

May 29, 2025

Attn: Faisal D'Souza, NCO Networking and Information Technology Research and Development National Coordination Office National Science Foundation 2415 Eisenhower Avenue, Alexandria, Virginia 22314

RE: Comments of ACT | The App Association to the Office of Science and Technology Policy on its Request for Information to the Updated National Artificial Intelligence Research and Development Strategic Plan

ACT | The App Association (App Association) appreciates the opportunity to submit views to the National Science Foundation's Networking and Information Technology Research and Development (NITRD) National Coordination Office (NCO) in response to its request for information on behalf of the Office of Science and Technology Policy (OSTP) on updates to the National Artificial Intelligence Research and Development Strategic Plan, which provides guidance to federal agencies to inform the development of regulatory and non-regulatory approaches regarding technologies and industrial sectors empowered or enabled by artificial intelligence (AI), and ways for agencies to reduce barriers to the development and adoption of AI technologies.¹ The App Association supports updating the National Artificial Intelligence Research and Development Strategic Plan to support and facilitate AI research and development by prioritizing and providing sufficient funding while also ensuring adequate incentives (e.g., streamlined availability of data to developers, tax credits) are in place to encourage private and non-profit sector research. Transparency research should be a priority and involve collaboration among all affected stakeholders who must responsibly address the ethical, social, economic, and legal implications that may result from AI applications.

The App Association represents thousands of small business software application development companies and technology firms that create the technologies that drive internet of things (IoT) use cases across consumer and enterprise contexts. Today, the value of the ecosystem the App Association represents—which we call the app economy—is valued at \$1.8 trillion and is responsible for 6.1 million American jobs, while serving as a key driver of the \$8 trillion IoT revolution.² Alongside the world's rapid embrace of mobile technology, our members create the innovative solutions that power

² ACT | The App Association, State of the App Economy (2022), <u>https://actonline.org/wpcontent/uploads/APP-Economy-Report-FINAL.pdf</u>.

ACTonline.org

C 202.331.2130



Actonline.org

¹ https://www.federalregister.gov/documents/2025/04/29/2025-07332/request-for-information-on-thedevelopment-of-a-2025-national-artificial-intelligence-ai-research.

IoT across modalities and segments of the economy. The National Artificial Intelligence Research and Development Strategic Plan, and the efforts of numerous agencies with respect to AI policy and regulation, directly impact the app economy. We support the Administration's goal of sustaining and enhancing America's AI dominance in order to promote human flourishing, economic competitiveness, and national security.³

Al is an evolving constellation of technologies that enable computers to simulate elements of human thinking – learning and reasoning among them. An encompassing term, Al entails a range of approaches and technologies, such as machine learning (ML) and deep learning, where an algorithm based on the way neurons and synapses in the brain change due to exposure to new inputs, allowing independent or assisted decision making. Al-driven algorithmic decision tools and predictive analytics are having, and will continue to have, substantial direct and indirect effects on Americans. Some forms of Al are already in use to improve American consumers' lives today – for example, Al is used to detect financial and identity theft and to protect the communications networks upon which Americans rely against cybersecurity threats.

Moving forward, across use cases and sectors, AI has incredible potential to improve American consumers' lives through faster and better-informed decision making, enabled by cutting-edge distributed cloud computing. As an example, AI-enabled healthcare treatments stand poised to improve disease prevention and conditions, as well as efficiently and effectively treat diseases through automated analysis of X-rays and other medical imaging. AI will also play an essential role in self-driving vehicles and could drastically reduce roadway deaths and injuries. From a governance perspective, AI solutions will derive greater insights from infrastructure and support efficient budgeting decisions.

Today, Americans encounter AI in their lives incrementally through the improvements in computer-based services, typically in the form of streamlined processes, image analysis, and voice recognition (we urge consideration of these forms of AI as "narrow" AI). The App Association notes that this "narrow" AI already provides great societal benefit. For example, AI-driven software products and services revolutionizes the ability of countless Americans with disabilities to achieve experiences in their lives far closer to the experiences of those without disabilities.

Nonetheless, AI also has the potential to raise a variety of unique considerations for policymakers. The App Association appreciates the efforts to develop a policy approach to AI that will bring its benefits to all, balanced with necessary safeguards to protect consumers. To assist the Administration, the App Association offers a comprehensive set of AI policy principles below for consideration that we strongly encourage alignment of the National Artificial Intelligence Research and Development Strategic Plan with the following:

³ Id.

- 1. Al Strategy: Many of the policy issues raised below involve significant work and changes that will impact a range of stakeholders. The cultural, workforce training and education, data access, and technology-related changes associated with AI will require strong guidance and coordination. An AI strategy incorporating guidance on the issues below will be vital to achieving the promise that AI offers to consumers and our economy. We believe it is critical to take this opportunity to encourage civil society organizations and private sector stakeholders to begin similar work. The National Artificial Intelligence Research and Development Strategic Plan is, and should remain, a key part of the United States' overall strategy for global leadership in this critical area of technology.
- 2. Research and Transparency: The National Artificial Intelligence Research and Development Strategic Plan 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.

We appreciate President Trump's acknowledgment, in the January 23 Executive Order establishing the President's Council of Advisors on Science and Technology, of the critical research and innovation enabled by initiatives such as the National Science Foundation's National Artificial Intelligence Research Resource (NAIRR). Launched in 2024, NAIRR provides researchers with access to datasets, models, training, cloud computing, and AI credits to drive groundbreaking advancements in AI applications across defense, healthcare, energy, and other sectors vital to U.S. competitiveness. However, the technology developer-donated credits that support NAIRR will expire at the end of the two-year pilot. While Congress has allocated some funding for the program's administration, NAIRR's continuation depends on congressional appropriations for researcher technology credits. The NAIRR Task Forceformed under the National AI Initiative Act of 2020, signed into law by President Trump—estimated that sustaining NAIRR requires \$2.25 billion in federal appropriations over six years to ensure researchers have the resources needed to develop transformative AI solutions and address society's most pressing challenges. The task force recommended congressional appropriations of \$750 million every two years, and we urge the Administration to incorporate this essential funding into future budget proposals to Congress.

3. **Quality Assurance and Oversight:** The National Artificial Intelligence Research and Development Strategic Plan, and the U.S. approach to Al generally, should utilize risk-based approaches to ensure that the use of Al aligns with any relevant recognized standards of safety and efficacy. Small software and device companies benefit from understanding the distribution of risk and liability in building, testing, and using AI tools. To the extent the strategic plan addresses liability, it should ensure the appropriate distribution and mitigation of risk and liability. 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 and efficacious.
- 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.

The App Association also urges OSTP to align with our recommendations on the roles and interdependencies in the AI value chain, which support the theme of a shared responsibility for safety and efficacy.⁴ In this framework, the App Association proposes clear definitions of stakeholders across the AI value chain, from development to distribution, deployment, and end use; discusses roles for supporting safety, ethical use, and fairness for each of these important stakeholder groups that are intended to illuminate the interdependencies between these actors, thus advancing the shared responsibility concept.

- 4. Thoughtful Design: The National Artificial Intelligence Research and Development Strategic Plan, and the U.S. approach to AI generally, 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 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, Infrastructure, and Affordability: The National Artificial Intelligence Research and Development Strategic Plan, and the U.S. approach to Al generally, should enable products and services that involve Al 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 Al-driven offerings and avoid policies that limit their accessibility options.
- 6. **Modernized Privacy and Security Frameworks:** While the types of data items analyzed by AI and other technologies are not new, this analysis will provide greater potential utility of those data items to other individuals, entities, and machines. Thus, there are many new uses for, and ways to analyze, the

⁴ This framework is included as **Appendix A** to this comment.

collected data. This raises privacy issues and questions surrounding consent to use data in a particular way (e.g., research, commercial product/ service development). It also offers the potential for more powerful and granular access controls for consumers. Accordingly, the National Artificial Intelligence Research and Development Strategic Plan, and the U.S. approach to AI generally, should address the topics of privacy, consent, and modern technological capabilities as a part of the policy development process. Risk management policy frameworks must be scalable and ensure that an individual's data is properly protected, while also allowing the flow of information and responsible evolution of AI. This information is necessary to provide and promote high-quality AI applications. Finally, with proper protections in place, policy frameworks should also promote data access, including open access to appropriate machine-readable public data, development of a culture of securely sharing data with external partners, and explicit communication of allowable use with periodic review of informed consent.

- 7. **Collaboration and Interoperability:** The National Artificial Intelligence Research and Development Strategic Plan, and the U.S. approach to AI generally, should enable eased data access and use through creating a culture of cooperation, trust, and openness among policymakers, AI technology developers and users, and the public.
- 8. **Bias:** The bias inherent in all data, as well as errors, will remain one of the more pressing issues with AI systems that utilize machine learning techniques in particular. 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. The National Artificial Intelligence Research and Development Strategic Plan, and the U.S. approach to AI generally, should:
 - Require the identification, disclosure, and mitigation of bias while encouraging access to databases and promoting inclusion and diversity.
 - Ensure that data bias does not cause harm to users or consumers.
- 9. Education: The National Artificial Intelligence Research and Development Strategic Plan, and the U.S. approach to AI generally, 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.
 - Consumers should be educated as to the use of AI in the service they are using.
 - Academic education should include curriculum that will advance the understanding of and ability to use AI solutions.
- 10. **Intellectual Property:** The protection of intellectual property (IP) rights is critical to the evolution of AI. In developing approaches and frameworks for AI governance, the National Artificial Intelligence Research and Development

Strategic Plan, and the U.S. approach to AI generally, should ensure that compliance measures and requirements do not force the disclosure or violation or IP or trade secrets.

11. **Tax Policy and R&D Incentives:** The Strategic Plan must recognize that tax policy is innovation policy. Since 2022, changes to Section 174 of the Internal Revenue Code have forced businesses to amortize, rather than immediately deduct, research and development (R&D) expenses. This shift is devastating for early-stage software developers and small AI firms whose most significant cost—engineering talent—is now effectively penalized.

The App Association strongly urges OSTP to endorse restoring **full and immediate expensing of R&D costs** as part of its innovation strategy. This policy is central to ensuring small businesses can reinvest in product development, scale operations, and remain globally competitive.

The policy issues raised by the National Artificial Intelligence Research and Development Strategic Plan involve significant work and changes that will impact a range of stakeholders. The cultural, workforce training and education, data access, and technology-related changes associated with AI will require strong guidance and coordination across U.S. federal agencies. The App Association supports the development of national AI strategies for federal agencies, which will be vital to achieving the promise that AI offers to consumers and entire economies.

Noting our general support for the current National Artificial Intelligence Research and Development Strategic Plan, we offer the following suggested revisions:

- Alignment with Other Leading Federal Policies for AI: The National Artificial Intelligence Research and Development Strategic Plan should align with other federal efforts to develop AI policy, such as the National Institute of Standards and Technology's (NIST) Artificial Intelligence Risk Management Framework, a policy that was developed in close collaboration with the private sector, academia, and others for voluntary use with the goal of improving the ability to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems.⁵
- Require Agencies to Advance Thoughtful Design Principles Across Al Use Cases: The National Artificial Intelligence Research and Development Strategic Plan should require design of Al systems informed by real-world workflows, human-centered design and usability principles, and end-user needs. Al systems should facilitate a transition to changes in the delivery of goods and services that benefit consumers and businesses. The design, development, and success of Al should leverage collaboration and dialogue among users, Al technology developers, and other stakeholders in order to have all perspectives reflected in

⁵ <u>https://www.nist.gov/itl/ai-risk-management-framework.</u>

Al solutions. As this concept must run across sectors and Al use cases, the National Artificial Intelligence Research and Development Strategic Plan should incorporate guidance for agencies to advance thoughtful design principles through their approaches and actions related to Al.

- Advance an Appropriate Distribution of Responsibility in the Al Value Chain: Technology developers, deployers, end users, and other stakeholders will all benefit from a shared understanding of the distribution of risk and responsibility in the building, testing, explaining, and using of Al tools. The National Artificial Intelligence Research and Development Strategic Plan should advance the appropriate distribution of responsibility to ensure the appropriate distribution and mitigation of risk and liability (namely, 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). We urge for alignment with the App Association's *Al Roles and Interdependencies Framework*, which describes a typical Al value chain and its actors and recommends steps for each actor to take in addressing safety, transparency, and explainability based on their knowledge, intended uses, and ability to mitigate known harms (mapped to the NIST Al Risk Management Framework's functions).⁶
- Support the Development of, and Access to, Open Standards Needed to Drive U.S. Leadership in AI: The National Artificial Intelligence Research and Development Strategic Plan should support the development and use of voluntary consensus standards that concern AI application, and should be updated to advance open standards, consistent with OMB-A119 ("Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities").⁷

It is also critical that the United States should ensure that AI standards are able to be used by American small business innovators by promoting a balanced approach to standard-essential patent (SEP) licensing. AI technical standards, built on contributions through an open and consensus-based process, bring immense value to consumers by promoting interoperability while enabling healthy competition between innovators and often include patented technology. When an innovator gives its patented technology to a standard, this can represent a clear path to reward in the form of royalties from a market that likely would not have existed without the standard being widely adopted. To balance this potential with the need for access to the patents that underlie the standard, many standardsetting organizations (SSOs) require holders of patents on standardized technologies to license their patents on fair, reasonable, and non-discriminatory

⁶ The App Association's AI Roles & Interdependency Framework is included as **Appendix A**, and is also available at <u>https://actonline.org/wp-content/uploads/ACT-AI-Roles-Interdependencies-Framework-final-text-May-2024-UK-English.pdf</u>.

⁷ <u>https://www.nist.gov/system/files/revised_circular_a-119_as_of_01-22-2016.pdf.</u>

(FRAND) terms. FRAND commitments prevent the owners of patents used to implement the standard from exploiting the unearned market power that they otherwise would gain as a consequence of the broad adoption of a standard. Once patented technologies are incorporated into standards, manufacturers are compelled to use them to maintain product compatibility. In exchange for making a voluntary FRAND commitment with an SSO, SEP holders gain the ability to obtain reasonable royalties from a large number of standard implementers that might not have existed absent the standard. Without the constraint of a FRAND commitment, SEP holders would have the same power as a monopolist that faces no competition.

Unfortunately, a handful of owners of FRAND-committed SEPs are flagrantly abusing their unique position by reneging on those promises with unfair, unreasonable, or discriminatory licensing practices. These practices, under close examination by antitrust and other regulators in many jurisdictions, not only threaten healthy competition and unbalance the standards system but also impact the viability of new markets such as AI. This amplifies the negative impacts on small businesses because they can neither afford years of litigation to fight for reasonable royalties nor risk facing an injunction if they refuse a license that is not FRAND compliant.

SEP licensing abuse now constitutes a national economic and security threat.⁸ Abusive SEP holders are exploiting ambiguous FRAND terms to overcharge or exclude U.S. companies from essential standards, undermining innovation and disrupting critical supply chains across industries like automotive, telecommunications, and healthcare. Foreign actors—particularly Chinese statebacked enterprises like Huawei—have used SEP portfolios and patent pools to lock U.S. businesses out of global markets, extract supra-FRAND royalties, and weaponize foreign courts to impose coercive global licensing terms. These abuses not only destabilize the integrity of U.S. technical standards but also jeopardize the resilience of American manufacturing and the nation's ability to maintain secure, interoperable systems in strategic sectors. SEP abuse has become a well-practiced tactic that empowers adversaries to hold technical standards hostage and shut down U.S. supply chains, demanding urgent policy action to restore fair access and protect American leadership in emerging technologies.

Therefore the National Artificial Intelligence Research and Development Strategic Plan should align with the following principles to support American leadership and innovation in AI standards and R&D:

⁸ ACT | The App Association, Memorandum RE: Promoting a Competitive Standard-Essential Patent Landscape (2024), <u>https://actonline.org/wp-content/uploads/App-Assn-Transition-Memo-re-SEP-and-US-Econ-Nat-Security-3-Jan-2025.pdf</u>.

- The FRAND Commitment Means All Can License A holder of a FRAND-committed SEP must license that SEP to all companies, organizations, and individuals who use or wish to use the standard on FRAND terms.
- Prohibitive Orders on FRAND-Committed SEPs Should Only Be Allowed in Rare Circumstances – Prohibitive orders (federal district court injunctions and U.S. International Trade Commission exclusion orders) should not be sought by SEP holders or allowed for FRANDcommitted SEPs except in rare circumstances where monetary remedies are not available.
- FRAND Royalties A reasonable rate for a valid, infringed, and enforceable FRAND- committed SEP should be based on the value of the actual patented invention itself, which is separate from purported value due to its inclusion in the standard, hypothetical uses downstream from the smallest saleable patent practicing unit, or other factors unrelated to invention's value.
- FRAND-committed SEPs Should Respect Patent Territoriality Patents are creatures of domestic law, and national courts should respect the jurisdiction of foreign patent laws to avoid overreach with respect to SEP remedies. Absent agreement by both parties, no court should impose global licensing terms on pain of a national injunction.
- The FRAND Commitment Prohibits Harmful Tying Practices While some licensees may wish to get broader licenses, a SEP holder that has made a FRAND commitment cannot require licensees to take or grant licenses to other patents not essential to the standard, invalid, unenforceable, and/or not infringed.
- The FRAND Commitment Follows the Transfer of a SEP As many jurisdictions have recognized, if a FRAND-committed SEP is transferred, the FRAND commitments follow the SEP in that and all subsequent transfers.

The App Association appreciates the Administration's consideration of the above views. We urge OSTP to contact the undersigned with any questions or ways that we can assist moving forward.

This document is approved for public dissemination. The document contains no business-proprietary or confidential information. Document contents may be reused by the government in developing the 2025 National AI R&D Strategic Plan and associated documents without attribution.

Sincerely,

Brian Scarpelli Senior Global Policy Counsel

dolvaly

Kedharnath Sankararaman Policy Associate

ACT | The App Association 1401 K St NW (Ste 501)

Washington, DC 20005 202-331-2130

Overview: Artificial Intelligence (AI), especially generative AI, is already a powerful tool for consumers and companies. App Association small business members have a vital role in advancing AI's positive impacts by identifying new and novel opportunities where the responsible use of AI can solve expensive problems and provide new efficiencies for consumers and businesses.

While AI capabilities are already positively transforming American society, the App Association also recognizes that the same capabilities raise unique challenges that the government, private sector, and others have an important role in addressing across development, distribution, deployment, and end use phases. The App Association has worked proactively with its diverse and innovative community of small businesses to develop this consensus taxonomy, which describes the roles and interdependencies of various actors in the value (or supply) chain of AI solutions. These roles include several AI/ML developer subgroups, deploying organizations, end users, standard-setting organizations, certification and test beds, specialty boards and licensing bodies, and academic institutions. Many of these stakeholders map to actors in the National Institute for Standards and Technology's (NIST's) AI Risk Management Framework (RMF), which we indicate on the far right of the matrix below.

While the App Association has created comprehensive policy principles for Al governance, there we have several recommendations from this roles and interdependencies document. The App Association recommends: (1) that requirements placed on small business Al developers and users be based on demonstrated harms; (2) the leveraging of a risk-based approach to Al harm mitigation where the level of review, assurance, and oversight is proportionate to those demonstrated harms; and (3) that those in Al value chains with the ability to minimize risks based on their knowledge and ability have appropriate responsibilities and incentives to do so.

Stakeholder Group	Definition	Roles	NIST AI RMF
Al/ML Developers	Someone who designs, codes, researches, or produces an AI/ML system or platform for internal use or for use by a third party. See below for defined Subgroups of this Stakeholder Group along with recommendations specific to that Subgroup.	 Informing deployers and users of data requirements/definitions, intended use cases/populations and applications (e.g., disclosing sufficient detail allowing providers to determine when an Al-enabled tool should reasonably apply to the individual they are treating), including whether the Al/ML tools are intended to augment human work versus automate workflows, and status of/compliance with all applicable legal and regulatory requirements. Prioritizing safety, effectiveness, transparency, and data privacy and security from the earliest stages of design, leveraging (and, where appropriate, updating) existing Al/ML guidelines on research and ethics, leading standards, and other resources. Employing algorithms that produce repeatable results and, when feasible, are auditable, and make decisions that comply with relevant sector-specific requirements. Using risk management approaches that scale to the potential likely harms posed in intended use scenarios to support safety, protect privacy and security, avoid harmful outcomes. Providing information that enables those further down the value chain can assess the quality, performance, and utility of Al/ML tools. Aligning with relevant ethical obligations and international conventions on human rights and supporting the development of new ethical guidelines to address emerging issues. 	Al Deployment; Operation and Monitoring; Test, Evaluation, Verification, and Validation (TEVV); Human Factors; Domain Expert; Al Impact Assessment; Governance and Oversight

Stakeholder Subgroup	Definition	Roles	NIST RMF Actor Tasks
Foundation Model Developer	Someone who creates or modifies large and generalizable machine learning models that can be	 Building on the cross-Al/ML Developer roles noted above: Assessing what efficacy and safety issues might be present in its Foundation Model, 	Al Deployment; Operation and Monitoring; Test, Evaluation, Verification, and Validation (TEVV); Human Factors;

Stakeholder Subgroup	Definition	Roles	NIST RMF Actor Tasks
	used/adapted for various downstream tasks and applications, such as natural language processing, computer vision, or software development.	 and documenting steps taken to mitigate those issues in its Transparency Documentation (e.g., Transparency Notes, System Cards and product documentation). Providing clear guidance on (1) how to use and adapt its Foundation Model for various foreseeable downstream tasks and applications, and (2) what limitations or risks may arise from doing so based on challenges discovered during testing and deployment. 	Domain Expert; Al Impact Assessment; Governance and Oversight
Al Platform Developer	Someone who leverages existing foundation models and builds an industry- agnostic platform that enables other developers to access, customize, and deploy these models for various use cases and applications, such as natural language processing, computer vision, and/or software development.	 Building on the cross-Al/ML Developer roles noted above: Testing for, identifying, and mitigating safety issues that may arise from using or modifying existing foundation models for its Al Platform, and documenting these issues and steps taken to address them in its transparency documentation (e.g., transparency notes, system cards and product documentation). 	Al Deployment; Operation and Monitoring; Test, Evaluation, Verification, and Validation (TEVV); Human Factors; Domain Expert; Al Impact Assessment; Governance and Oversight
Use Case Al Platform Developer	Someone who creates or uses AI-powered platforms that are tailored for a particular domain or sector. These platforms may leverage foundation models (or other types of machine learning models or solutions), such as AI platforms, that are suitable for domain-specific	 Building on the cross-Al/ML Developer roles noted above: Meeting specific requirements and standards of the domain to address unique accuracy, efficacy, explainability, and compliance needs. Testing for, identifying, and mitigating any efficacy and safety issues that may affect domain-specific outcomes or performance needs, and documenting these issues and the steps it has taken to address them in its transparency 	Al Deployment; Operation and Monitoring; Test, Evaluation, Verification, and Validation (TEVV); Human Factors; Domain Expert; Al Impact Assessment; Governance and Oversight

Stakeholder Subgroup	Definition	Roles	NIST RMF Actor Tasks
	problems and data sources.	documentation (e.g., transparency notes, system cards and product documentation).	
Al Solution Developer	Someone who creates complete digital tools and technologies for a domain. They may build or incorporate AI solutions with both use case AI platforms, which are specialized for the domain, and AI platforms, which are more general and adaptable for various use cases and applications.	 Building on the cross-Al/ML Developer responsibilities noted above: Specifying appropriate uses for its solution to avoid amplifying safety issues that may exist in the underlying foundation models, AI platforms, or domain-specific AI platforms. Designing user interfaces to enable an end user to safely and effectively act upon the output of the tool, such as providing explanations, feedback mechanisms, or human oversight options, providing clear documentation to Deploying Organizations and Users to help them avoid safety issues. 	Al Deployment; Operation and Monitoring; Test, Evaluation, Verification, and Validation (TEVV); Human Factors; Domain Expert; Al Impact Assessment; Governance and Oversight

Stakeholder Group	Definition	Roles	NIST AI RMF Actor Tasks
Deploying Organization	Someone who is deploying solutions built by AI Solution Developers. They may also have their own internal IT staff that employ use case AI platforms or general AI platforms to develop their own custom AI solutions.	 Respecting that managing Al/ML risks will be more challenging for small to medium-sized organizations depending on their capabilities and resources: Adopting Al/ML Developer instructions for use, specifying appropriate uses for Users through governance policies to avoid safety issues that may exist in the underlying foundation models, Al platforms, or use case Al platforms. Developing and leveraging solutions that augment efficiencies in automation, facilitate administrative simplification/reduce workflow burdens, and are fit for purpose. Setting organization policy/designing workflows to reduce the likelihood that a User will act upon the output 	Al Deployment; Operation and Monitoring; Domain Expert; Al Impact Assessment; Procurement; Governance and Oversight

Stakeholder Group	Definition	Roles	NIST AI RMF
		 of the tool in a way that would cause efficacy or safety issues (tailored explanations, feedback mechanisms, and/or human oversight options). Assuring that AI/ML systems allow for the individualized assessment of domain-specific circumstances and flexibility to override automated decisions, ensuring that use of AI/ML does not improperly reduce or withhold intended benefits or inappropriately override human judgement. Developing support mechanisms for the use of AI/ML by providers based on validation, aligning with decision-making processes familiar to the domain and high-quality evidence. Developing organizational guidance on how the AI solution should and should not be used. Creating engagement pathways to support dialogue with AI use case developers, AI solution developers, or any other applicable AI/ML developer, to enable ongoing updates to address evolving risks and benefits of AI solution uses. Creating risk-based, tailored communications and engagement plans to enable easily understood explanations to customers about how the AI solution was developed, its performance and maintenance, and how it aligns with the latest best practices and regulatory requirements 	
Al End Users	Someone who directly interacts with or benefits from the AI solutions that are built by AI Solution Developers or by the internal IT staff of the Deploying Organization.	 Respecting that managing AI/ML risks will be more challenging for small to medium-sized organizations depending on their capabilities and resources: Aligning with consensus AI/ML definitions, present-day and future AI/ML solutions, the future of AI/ML changes and trends. Taking required training and incorporating employer guidance about use of AI/ML solutions. Documenting (through automated processes or otherwise) and reporting any issues or feedback to the 	AI Deployment; Operation and Monitoring; Domain Expert; AI Impact Assessment; Procurement; Governance and

Stakeholder Group	Definition	Roles	NIST AI RMF
		 developer, such as errors, vulnerabilities, or harms (where AI/ML's use is known by the User). Ensuring there is appropriate review of the output or recommendations from each AI solution prior to acting on it to make decisions, if relevant (where AI/ML's use is known by the User). Raising awareness of and acting according to customers' rights and choices when using AI solutions, such as consent, access, correction, or deletion of their personal data. 	Oversight; Human Factors
Standard-Setting Organizations	An organization whose primary function is developing, coordinating, promulgating, revising, amending, reissuing, interpreting, or otherwise contributing to the usefulness of technical standards to those who employ them.	 Developing and promoting adoption of international voluntary/non-regulatory consensus standardized approaches and resources to steward a shared responsibility approach to technology standards that include or are otherwise related to AI. 	Human Factors; Domain Expert; Al Impact Assessment; Governance and Oversight
Certification Bodies & Test Beds	A certification body is a third-party organization that assures the conformity of a product, process or service to specified requirements. A test bed is a platform for conducting rigorous, transparent, and replicable testing of scientific theories, computing tools, and new technologies to a standard.	 Creating and making available transparent and reliable processes for the assurance of conformity to voluntary AI standards. Creating and making available voluntary sandbox environments to help evaluate the usability and performance of AI/ML-based high-performance computing applications to advance the understanding of how reliable and efficacious AI, and to provide an appropriate assurance of reliability and efficacy. 	Test, Evaluation, Verification, and Validation (TEVV); Human Factors; Domain Expert; Al Impact Assessment; Governance and Oversight
Accrediting and Licensing Bodies, Specialty Societies and Boards	Accrediting and licensing bodies are governing authorities that establish the suitability of any participating certification body. Notably, state-level boards serve	 Based on needs and expertise, developing and setting the standard of practice/behavior and ethical guidelines to address emerging issues with the use of AI/ML in the relevant domain. Identifying the most appropriate uses of AI-enabled technologies and developing and disseminating 	Test, Evaluation, Verification, and Validation (TEVV); Human

Stakeholder Group	Definition	Roles	NIST AI RMF Actor Tasks
	this purpose for certain professions to standards set by each state. Specialty societies are organizations for specialized professionals.	guidance and education on the responsible deployment of AI/ML, both generally and for specialty-specific uses.	Factors; Domain Expert; Al Impact Assessment; Governance and Oversight
Academic Education Institutions	Tertiary educational institutions, professional schools, or forms a part of such institutions, that teach and award professional degrees.	 Developing and teaching curriculum that will advance understanding of and ability to use AI/ML solutions responsibly, which should be assisted by inclusion of data scientists and engineers as instructors as needed. Developing curriculum to advance the understanding of data science research to help inform ethical bodies. 	Human Factors; Domain Expert; Al Impact Assessment