March 09, 2026

Re: Request for Information on Practices and Methodologies for Measuring and Improving the Secure Development and Deployment of AI Agent Systems

ACT | The App Association (ACT) appreciates the opportunity to respond to this Request for Information (RFI) from the Center for AI Standards and Innovation (CAISI) at the National Institute of Standards and Technology (NIST). ACT represents small and medium-sized software developers and technology firms that are building and deploying AI-enabled solutions across scientific, industrial, and commercial domains.

The small businesses ACT represents are frequently the source of novel technical approaches and real-world applications, yet they face persistent structural barriers in translating early-stage research into scalable, deployable technologies. For emerging technologies, the policy design choices that agencies make, including the structure, accessibility, and usability of technical guidance, can determine whether best practices become widely adopted or remain aspirational artifacts for only the largest firms.

ACT strongly supports CAISI's work to improve the secure development and deployment of AI agent systems. We offer the comments below to encourage CAISI to prioritize scalable, risk-based guidance, adoption-ready outputs for smaller developers, and measurement approaches that reduce friction in the path to secure deployment.

I. Statement of Interest

ACT is a global policy trade association for the small business technology developer community. Our members are entrepreneurs, innovators, and independent developers within the global app ecosystem that engage with verticals across every industry. We work with and for our members to promote a policy environment that rewards and inspires innovation while providing resources that help them raise capital, create jobs, and continue to build incredible technology. The small businesses and startups we represent both participate and seek to participate in the international standard-setting process to contribute and build on important technical standards. ACT members are active in new platforms, like Web3, develop using indispensable technologies (i.e., artificial intelligence), and innovate on top of technical standards. The value of the domestic ecosystem ACT represents, which we call the app economy, is approximately $1.8 trillion and is responsible for 6.1 million American jobs, while serving as a key driver of the $8 trillion internet of things (IoT) revolution.[1] As leading innovators in critical and emerging technologies (CET), small businesses must be a part of this conversation.

[1] https://actonline.org/wp-content/uploads/Fast-FactS_Website.pdf.

II. ACT Comments

Core Principles for CAISI Agent Security Guidance

Include small developers by design. Small and medium-sized businesses (SMBs) are often excluded not by explicit rules, but by program structures that assume hyperscale capacity or extensive administrative resources.[2] CAISI should ensure its outputs are adoptable by smaller teams that do not have dedicated compliance staff or large security budgets.

Use a risk-based approach tied to operational authority.[3] For agent systems, the most policy relevant risk signal is not the underlying model label, because the same model can present materially different risks depending on how it is deployed. The key governance question is how much real-world authority the agent is granted, including what actions it can execute, whether it operates in a bounded or unbounded environment, and what credentials, tools, and system integrations it can access. Security expectations and oversight should scale in proportion to that authority, with the strongest controls reserved for agents deployed with elevated privileges, consequential capabilities, or operation in high-impact environments.

Prioritize practical adoption and reduce uncertainty. Small developers benefit when government guidance provides clarity on what acceptably efficacious and safe AI looks like, lowers fixed compliance costs, and shortens the path from experimentation to secure deployment.

Protect sensitive information and IP. CAISI should be cautious about approaches that would require sensitive disclosures as a condition of participation or trust. ACT has emphasized in other contexts that programs should not require disclosure of proprietary information, trade secrets, or commercially sensitive assets.

Targeted Responses to CAISI Questions

1(a) and 1(d) Security risks unique to AI agent systems and how they are evolving

[2] *See, e.g.,* Allison Lampert & Mike Stone, *New Cybersecurity Rules for U.S. Defense Industry Create Barrier for Some Small Suppliers,* Reuters (Feb. 20, 2026) (reporting that DoD's Cybersecurity Maturity Model Certification requirements create high compliance costs for small suppliers and that limited audit capacity and uncertainty are prompting some small firms to reconsider or exit the industry).

[3] *See* Madkour, N. et al. "Agentic AI Risk-Management Standards Profile." UC Berkeley Center for Long-Term Cybersecurity (CLTC), February 2026. https://cltc.berkeley.edu/publication/agentic-ai-risk-profile/; Lewis, C., & Isenberg, R. "Deploying Agentic AI with Safety and Security: A Playbook for Technology Leaders." McKinsey & Company, October 16, 2025. https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/deploying-agentic-ai-with-safety-and-security-a-playbook-for-technology-leaders.

AI agent systems can unlock major productivity and safety gains because they can move beyond generating outputs to completing tasks in the real world. That same capability is why security needs to be approached in a way that encourages responsible deployment rather than slowing adoption.

From a policy perspective, the most innovation-friendly way to manage these risks is to focus on the combination of autonomy and access. The same underlying model can be deployed safely or unsafely depending on what tools it can use, what permissions it has, and how tightly its operating environment is bound. As integrations expand and autonomy increases, the goal should be to scope authority appropriately, build in proportionate oversight for higher consequence actions, and ensure accountability across the value chain.

As deployment grows, these needs will become more common, not because agent systems are inherently unsafe, but because they will be used in more settings and with more connections. CAISI can accelerate secure adoption by publishing risk-based profiles and practical templates that help developers and deployers set clear boundaries, implement escalation pathways where needed, and provide straightforward documentation so downstream users can match an agent system to an appropriate environment.

2(a) Security practices for AI agent systems

ACT encourages CAISI to focus on practical, scalable practices that can be adopted by small developers without turning agent security into a compliance regime that only the largest firms can meet.

CAISI should prioritize guidance that helps developers and deployers do three things well:

- Scope an agent's authority based on context and consequence.
- Implement appropriate oversight for higher consequence actions, including clear approval and accountability pathways.
- Produce documentation and testing practices that communicate how the system is intended to operate and what risk management steps are in place.

CAISI can strengthen adoption by structuring its outputs as modular profiles that scale. A baseline set of practices should be feasible for most developers. Enhanced expectations should be reserved for higher risk contexts.

2(e) Cybersecurity guidelines, frameworks, and best practices relevant to agent systems

CAISI should anchor its work in harmonization and alignment, rather than creating a parallel framework that adds complexity for resource constrained developers. ACT has urged

policymakers in other contexts to align AI governance with NIST's AI Risk Management Framework and its scaled, risk-based approach.

CAISI's output can be most useful when they translate existing security concepts into agent relevant guidance that is easier to adopt. For SMBs, the main impediments are typically uncertainty, fragmentation, and high fixed costs rather than unwillingness to manage risk. CAISI can reduce these impediments by providing clear, risk-based profiles and practical templates.

CAISI should also take care with disclosure expectations. ACT's AI Policy Principles[4] emphasize that research on the costs and benefits of transparency should be a priority and should involve collaboration among affected stakeholders to understand when transparency mandates help address risks.

That principle is particularly important in agent security, where overly broad disclosures can create new security risks.

3(b) Assessing the security of a particular AI agent system

CAISI should provide a simple, standardized way to assess an agent system that recognizes not all threats apply to all systems. A practical assessment should focus on a small set of factors that matter most for deployment decisions, including:

- What external actions the agent can take
- How bounded or unbounded the deployment environment is
- What privileges, tools, and sensitive data the agent can access
- What oversight mechanisms exist for higher consequence actions
- What documentation exists to support safe deployment and incident response

This type of risk-based profile will be more useful to small deployers and customers than complex or highly technical disclosures. It also emphasizes risk-based oversight and appropriate distribution of risk and incentives across the AI value chain.

III. Conclusion

AI agent systems are expected to deliver meaningful benefits across sectors, but secure adoption depends on guidance that is practical, scalable, and aligned with real world deployment conditions. CAISI can play a critical role by publishing risk based, modular resources that are adoptable by the full innovation ecosystem, including the small and medium-sized developers that drive much of America's technological leadership.

[4] Appendix.

Sincerely,

Brian Scarpelli
Senior Global Policy Counsel

Kedharnath Sankararaman
Policy Associate

ACT | The App Association
1401 K St NW (Ste 501)
Washington, DC 20005

ACT | The App Association
Policy Principles for AI

(Washington, DC) – Shaping the future of artificial intelligence (AI) has become a priority for policymakers around the world, with nearly 40 AI-related laws enacted globally, more than 80 pending legislative proposals at the state level, and multiple bills at the federal level. To understand and shape rules for this complex and evolving technology, the voice of small businesses, including members of ACT| The App Association, must be at the table.

When we initially released these principles in 2021 we knew that a future where data would drive ever more powerful computers was a reality. We also note that if policymakers were to enact an overwhelming regulatory framework governing the use or development of AI based on what we know about it today, it will likely be fundamentally out of date in the next five to 10 years. With the direct exposure to consumers of generative AI tools like ChatGPT, AI in general jumped to the forefront of the global consciousness. This updated set of recommendations reflects new developments in privacy and data security laws around the world and new learnings about the benefits, risks, and challenges presented by evolving AI tools in use cases from healthcare and education to software development and cybersecurity.

A successful policy approach to AI will align with the following principles:

1. **Understand the Technology:** App Association members have been using AI for decades and even with the more advanced technologies in use today, humans are always responsible for its output. Government intervention designed for some of these AI tools is likely a poor fit for other kinds, especially when they are used in different industries and existing regulatory verticals. Although there is no consensus definition that encompasses the various types of AI in use today, basic taxonomies are nonetheless helpful for anybody seeking an understanding of the technology. Policymakers should understand that AI refers to a wide range of tools, which may be broken down in numerous ways, including as follows:
    - ***Machine learning (ML).*** ML has generally been the most widely-adopted and versatile subset of AI tools. Examples of its use include clinical decision support (CDS) for clinicians to better diagnose health conditions and develop a treatment plan, real-time GPS navigation, inventory management to improve on-shelf availability, and digital tutoring.
    - ***Deep Learning (DL).*** Unlike machine learning, deep learning tools teach themselves new ways of learning from and producing outputs based on the data they analyze. Some current examples of how deep learning is used include drug discovery that avoids some of the enormous costs of human clinical trials, self-driving vehicles, and computer vision (now relying more on deep learning than statistical models).
    - ***Generative AI.*** ChatGPT and DALL-E are examples of generative AI, which refers to a set of ML tools that draw on foundation models (FMs). These technologies can generate content like text compositions, images, videos, or other materials that mimic the style and other attributes of the source material. Unlike reactive DL and ML, generative AI can produce materials that can sometimes appear indistinguishable from genuine versions of them. Some examples of generative AI uses include software coding, product descriptions, and media editing.

2. **Understand How Current Law Applies to the Use of AI**: A wide range of federal, local, and state laws prohibit harmful conduct regardless of whether the use of AI is involved. For example, the Federal Trade Commission (FTC) Act prohibits a wide range of unfair or deceptive acts or practices, and states also have versions of these prohibitions in their statute books. The use of AI does not shield companies from these prohibitions. However, federal and state agencies alike must approach the applicability of these laws in AI contexts thoughtfully and with great sensitivity to the novel or evolving risks AI systems present. Congress and other policymakers must first understand how existing frameworks apply to activities involving AI to avoid creating sweeping new authorities or agencies that awkwardly or inconsistently overlap with current policy frameworks.

3. **Quality Assurance and Oversight:** Policy frameworks should utilize risk-based approaches to ensure that the use of AI aligns with any relevant recognized standards of safety, efficacy, and equity. Small software and device companies benefit from understanding the distribution of risk and liability in building, testing, and using AI tools. Risk management frameworks addressing liability should ensure the appropriate distribution and mitigation of risk and liability across the value chain and appropriately assess how it may apply to distinct categories such as developers, deployers, and users of AI. Specifically, those in the value chain with the ability to minimize risks based on their knowledge and ability to mitigate should have appropriate incentives to do so. Some recommended areas of focus include:
    - Ensuring AI is safe, efficacious, and equitable.
    - Encouraging AI developers to consistently utilize rigorous procedures and enabling them to document their methods and results.
    - Encouraging those developing, offering, or testing AI systems intended for consumer use to provide truthful and easy-to-understand representations regarding intended use and risks that would be reasonably understood by those intended, as well as expected, to use the AI solution.

4. **Thoughtful Design:** Policy frameworks should encourage design of AI systems that are informed by real-world workflows, human-centered design and usability principles, and end-user needs. AI systems should facilitate a transition to changes in the delivery of goods and services, including increases in efficiency, that benefit consumers and businesses. The design, development, and success of AI should leverage collaboration and dialogue among users, AI technology developers, and other stakeholders to have all perspectives reflected in AI solutions.

5. **Access and Affordability:** Policy frameworks should enable products and services that involve AI systems to be accessible and affordable. Significant resources may be required to scale systems. Policymakers should also ensure that developers can build accessibility features into their AI-driven offerings and avoid policies that limit their accessibility options.

6. **Research and Transparency:** Policy frameworks should support and facilitate research and development of AI by prioritizing and providing sufficient funding while also maximizing innovators' and researchers' ability to collect and process data from a wide range of sources. Research on the costs and benefits of transparency in AI should also be a priority and involve collaboration among all affected stakeholders to develop a better understanding of how and under which circumstances transparency mandates would help address risks arising from the use of AI systems.

7. **Modernized Privacy and Security Frameworks:** The many new AI-driven uses for data, including sensitive personal information, raise privacy questions. They also offer the potential for more powerful and granular privacy controls for consumers. Accordingly, any policy framework should address the topics of privacy, consent, and modern technological capabilities as a part of the policy development process. Policy frameworks must be scalable and assure that an individual's data is properly protected, while also allowing the flow of information and responsible evolution of AI. A balanced framework should avoid undue barriers to data processing and collection while imposing reasonable data minimization, consent, and consumer rights frameworks.

8. **Bias:** The potential for harmful bias (whether racial, political, geographic, etc.), as well as errors, will remain one of the more pressing issues with AI systems that utilize machine learning techniques in particular. Stakeholders throughout the chain, including AI deployers, developers, and consumers, should develop and measure adherence to diversity, equity, and inclusion best practices to identify, address, and mitigate the risks of harmful bias that may affect marginalized and underrepresented groups. Regulatory agencies should examine data provenance and bias issues present in the development and uses of AI solutions to ensure that bias in datasets does not result in harm to users or consumers of products or services involving AI, including through unlawful discrimination.

9. **Ethics:** The success of AI depends on ethical use. A policy framework must promote many of the existing and emerging ethical norms for broader adherence by AI technologists, innovators, computer scientists, and those who use such systems. Relevant ethical considerations include:
    - Applying ethics to each phase of an AI system's life, from design to development to use.
    - Maintaining consistency with international conventions on human rights.
    - Prioritizing inclusivity such that AI solutions benefit consumers and are developed using data from across socioeconomic, age, gender, geographic origin, and other groupings.
    - Reflecting that AI tools may reveal extremely sensitive and private information about a user and ensure that laws require the protection of such information.

10. **Education:** Policy frameworks should support education for the advancement of AI, promote examples that demonstrate the success of AI, and encourage stakeholder engagements to keep frameworks responsive to emerging opportunities and challenges.
    - AI users and consumers should be educated as to the use of AI in the service(s) they are using. Policy frameworks should encourage or provide education for AI users and consumers as to how copyright law applies to generative AI and in other contexts involving AI tools, including intellectual property (IP) licensing best practices.
    - Academic education should include curriculum that will advance the understanding of and ability to use AI solutions, including with a focus on understanding and applying ethical frameworks to AI use and development.

11. **Intellectual Property:** The protection of IP rights is critical to the evolution of AI. In developing approaches and frameworks for AI governance, policymakers should be mindful of how current legal protections apply in circumstances involving AI and ensure that compliance measures and requirements do not undercut IP or trade secrets.

About ACT | The App Association: The App Association is a global trade association for small and medium-sized technology companies. Our members are entrepreneurs, innovators, and independent developers within the global app ecosystem that engage with verticals across every industry. We work with and for our members to promote a policy environment that rewards and inspires innovation while providing resources that help them raise capital, create jobs, and continue to build incredible technology.