeatable processes for companies to identify and mitigate potential risks that can arise throughout the lifecycle of an AI system. Risk management programs have two key components: (1) a governance framework of policies, procedures, and personnel that support the company's risk management function; and (2) a scalable process for performing impact assessments that identify and mitigate risks of an AI system.

One way for companies to establish risk management programs is by using the AI Risk Management Framework (AI RMF), which was released in 2023 by the National Institute of Standards and Technology (NIST).[4] The AI RMF builds on NIST's work creating frameworks for managing cybersecurity and privacy risks.[5] The AI RMF helps companies incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products. Ultimately, effective AI risk management programs should support cross-company coordination to promote the identification and mitigation of risks across the lifecycle of an AI system.

b. Impact Assessments

BSA recognizes that performing impact assessments for high-risk uses of AI is a key part of creating a meaningful risk management program.

Both developers and deployers should use impact assessments as a tool for the responsible development and use of high-risk AI systems—and each type of company should conduct an impact assessment that reflects their role in developing or deploying the AI system. Impact assessments have three purposes: (1) identifying potential risks that an AI system may pose; (2) quantifying the degree of potential harms the system could generate; and (3) documenting steps taken to mitigate those risks.[6]

Impact assessments are already widely used in a range of other fields, including privacy, as an accountability mechanism that demonstrates a product or system has been designed in a manner that accounts for the potential risks it may pose to the public. Because impact assessments already exist today in other fields, they can be readily adapted to help companies identify and mitigate AI-related risks.[7]

Risk management programs and impact assessments also avoid the pitfalls of other approaches, such as third-party audits. The AI audit ecosystem is still immature and not fully equipped to assure the

---

[4] *See* NIST AI Risk Management Framework, *available at* https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.
[5] *See* NIST, Cybersecurity Framework, Questions and Answers, (discussing federal agency use of the NIST CSF), *available at* https://www.nist.gov/cyberframework/faqs.
[6] *See* Business Software Alliance, Impact Assessments: A Key Part of AI Accountability, *available at* https://www.bsa.org/files/policy-filings/08012023impactassess.pdf.
[7] Many state privacy laws also require impact assessments. *See*, e.g., Colorado Privacy Act, Colo. Rev. Stat. Tit. 6, Art. 1, Pt. 13 §§ 6-1-1301–6-1-1313; Connecticut Data Privacy Act Conn. Gen. Stat. Tit. 42, Ch. 743jj, Sec. 42-515-525; Florida Digital Bill of Rights, Fla. Stat. Tit. XXXIII, Ch. 501, Pt. V, Sec. 501.701-722; Oregon Consumer Privacy Act, Or. Rev. Stat. 646A.570-646A.589; Texas Data Privacy and Security Act, Tex. Bus. & Com. Code Ann. § 541.001-205; Virginia Consumer Data Protection Act; Va. Code Tit. 59.1, Ch. 53, § 59.1-575-585. Globally, privacy and data protection laws worldwide use impact assessments as a tool for improving accountability.

performance of AI systems.[8] Unlike regulated industries, such as accounting and financial services, there is not a professional body to certify AI auditors and standards are still in development. As a result, the range of qualified and available auditors may vary greatly—and undermine the goal of implementing audits consistently across different AI systems.

Moreover, third-party audits often require access to confidential business information, creating concerns around the treatment of trade secrets and other sensitive information. In some cases, audits may require access to consumers' personal information, and providing large data sets associated with AI systems to a third-party auditor can create significant privacy concerns. At the same time, if an auditor is provided access to the AI system itself, it can create security concerns. Engaging a third-party auditor would require working through these and other issues that are not present in the context of internally focused accountability tools, such as risk management programs and impact assessments.

III. AI Legislation Should Distinguish Between Different Entities in the AI Ecosystem

The AI supply chain is evolving, and AI legislation should not create one-size-fits-all requirements when companies have very different roles.

For example, one company may develop an AI model, a second company may integrate that AI model into an application, and yet another company may use that application to make decisions about individual consumers.

All companies that develop and use high-risk systems have responsibilities to manage AI risks, but obligations must reflect the role of each type of company, since each will know different information about the AI system and will be able to take different actions to identify and mitigate risks. Any AI legislation must recognize these differences to be workable in practice and to safeguard consumers.

AI legislation that clearly differentiates between companies that develop, integrate, or deploy AI and assigns responsibilities based on each type of company's role in the AI value chain should also hold companies accountable when they fail to fulfill their obligations. Not all approaches to assigning liability, however, are workable in AI policy.

a. Approaches to AI Liability

The most straightforward approach to ensuring that companies develop and use AI responsibly is to place clear obligations on them, based on their role in the AI value chain, and to hold them accountable when they fail to comply. This approach creates clarity for businesses in understanding their responsibilities and provides robust protections for consumers.

At the state level, we've also seen interest in ensuring companies develop and deploy AI responsibly by assigning them a duty of care. The concept of a "duty of care" is deeply rooted in tort law, which governs

[8] *See* Business Software Alliance, Enhancing AI Accountability: Effective Policies for Assessing Responsible AI, *available at:* https://www.bsa.org/policy-filings/enhancing-ai-accountability-effective-policies-for-assessing-responsible-ai.

civil wrongs and personal injury. Courts frequently impose a duty of care on individuals or organizations that have the power to prevent foreseeable harm to others. For example, drivers must operate their cars safely to avoid injuring pedestrians; a doctor must act as a reasonably competent physician would under similar circumstances; a company must ensure that its products are safe for ordinary use. These duties are not static rules—they evolve with context, technology, and social expectations. The standard is flexible, focusing on whether an actor took reasonable steps to prevent foreseeable harm given their role, expertise, and resources. It is also not tied to strict liability or products liability regimes. As a result, that flexibility can promote responsible development and the use of fast-changing technologies like AI, especially when paired with a specific list of actions that companies can take to meet the standard.

Policymakers focused on AI issues have occasionally looked to other liability systems, such as products liability or strict liability. This is problematic, as those systems assign liability based on outcomes rather than conduct.

Under products liability, for instance, a manufacturer can be held liable for harm even if it took all reasonable precautions. That approach is a poor fit for AI systems, however, because outcomes depend heavily on how an AI tool is deployed. For example, a developer may create an AI system that is well-suited to specific uses, but a deployer might then create significant risks if they use it in other settings. Each business should be held responsible for what it can control — and not for outcomes that result from others' actions.

In contrast, a straightforward approach to assigning responsibilities to different companies and holding each company accountable for their obligations emphasizes responsible behavior — encouraging both developers and deployers to identify and address risks, conduct robust testing, and act promptly when problems emerge.

IV. AI Legislation Should Ensure Strong Enforcement

Strong enforcement is needed in any AI legislation. Exclusive enforcement authority by the Attorney General can help that office establish clear guidance and a consistent approach to enforcing the bill's requirements. Exclusive governmental enforcement by a single regulator ensures companies know how to implement the legislation's obligations—and avoids the conflicting interpretations and confusion likely to arise if courts reach different conclusions about how companies are to apply obligations in AI legislation.

V. AI Legislation Should Promote Interoperability and Incorporate Stakeholder Feedback

Illinois is home to global companies, and your legislation will be most effective when it is interoperable with other approaches to AI regulation. Global companies can better serve their customers when they build strong compliance programs that work across markets. We also encourage you to continue working with stakeholders as you develop your legislation, to understand how your AI law will work in practice, across a range of different industries and uses.

VI. Provisions in Pending Privacy Bills Undermine Workability

In addition to AI-specific proposals, the subcommittee is also considering broader privacy legislation, including SB 2875 and SB 3890. These bills would directly shape how enterprise software companies

develop and deploy AI-enabled services by regulating the data practices on which those services rely. For that reason, we briefly highlight a few provisions in these bills that raise operational concerns and could have unintended consequences for both business practices and effective privacy protection.

Both SB 2875 and SB 3890 include provisions that would allow a controller to object before a processor can engage subprocessors. While intended to increase oversight, this approach would disrupt standard business practices in the enterprise software ecosystem. Processors routinely rely on specialized subprocessors to provide secure, scalable, and resilient services, including cloud infrastructure, cybersecurity tools, and data storage. Requiring advanced objection rights for each engagement involves friction, delays deployment, and creates legal uncertainty in dynamic environments where vendors must respond quickly to security threats and operational needs. It also cuts against well-established, risk-based contracting models that already require processors to flow down data protection obligations to subprocessors. In practice, this provision would not meaningfully improve privacy outcomes but would instead reduce efficiency and limit the ability of companies to implement best-in-class security and compliance measures.

Both bills also grant consumers the right to obtain a list of specific third parties with whom a controller has shared personal data. Requiring disclosure of individualized third-party recipients is not workable in practice, particularly in complex B2B environments. Modern data ecosystems involve large, constantly changing networks of service providers, partners, and infrastructure vendors. Moreover, such disclosures could expose sensitive business information, including proprietary vendor relationships and security architectures, without delivering meaningful benefits to consumers. A more practical approach, consistent with other state privacy laws, is to require disclosure of categories of third parties, which provides transparency while remaining feasible to implement and preserving strong privacy and security protections.

We encourage the subcommittee to address these concerns as you consider SB 2875 and SB 3890, to help ensure privacy legislation remains workable in practice while continuing to provide meaningful protections for consumers.

* * *

I want to conclude by emphasizing the tremendous opportunity you have to develop AI legislation that builds trust and drives the responsible adoption of AI across Illinois's economy. The economic and societal benefits of AI will be captured by states that adopt enterprise AI tools securely, responsibly, and at scale. In recent years, states have been at the forefront of enacting technology laws, and I commend you for your focus on AI. Thoughtful AI legislation can cultivate trust in AI technologies and help ensure that consumers and companies benefit as AI continues to evolve.

Thank you for the opportunity to testify. I look forward to your questions.