BSA | The Software Alliance Comments on the White House Office of Science and Technology Policy's Request for Information on National Priorities for Artificial Intelligence
July 6, 2023

Document ID: OSTP-TECH-2023-0007-0001

BSA | The Software Alliance appreciates the opportunity to provide comments in response to the White House Office of Science and Technology Policy's Request for Information on National Priorities for Artificial Intelligence.

BSA is the leading advocate for the global software industry.[1] Our members are enterprise software companies that create business-to-business technologies that help other businesses innovate and grow.[2] For example, BSA members provide tools including cloud storage and data processing services, customer relationship management software, human resource management programs, identity management services, and collaboration software. BSA members are on the leading edge of providing AI-enabled products and services, and tools used by others in the development of AI systems and applications. As a result, they have unique insights into the technology's tremendous potential to spur digital transformation and the policies that can best support the responsible use of AI.

BSA's views are informed by our recent 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 to identify risks of AI bias and highlights corresponding risk mitigation best practices. BSA has testified before the United

[1] BSA's members include: Adobe, Alteryx, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, Databricks, DocuSign, Dropbox, Elastic, Graphisoft, IBM, Informatica, Juniper Networks, Kyndryl, MathWorks, Microsoft, Okta, Oracle, Prokon, PTC, Rubrik, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Unity Technologies, Inc., Workday, Zendesk, and Zoom Video Communications, Inc.

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

[3] *See* BSA | The 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.

States Congress and the European Parliament on the Framework and its approach to mitigating AI-related risks.[4] Our extensive experience on these issues informs our response to your questions below.

***Question 1**: What specific measures – such as standards, regulations, investments, and improved trust and safety practices – are needed to ensure that AI systems are designed, developed, and deployed in a manner that protects people's rights and safety? Which specific entities should develop and implement these measures?*
***Question 2**: How can the principles and practices for identifying and mitigating risks from AI, as outlined in the Blueprint for an AI Bill of Rights and the AI Risk Management Framework, be leveraged most effectively to tackle harms posed by the development and use of specific types of AI systems, such as large language models?*
***Question 3**: Are there forms of voluntary or mandatory oversight of AI systems that would help mitigate risk? Can inspiration be drawn from analogous or instructive models of risk management in other sectors, such as laws and policies that promote oversight through registration, incentives, certification, or licensing?*

Response to Questions 1-3: BSA supports a range of measures that can improve the responsible development and deployment of AI. Ultimately, the success of AI will depend on public trust and confidence in the technology.

Legislation for High-Risk AI Systems. BSA supports legislation requiring organizations to conduct impact assessments if they develop or deploy high-risk AI systems.

Impact assessments are important accountability tools and can be used to help AI developers and deployers identify and mitigate risks throughout the lifecycle of an AI system. Impact assessments are already used today in a range of other fields, including environmental protection and data protection. BSA supports leveraging these existing tools to create important guardrails around the development and deployment of high-risk AI systems.

There are a wide array of AI use cases, many of which are mundane and pose little risks to individuals, such as filtering background noise in video calls or identifying suspicious IP addresses to enhance cybersecurity. Focusing on high-risk uses

[4] *See* Testimony of Victoria Espinel, Public Hearing on AI & Bias, Special Committee on Artificial Intelligence in a Digital Age, European Parliament, Nov. 30, 2021, *available at* https://www.europarl.europa.eu/cmsdata/244265/AIDA_Verbatim_30_November_2021_EN.pdf; Testimony of Aaron Cooper, Task Force on Artificial Intelligence: Beyond I, Robot: Ethics, Artificial Intelligence and the Digital Age, before the House Financial Services Committee (Oct. 13, 2021), *available at* https://www.congress.gov/117/meeting/house/114125/witnesses/HHRG-117-BA00-Wstate-CooperA-20211013.pdf.

allows policymakers to address those situations that are most likely to result in consequential decisions for individuals without unduly restricting low-risk activities.

Importantly, such legislation should recognize that different companies play different roles in developing and using AI systems. Both the company that develops a high-risk AI system and the company that uses that system should have obligations to ensure responsible AI innovation, but those obligations should be tailored to their different roles in the ecosystem.[5] For example, developers that design an AI system are well-positioned to have access to information about the type of data used to train an AI system, the system's known limitations, and the system's intended use cases. In contrast, a deployer using an AI system is best-positioned to have access to information regarding the specific ways in which it uses that system. Any policies focused on supporting AI accountability should reflect these different roles and assign obligations accordingly.

An important feature of impact assessments for high-risk AI systems is that they facilitate documentation of key aspects of those AI systems. The relevant documents are important reference points for understanding the operation of AI systems and will be different for developers that design an AI system than for deployers using an AI system.

Developers of high-risk AI systems should maintain documentation for a reasonable time period in light of the intended use regarding:

- *The intended purpose of the AI system;*
- *Known limitations of the AI system;*
- *Known, likely, and specific high risks that could occur and steps taken to mitigate those risks;*
- *An overview of the data used to train the AI system; and*
- *A summary of how the AI system was evaluated prior to sale.*

Deployers of high-risk AI systems should maintain documentation for a reasonable time period in light of the intended use regarding:

- The purpose for which the deployer intends to use the AI system;
- Transparency measures, including notices to impacted individuals about the AI system's use;
- A summary of how the AI system is evaluated, if applicable;
- Known, likely, and specific high risks that could occur and steps taken to mitigate those risks; and
- Post-deployment monitoring and user safeguards, if applicable.

[5] *See* BSA | The Software Alliance, AI Developers and Deployers: An Important Distinction, *available at* https://www.bsa.org/files/policy-filings/03162023aidevdep.pdf.

Requiring companies to document how they have assessed and mitigated risks helps ensure AI systems are developed and deployed responsibly.

Use of Risk Management Frameworks. BSA's AI Risk Management Framework was published in 2021. As discussed above, BSA's Framework identifies specific practices that AI developers and deployers can implement to ensure that AI is developed and used responsibly. For example, the BSA Framework highlights a range of issues that impact assessments should address, including identifying fairness metrics that will be used to assess bias in the AI system, ensuring that senior leadership has been briefed on potential high-risk AI systems, scrutinizing training data for bias, and documenting how testing was performed.

BSA has also strongly supported NIST's work to develop the AI Risk Management Framework (AI RMF), in line with our longstanding recognition that risk management is a key component of promoting trust in AI. The NIST AI RMF identifies specific practices that can be implemented to develop and deploy trustworthy AI.[6] Both the NIST and BSA frameworks recommend:

- Consultation with a diverse group of stakeholders;
- Establishing processes to identify, assess, and mitigate risks;
- Defining individual roles and responsibilities to people throughout an organization;
- Identifying metrics for evaluation;
- Evaluating fairness and bias;
- Maintaining post-deployment feedback mechanisms; and
- Establishing detailed plans for responding to incidents.

BSA recently prepared a crosswalk between the two frameworks, which illustrates the significant alignment between the two approaches.[7] These frameworks are appropriately voluntary, and the flexibility to utilize different components allows organizations to tailor the frameworks to their business needs.

Increased Investment in Research and Development. BSA supports increased investment in research and development to help promote the development of trustworthy AI, spark innovation, ensure economic competitiveness, and maintain America's historical scientific and technological edge. Efforts to establish shared resources for advanced AI research, such as the National AI Research Resource (NAIRR), offer a path for leveraging the unique capabilities of the public, private, and academic sectors. By aiming to provide AI researchers with access to computational resources and high-quality data, the NAIRR has the potential to democratize and enhance America's AI capabilities. The long-term success of the

[6] *See* National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework, *available at* https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.

[7] *See* BSA | The Software Alliance, Crosswalk Between BSA Framework to Build Trust in AI and NIST AI Risk Management Framework, *available at* https://www.bsa.org/files/policy-filings/04122023aiframeworknistcrosswalk.pdf.

NAIRR can be enhanced by encouraging and enabling research partnerships between universities and the public and private sectors. The United States should also advance the establishment of the NAIRR as part of the effort to develop a broader understanding of how to address sociotechnical issues implicated by AI. These investments will help facilitate responsible AI development and help the United States maintain global leadership on AI.

Question 5*: How can AI, including large language models, be used to generate and maintain more secure software and hardware, including software code incorporating best practices in design, coding and post deployment vulnerabilities?*
Question 6*: How can AI rapidly identify cyber vulnerabilities in existing critical infrastructure systems and accelerate addressing them?*

Response to Questions 5-6: AI has significant potential to improve cybersecurity and is already improving secure software development and maintenance.

For example, AI can be used for code analysis, automated bug detection, secure coding suggestions and security-aware coding assistance, threat modeling, automated security testing, automated vulnerability scanning, vulnerability prediction, automated patching, penetration testing, anomaly detection, intrusion detection, security analytics, user behavior analysis, threat intelligence, and automated patch management.

These uses can improve software security broadly. As described in the BSA Framework for Secure Software, which is heavily cited by the National Institute of Standards and Technology's Secure Software Development Framework, the three functions of software security are secure development, secure capabilities, and secure lifecycle.[8] Each function can be improved by integrating AI.

Moreover, the speed of AI increases its ability to improve cybersecurity. AI systems are well suited to rapidly identify vulnerabilities and expedite incident response. That is because AI moves at the speed of data, which is much faster than the speed of human fingers on a keyboard.

Using AI systems in these ways can free up human experts to focus on higher value activities, further improving cybersecurity.

Question 9*: What are the opportunities for AI to enhance equity and how can these be fostered? For example, what are the potential benefits for AI in enabling broadened prosperity, expanding economic and educational opportunity, increasing access to services, and advancing civil rights?*
Question 10*: What are the unique considerations for understanding the impacts of AI systems on underserved communities and particular groups, such as minors and*

[8] BSA | The Software Alliance, The BSA Framework for Secure Software: A New Approach to Securing the Software Lifecycle (Sept. 2020), *available at* https://www.bsa.org/files/reports/bsa_framework_secure_software_update_2020.pdf.

people with disabilities? Are there additional considerations and safeguards that are important for preventing barriers to using these systems and protecting the rights and safety of these groups?
***Question 12**: What additional considerations or measures are needed to assure that AI mitigates algorithmic discrimination, advances equal opportunity, and promotes positive outcomes for all, especially when developed and used in specific domains (e.g., in health and human services, in hiring and employment practices, in transportation)?*
***Question 17**: What will the principal benefits of AI be for the people of the United States? How can the United States best capture the benefits of AI across the economy, in domains such as education, health, and transportation? How can AI be harnessed to improve consumer access to and reduce costs associated with products and services? How can AI be used to increase competition and lower barriers to entry across the economy?*

Response to Questions 9-10, 12, and 17: AI offers immense benefits to individuals, businesses, and society. These benefits include improving healthcare services, responding to natural disasters, optimizing manufacturing, refining weather and climate forecasts, increasing power grid reliability, improving cybersecurity, and making transit more efficient.[9] For example, AI has been leveraged in efforts to model the structure of the COVID-19 virus and accelerate discovery of treatment options. AI can also be used to analyze geospatial and weather data along with historic patterns to predict where hurricanes might cause the most damage. After a weather incident, AI can be used to analyze satellite photos of the area to identify damage. Rescuers and aid groups can then use this information to distribute aid more efficiently.[10] AI can also be used by businesses to monitor their fleet of vehicles to detect when potential safety issues could arise, allowing them to perform predictive maintenance.[11]

AI also has the potential to make services more inclusive by providing the disabled with critical tools and enabling more efficient access to social services. For example, the Seeing AI app helps blind and visually impaired people navigate the world by providing auditory descriptions of objects in photographs they take with their smartphones.[12] AI is also improving education, for example, by giving teachers access to math resources in seconds, including lesson plans and

[9] *See* BSA | The Software Alliance, Artificial Intelligence in Every Sector, *available at* https://www.bsa.org/files/policy-filings/06132022bsaaieverysector.pdf.
[10] *See* World Economic Forum, Natural Disasters Are Increasing in Frequency and Ferocity. Here’s How AI Can Come to the Rescue, Jan. 14, 2020, *available at* https://www.weforum.org/agenda/2020/01/natural-disasters-resilience-relief-artificial-intelligence-ai-mckinsey/.
[11] *See* https://www.oracle.com/artificial-intelligence/anomaly-detection/#industry.
[12] *See* https://www.microsoft.com/en-us/ai/seeing-ai.

teaching strategies for students with varying levels of preparation and ability.[13] AI is also improving access to critical services. For example, using an AI-assisted interactive voice response menu, the United Way is able to route inbound calls and texts from people using their 211 system to gain access to essential housing, financial assistance, food, childcare, and transportation.[14]

Although AI can enhance inclusivity, there is also a potential for bias to emerge in AI systems. As outlined in the BSA Framework, bias could occur at multiple stages in the development and deployment of an AI system. For example, with respect to training data, there are risks of perpetuating historical biases reflected in the data or of sampling bias – where the data is misrepresentative of the population in which it will be used. The process of selecting the input variables (i.e., features) that the model will weigh as it is being trained is another critical decision point that can introduce bias. Even when sensitive demographic data is excluded, bias may be introduced if the system relies on features that are closely correlated to those traits, called proxies. Bias can also arise in various ways after a system has been deployed, including when the data used to train or evaluate an AI system differs materially from the population the system encounters when it is deployed.

One way to identify and address these risks in high-risk systems is conducting an impact assessment. For example, the BSA Framework recommends that in the data acquisition phase of the design of an AI model, companies should evaluate the representativeness of the data as part of conducting an impact assessment. To do so, a company can compare the demographic distribution of training data to the population in which the system will be deployed and assess whether there is sufficient representation of subpopulations that are likely to interact with the system. To mitigate issues that arise, companies can consider rebalancing the dataset with additional data or synthetic data, which involves oversampling data from underrepresented groups.

Similarly, in the data preparation and model definition phase, an impact assessment could include documenting a potential correlation between selected features and sensitive demographic attributes. For features that closely correlate to a sensitive class, companies can document the relevance to the target variable and the rationale for its inclusion in the model, consistent with laws prohibiting discriminatory actions. These practices, along with other measures, can help to prevent algorithmic discrimination while enabling innovation.

Question 11*: How can the United States work with international partners, including low- and middle-income countries, to ensure that AI advances democratic*

[13] *See* https://www.ibm.com/ibm/responsibility/initiatives/activitykits/teacheradvisor/#:~:text=Teacher%20Advisor%20With%20Watson%20is,planning%20effective%20and%20aligned%20lessons.

[14] *See* https://customers.twilio.com/2156/unitedway/.

values and to ensure that potential harms from AI do not disproportionately fall on global populations that have been historically underserved?

Response to Question 11: The United States can take several steps to ensure responsible AI innovation globally. The United States should work with international partners to encourage organizations in their countries to adopt the NIST AI RMF. In doing so, the United States and its partners can ensure that trustworthy AI practices are implemented globally, including in low- and middle-income countries. The United States should also continue implementing the joint roadmap for cooperation on AI issued by the European Union and the United States as part of the U.S.-EU Trade and Technology Council (TTC) meetings. The TTC's work on developing shared terminology, standards, and risk measurement techniques will provide valuable contributions to AI development and policy discussions. U.S. engagement in international standards-setting bodies will also continue to be important, as AI standards are currently under development in the International Organization for Standardization. Finally, the United States should work in the G7 to develop shared approaches on AI, including elevating frameworks like the NIST AI RMF.

***Question 20**: What are potential harms and tradeoffs that might come from leveraging AI across the economy? How can the United States promote quality of jobs, protect workers, and prepare for labor market disruptions that might arise from the broader deployment of AI in the economy?*
***Question 22**: What new job opportunities will AI create? What measures should be taken to strengthen the AI workforce, to ensure that Americans from all backgrounds and regions have opportunities to pursue careers in AI, and otherwise to prepare American workers for jobs augmented or affected by AI?*

Response to Questions 20, 22: From retail to manufacturing, AI is powering digital transformation across industries, providing tremendous benefits across sectors.

AI will create shifts in the labor market, but these changes will include creating new jobs and enabling workers to do existing jobs more efficiently. To ensure that American workers are prepared for the jobs of tomorrow, the United States should expand workforce retraining programs and alternative pathways and improve access to and support for STEM education. Industry and government should invest in programs to support the creation of alternative paths to AI careers that enable workers to develop high-demand technology skills without the need for a bachelor's or graduate degree. Programs like apprenticeships, partnerships with community colleges, "boot camps," and public service opportunities are all important gateways for helping new and mid-career workers develop in-demand digital skills. In addition, broadening opportunities, improving training programs, and expediting the development of a diverse workforce is needed to help people take advantage of the numerous opportunities available to and demand for skilled STEM workers.

***Question 24**: How can the Federal Government effectively and responsibly leverage AI to improve Federal services and missions? What are the highest priority and most cost-effective ways to do so?*
***Question 26:** How can the Federal Government work with the private sector to ensure that procured AI systems include protections to safeguard people's rights and safety?*

Response to Questions 24, 26: The federal government can use AI in myriad ways, including to assist federal employees who have limited resources to perform their job functions. For example, civilian federal acquisition professionals are often trying to do more with less. The number of professionals over the age of 60 is now four times higher than the number under the age of 30, and as retirements move these experts out of government service, the same workload is being borne by fewer people. AI can ease the load of repetitive work on acquisition staff so that they can focus on more complex items. For example, at GSA's Public Building Service, AI is being used for contract closeout activities, such as automatically sending reminders for self-assessments to federal contractors instead of requiring federal personnel to do so. In addition, AI is powering end-of-year spending estimates, which has enabled agencies to spend additional funding by the end of the year. These AI-assisted improvements provide greater productivity and allow employees to focus on mission-oriented tasks.

The federal government can work with the private sector to ensure that procured AI systems include protections to safeguard people's rights and safety. For example, the NIST AI RMF has the potential to help a broad range of companies implement processes for managing risks of AI systems. Agencies may encourage companies from which they procure AI systems to implement such risk management frameworks, including the NIST AI RMF or BSA's Framework. These frameworks provide a common methodology for both agencies and companies to identify and mitigate AI risks, define roles and responsibilities for relevant personnel, document testing, and consult with a diverse group of stakeholders. Implementation of these frameworks can help ensure responsible AI innovation.