BSA | The Software Alliance Comments on the National Information and Telecommunications Administration AI Accountability Policy Request for Comment

Regulations.gov Docket Number NTIA-2023-0005

BSA | The Software Alliance appreciates the opportunity to submit this feedback in response to the request for comments (RFC) on AI Accountability Policy by the National Telecommunications and Information Administration (NTIA).

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 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.

---

[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.

[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),

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

Our comments highlight four key themes:

- **AI Policies on Accountability Should Distinguish Between Developers and Deployers**. The AI supply chain includes developers of AI systems and deployers that use those AI systems. Both types of entities should have obligations to ensure responsible AI innovation, though 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 its 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.

- **The AI Audit Ecosystem Is Not Mature Enough to Support Mandatory Third-Party Audits**. The RFC makes several references to the potential for third-party auditors to conduct audits of AI systems. However, the process of developing auditable standards for AI is nascent. There are few existing procedures or best practices for companies to either: (1) choose a reputable company capable of auditing an AI system, or (2) determine what standards any such auditing company should apply. Indeed, although the International Organization for Standardization has issued some AI-related standards, including guidance on risk management practices, several other standards are still under development, and more broadly there is a lack of sufficient voluntary consensus-based standards addressing AI systems. Without common standards, the quality of any audits will vary significantly because different audits may measure against different benchmarks, undermining the goal of obtaining an evaluation based on an objective benchmark. Moreover, there is no clear way to ensure audits are conducted by reputable companies that have met specific criteria demonstrating they are qualified to conduct audits of AI systems. The variation among existing auditing companies allows organizations to select auditors based on their own preferred criteria, methods, and scope, which also makes the resulting audits less reliable. This concern is exacerbated by the lack of professional bodies governing AI auditors, which is important for ensuring auditors adhere to ethical standards. As discussed in more detail below, these and other concerns can create significant limitations in the use of third-party audits as an AI accountability tool.

- **Federal Legislation Should Require the Use of AI Impact Assessments for High-Risk Uses of AI Systems**. Federal law should establish clear rules of the road for AI developers and deployers. Specifically, federal legislation should require both developers and deployers to conduct impact assessments for high-risk uses of AI systems to ensure they are examining risks posed by an AI system throughout its lifecycle. Impact assessments are critical tools that companies can readily leverage to help ensure AI accountability by enabling organizations to identify and mitigate risks.

---

available at https://www.congress.gov/117/meeting/house/114125/witnesses/HHRG-117-BA00-Wstate-CooperA-20211013.pdf.

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

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

- **AI Policies Should Be Designed to Permit the Continued Growth and Use of Open Source AI Systems.** An increasingly important part of the AI supply chain consists of open source AI, including the use of 'fine-tuned" (retrained) open source large language models (LLMs) in a wide variety of business applications. It is important that any AI framework facilitates the continued growth of the open source ecosystem, while also addressing risks that may emerge. Generally, it will likely not make sense to place obligations on developers or contributors to open source projects to have knowledge of or connection with downstream participants. This is an important aspect of the ongoing discussions around the development and use of AI systems.

***Question 1**. What is the purpose of AI accountability mechanisms such as certifications, audits, and assessments?*

Accountability mechanisms enable organizations to demonstrate that they have developed or deployed AI responsibly, which is critical for cultivating public trust in AI systems. However, there are significant differences among different types of AI accountability mechanisms.

Impact assessments enable organizations to identify and mitigate risks throughout an AI system's lifecycle and should be conducted by developers and deployers with regard to high-risk uses of AI systems. By allowing personnel across the organization to examine the objectives, data preparation, design choices, and testing results, these assessments help drive internal changes to an organization's risk management program. Implementing these changes enables organizations to better address existing concerns and adapt to new risks as they emerge.

Third-party audits also aim to enhance accountability and are commonplace in some fields — such as cybersecurity — where there are well-established standards and professional norms. However, the same is not true of external audits in the AI context. AI standards development is in its early stages and, therefore, there is a lack of sufficient measurable standards to audit against. There is also a lack of professional bodies to oversee third-party auditors, which play a critical role in ensuring auditors adhere to ethical standards. In addition, third-party audits often require access to confidential business proprietary information, creating concerns around the treatment of trade secrets and other sensitive information. Moreover, 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 presented in the context of internally-focused impact assessments.

For these reasons, impact assessments are currently the most useful accountability mechanism. Indeed, a recent report on AI accountability concluded that impact assessments had several advantages over other accountability tools, noting that "(1) they are familiar to organizations already conducting impact assessments for privacy and data protection; (2) they are practical because they do not rely on technical standards, which are currently nascent; [and] (3) they are future-proof because they can adapt as AI systems and AI governance evolve."[6]

[6] Access Partnership, Workday, Impact Assessments: Supporting AI Accountability & Trust 11 (Jan. 28, 2023), *available at* https://accesspartnership.com/impact-assessments-supporting-ai-accountability/.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

***Question 2**. Is the value of certifications, audits, and assessments mostly to promote trust for external stakeholders or is it to change internal processes? How might the answer influence policy design?*

The principal value of an impact assessment is that it allows an organization to rigorously examine its practices, which drives change in internal processes. As discussed above, these changes help organizations adapt to new and emergent risks and implement changes across their products and services. The fact that assessments are being performed for high-risk uses of AI systems also promotes trust for external stakeholders because they will know that an organization is conducting a thorough examination of AI systems, and that the assessments are available to regulators upon request in the event of an investigation.

Because impact assessments are designed to drive change in internal processes, we strongly recommend the development of AI accountability policies that acknowledge the utility of internal assessments. There are several ways to enhance accountability within an organization. For example, a company can structure its impact assessment process to include multiple layers of independent review across an organization. Impact assessments can also be designed to ensure that relevant stakeholders are consulted at appropriate stages.

Importantly, these impact assessments should be treated as confidential to preserve the incentives for companies to implement them through rigorous processes that identify and mitigate a wide range of potential risks. Public disclosure would change an organization's incentives and result in less thorough examinations that do not surface as many issues. To enhance public trust, organizations may want to self-certify that they have completed an assessment, which would give stakeholders the assurance that risks have been identified and addressed.

***Question 3**. AI accountability measures have been proposed in connection with many different goals, including those listed below. To what extent are there tradeoffs among these goals? To what extent can these inquiries be conducted by a single team or instrument?*
a. The AI system does not substantially contribute to harmful discrimination against people.
b. The AI system does not substantially contribute to harmful misinformation, disinformation, and other forms of distortion and content-related harms.
c. The AI system protects privacy.
d. The AI system is legal, safe, and effective.
e. There has been adequate transparency and explanation to affected people about the uses, capabilities, and limitations of the AI system.
f. There are adequate human alternatives, consideration, and fallbacks in place throughout the AI system lifecycle.
g. There has been adequate consultation with, and there are adequate means of contestation and redress for, individuals affected by AI system outputs.
h. There is adequate management within the entity deploying the AI system such that there are clear lines of responsibility and appropriate skillsets.

The overarching goal of performing impact assessments is to identify and mitigate harms created by high-risk uses of AI systems. These potential harms can vary and may include discrimination or risks to privacy and safety. Because different systems will pose different risks, it is important to avoid a one-size-fits-all approach.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

As the National Institute of Standards and Technology's (NIST) AI Risk Management Framework (RMF) recognized, there are tradeoffs among the goals of preventing these individual harms. NIST acknowledged that "[a]ddressing AI trustworthiness characteristics individually will not ensure AI system trustworthiness; tradeoffs are usually involved, rarely do all characteristics apply in every setting, and some will be more or less important in any given situation."[7] For example, privacy-enhancing techniques could result in a loss of accuracy, which could affect decisions about fairness.[8] For these reasons, a context-based approach will be crucial to navigating these tradeoffs.

Question 5. *Given the likely integration of generative AI tools such as large language models (e.g., ChatGPT) or other general-purpose AI or foundational models into downstream products, how can AI accountability mechanisms inform people about how such tools are operating and/or whether the tools comply with standards for trustworthy AI?*

Transparency is an important goal in promoting AI accountability. As policy approaches to enhance transparency are developed, it is important to distinguish between the different roles that companies play in the AI ecosystem. Developers are the companies that design, code, or produce AI systems. In contrast, deployers are the companies that actually implement AI systems in specific use cases. While developers are best positioned to understand the intended uses and limitations of an AI system, deployers are the entities that typically interact with consumers and decide when and how to deploy the AI system for a particular use. Both types of companies have distinct roles to play in ensuring appropriate education and transparency for consumers.

It is important to note that AI accountability is heavily context-dependent. Developing an accountability framework based on a single use case — i.e., consumer-facing uses — can be detrimental to many applications used in other contexts, such as business-to-business settings that may present vastly different expectations and risk considerations. As an example, many companies are building business tools that enhance products and services using LLMs. Sometimes companies use open source versions of LLMs as a base, fine tuning (retraining) the LLM using technical information specific to the customer's internal use case, such as identifying particular corporate forms or creating transcriptions of internal meetings. This application of the original LLM presents significantly lower potential risks than some consumer-facing applications of an LLM, which may create specific risks to consumers that warrant a more rigorous accountability framework applicable to the deployer.

Question 7*. Are there ways in which accountability mechanisms are unlikely to further, and might even frustrate, the development of trustworthy AI? Are there accountability mechanisms that unduly impact AI innovation and the competitiveness of U.S. developers?*

Impact assessments can help achieve the important goal of identifying and mitigating AI risks for high-risk uses of AI systems. As discussed above, these assessments can help drive change in internal processes. In contrast, third-party audits do not necessarily offer the same benefits and have significant drawbacks. For example, third-party audits often require the disclosure of proprietary information and, in some cases, consumers' personal information, harming competitiveness and undermining efforts to protect privacy and security. In practice, companies that engage a third-party auditor may need to spend significant resources designing a set of privacy protections for the personal data provided

[7] National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework 12, *available at* https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf.
[8] *Id.* at 12-13.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001
P 202-872-5500
W bsa.org

to that third-party auditor to ensure the data is not misused and that the company has permission to provide the data to the auditor in the first place. Similar safeguards may be needed to protect confidential business information and to address cybersecurity concerns to ensure any audit is conducted in a way that does not frustrate the company's security and privacy objectives.

The lack of measurable AI auditing standards may also undermine trustworthiness, as the quality of audits will vary significantly without common standards. This variation frustrates efforts to provide assurances that AI systems have been evaluated based on a common objective benchmark. Further, the lack of common standards allows organizations to select auditors based on the most favorable criteria, methods, and scope, making the audits less reliable.

In addition, proactive disclosure to regulators or public disclosure of third-party audits or impact assessments could impede efforts to conduct a rigorous internal inquiry, as the preparation may vary once the purpose of the inquiry is to share with a regulator or the public. In lieu of proactive disclosure, organizations should be able to self-certify that they have conducted appropriate impact assessments.

In contrast, because impact assessments are designed to be conducted internally and to drive change across an organization, they do not present the same privacy, security, or business confidentiality concerns that are presented when sharing AI-related data with a third-party auditor. Moreover, many companies that operate across state lines already have existing processes in place to conduct impact assessments, particularly those companies subject to state and EU privacy laws that require privacy impact assessments. Companies can more readily adapt those processes to conduct AI-related impact assessments, given their familiarity with those requirements in the privacy context.

Accountability mechanisms should also account for different roles of developers and deployers of AI systems. These different companies will have access to different types of information and be positioned to assess different types of risks, given their different roles in either developing or deploying the AI system. These differences also occur with respect to open source AI systems. For example, an accountability mechanism that requires an originator of, or contributor to, an open source AI system to monitor or interact with downstream users would generally not be workable because of the fundamental nature of open source software. These considerations are important in designing accountability systems, to ensure a company's role in promoting accountability is in line with its role in developing and deploying an AI system.

***Question 9**. What AI accountability mechanisms are currently being used? Are the accountability frameworks of certain sectors, industries, or market participants especially mature as compared to others? Which industry, civil society, or governmental accountability instruments, guidelines, or policies are most appropriate for implementation and operationalization at scale in the United States? Who are the people currently doing AI accountability work?*

Accountability mechanisms are more mature in other fields, such as cybersecurity and financial regulation. For example, in the United States, businesses conduct audits or assessments of their cybersecurity practices to comply with a range of laws including:

- Sarbanes-Oxley Act (SOX), which requires publicly traded companies to maintain adequate controls, including cybersecurity controls, over their financial reporting;
- Health Insurance Portability and Accountability Act of 1996 (HIPAA), which requires organizations that possess patient health information to protect that information;

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

- Gramm-Leach-Bliley Act (GLBA), which requires financial institutions to secure customer information;
- Federal Acquisition Regulation (FAR), which requires organizations that sell solutions to the US Government to meet baseline cybersecurity practices; and
- Defense Federal Acquisition Regulations Supplement (DFARS), which requires organizations in the defense industrial base to meet baseline cybersecurity practices.

In addition to any legal requirements to conduct cybersecurity audits, customers often require their vendors to demonstrate strong cybersecurity practices — creating another layer of certifications and audit requirements. For example, customers frequently require vendors to certify they are compliant with the ISO 27000 series of standards (which govern information security management)[9] and Service Organization Control (SOC) 2 Type 2 requirements (which assess controls related to security, availability, processing integrity, confidentiality, or privacy of information).[10] Companies that offer multiple products may be required to obtain a certification for each product, compounding these requirements.

Privacy impact assessments are another important reference point and are required by state and EU privacy laws. These assessments are typically required where the data processing presents a heightened risk of significant harm. Privacy impact assessments require organizations to weigh the benefits of the data processing against the potential risk to consumers. The assessments help organizations understand how they are collecting information, how it is shared, and how to manage data risks. Privacy impact assessments are useful examples of accountability mechanisms.

The BSA Framework recommends the use of impact assessments for high-risk uses of AI systems and is appropriate for organizations to implement to enhance AI accountability. The NIST AI RMF is also appropriate for this purpose. Both 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.

We recently prepared a crosswalk between the two frameworks, which illustrates the significant alignment between the two approaches.

***Question 14**. Which non-U.S. or U.S. (federal, state, or local) laws and regulations already requiring an AI audit, assessment, or other accountability mechanism are most useful and why? Which are least useful and why?*

A range of state privacy laws require privacy impact assessments for data processing that presents a heightened risk of harm, including for profiling. For example, the Colorado

[9] *See* ISO/IEC 27001 and related standards, *available at* https://www.iso.org/isoiec-27001-informationsecurity.html.

[10] *See* Association of International Certified Professional Accountants, SOC for Service Organizations, *available at* https://us.aicpa.org/interestareas/frc/assuranceadvisoryservices/serviceorganizationsmanagement.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

Privacy Act requires a data protection impact assessment for profiling where the profiling presents a reasonably foreseeable risk of unfair or deceptive treatment of, or unlawful disparate impact on, consumers; financial or physical injury to consumers; an intrusion upon the seclusion or private affairs of a consumer; or other substantial injury to consumers.[11] The law requires companies to weigh the benefits of the data processing against the potential risks to consumers. It also requires companies to factor the context of the processing into this assessment.

Privacy laws in Connecticut, Indiana, Montana, Tennessee, and Virginia all similarly require companies to conduct privacy impact assessments for certain activities. California will also require such assessments after completion of its upcoming rulemaking. The EU's General Data Protection Regulation also requires companies to conduct such assessments, which must contain: (1) a systematic description of the envisioned processing and its purposes, (2) an assessment of the necessity and proportionality of the processing in relation to those purposes, (3) an assessment of the risks to the rights and freedoms of data subjects, and (4) measures envisioned to address those risks, including safeguards, security measures, and mechanisms to ensure the protection of personal data.

One example of an AI audit requirement that has proven challenging to implement is New York City's ordinance, Local Law 144, on automated employment decision tools.[12] The ordinance requires a bias audit to be conducted not more than one year prior to the use of the tool. Regulators have issued multiple rounds of regulations to attempt to address concerns raised by the ordinance's broad and vague language. Even with these implementing regulations, however, the ordinance poses challenges for several reasons.

First, the implementing regulations require independent, third-party audits. For the reasons expressed above, third-party audits create a range of policy concerns, as they not only implicate the sharing of confidential business information, but they also may require access to sensitive personal data. Second, the regulations also require the calculation of selection rates and impact ratios based on categories that do not align with the EEOC's categories for disparate impact testing. Third, the results of the bias audit, including the selection rates and impact ratios, must be posted on the company's website, which could actually undercut efforts to create a more diverse workforce because applicants may be discouraged to apply to companies with lower selection rates. For these reasons, the New York City ordinance and implementing regulations should not be a model that other jurisdictions considering similar issues should adopt.

Question 15*. The AI value or supply chain is complex, often involving open source and proprietary products and downstream applications that are quite different from what AI system developers may initially have contemplated. Moreover, training data for AI systems may be acquired from multiple sources, including from the customer using the technology. Problems in AI systems may arise downstream at the deployment or customization stage or upstream during model development and data training.*
a. Where in the value chain should accountability efforts focus?
b. How can accountability efforts at different points in the value chain best be coordinated and communicated?
c. How should vendors work with customers to perform AI audits and/or assessments? What is the role of audits or assessments in the commercial and/or public procurement

[11] Colo. Rev. Stat. § 6-1-1309(2)(a). A heightened risk of harm includes targeted advertising, the sale of personal data, and processing sensitive data.
[12] Title 20 of the Administrative Code of the City of New York, Chapter 5, Subchapter 25, *available at* https://codelibrary.amlegal.com/codes/newyorkcity/latest/NYCadmin/0-0-0-135598.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001
P 202-872-5500
W bsa.org

process? Are there specific practices that would facilitate credible audits (e.g., liability waivers)?
d. Since the effects and performance of an AI system will depend on the context in which it is deployed, how can accountability measures accommodate unknowns about ultimate downstream implementation?

As discussed above, the AI supply chain includes developers and deployers. Any policies that create obligations for companies that design and use AI systems should reflect these different roles and assign obligations accordingly.

Communication among the different AI actors is important to ensuring the successful operation of accountability frameworks across the lifecycle of an AI system. Developers that design an AI system should provide deployers using that AI system with the information reasonably necessary for the deployer to conduct an impact assessment. This may include the AI system's capabilities, known limitations, and guidelines for intended use. By providing this information, a deployer can then assess the use of an AI system in light of the developer's intended use for the system and its known limitations. At the same time, because developers will not have insight into the actual use of the AI system and do not have a relationship with the consumer, a deployer should be responsible for monitoring issues that arise in downstream implementation, including having a feedback mechanism to surface issues that arise.

An important and growing aspect of the AI supply chain consists of the use of open source AI models and components. This use case is another illustration of the importance of distinguishing between developers of AI systems and deployers of those systems in creating accountability mechanisms. As described above, differing regulatory obligations should apply to developers and deployers, with the developer obligations focusing on steps taken during the design and development stages of the AI system and the deployer obligations relating to how the AI system is used at launch and thereafter. In the context of open source AI, the developer has no way of monitoring the entities that eventually use the open source AI, and, like other contexts, no way to monitor how those entities are using the AI system. Developers of those systems should be responsible for activities prior to, and up to the point of, releasing the system as an open source project, such as pre-release testing, risk assessment and calibration, and documentation. Generally, regulatory obligations focused on the use of an AI system, including those based on open source, are best applied to the deployer.[13]

***Question 16**. The lifecycle of any given AI system or component also presents distinct junctures for assessment, audit, and other measures. For example, in the case of bias, it has been shown that "[b]ias is prevalent in the assumptions about which data should be used, what AI models should be developed, where the AI system should be placed — or if AI is required at all." How should AI accountability mechanisms consider the AI lifecycle?*

Impact assessments of high-risk AI systems should address the lifecycle of those systems and examine issues in all stages: project conception; data acquisition; data preparation and model definition; validation, testing, and revising the model; and preparing for deployment and use.

[13] Importantly the act of releasing a project as open source (or making a contribution to an existing open source project) does not, by itself, constitute a deployment; the deployer of an open source software AI system should be considered to be the party that actually puts such a system into use.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

Although one of the many benefits of AI systems is that they can be used to help detect bias, bias, as the question highlights, could occur at multiple stages in the development and deployment of an AI system. For example, problem formulation bias — where the basic assumptions underlying a proposed AI system may be inherently biased — could occur at the conception phase. 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.

At each stage in the lifecycle of a high-risk AI system, impact assessments enable companies to identify, examine, and mitigate risks, including bias. For example, the BSA Framework recommends that in the data acquisition phase, companies should evaluate the representativeness of the data as part of an impact assessment. To do so, they can compare demographic distribution of training data to the population where 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 the documentation of the 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. In sum, a lifecycle-based approach is necessary to identify and address risks throughout the various stages of development and deployment.

In regard to frequency, impact assessments for high-risk uses of AI systems should be performed prior to launch and annually post-launch, with additional updates when there is a material change to the intended purpose or type of input data used. Such an approach allows flexibility for a system to be updated routinely without additional requirements but also highlights the need to conduct an impact assessment when potentially significant changes warrant further examination.

***Question 20**. What sorts of records (e.g., logs, versions, model selection, data selection) and other documentation should developers and deployers of AI systems keep in order to support AI accountability? How long should this documentation be retained? Are there design principles (including technical design) for AI systems that would foster accountability-by-design?*

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.

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

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

- 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.

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.

***Question 21**. What are the obstacles to the flow of information necessary for AI accountability either within an organization or to outside examiners? What policies might ease researcher and other third-party access to inputs necessary to conduct AI audits or assessments?*

Organizations should have risk management processes in place that foster communication of AI risks with internal stakeholders. The BSA Framework calls for senior leadership to be adequately briefed on potential high-risk AI systems. It also recommends that validation and testing documentation should be reviewed by personnel who were not involved in the system's development. In addition, it encourages organizations to define and document who is responsible for the system's outputs, including details about how a system's decisions can be reviewed if necessary. Effective communication at each of these junctures is critical for accountability systems to work in practice.

However, providing information to outside examiners presents challenges because it often requires the disclosure of confidential business proprietary information and, in some cases, consumers' personal information, undermining competitiveness and efforts to protect privacy and security. Indeed, this is a key challenge with third-party audits. As discussed above, impact assessments are useful mechanisms for enhancing accountability within an organization and do not implicate the issues that arise when third parties are involved.

One aspect of open source AI relevant to accountability and transparency is that the model code and, typically, the training data are readily available for inspection by anyone, including outside examiners and any other interested party.

***Question 22**. How should the accountability process address data quality and data voids of different kinds? For example, in the context of automated employment decision tools, there may be no historical data available for assessing the performance of a newly deployed, custom-built tool. For a tool deployed by other firms, there may be data a vendor has access to, but the audited firm itself lacks. In some cases, the vendor itself may have intentionally limited its own data collection and access for privacy and security purposes. How should AI accountability requirements or practices deal with these data issues? What should be the roles of government, civil society, and academia in providing useful data sets (synthetic or otherwise) to fill gaps and create equitable access to data?*

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

Access to relevant data is an important issue in ensuring an effective accountability framework for AI systems. As the question highlights, the use of historical data may not always be possible. In lieu of historical data, organizations should be able to use test data for assessing performance of AI systems.

Organizations may also appropriately take privacy and security considerations into account by limiting access to data when selecting and designing an accountability mechanism. However, it is important to recognize that these considerations may require establishing processes that can impede third parties' ability to evaluate an AI system. One solution to this challenge is to encourage the use of impact assessments for high-risk AI systems. These assessments allow internal stakeholders to thoroughly examine an AI system without creating the privacy and security risks that may arise in providing data in an AI system to a third party.

In addition, assembling good datasets can be time-consuming and expensive. The government should build on the success of the OPEN Government Data Act and expand its efforts to make non-sensitive data sets publicly available, which would provide more equitable access to data.

***Question 24**. What are the most significant barriers to effective AI accountability in the private sector, including barriers to independent AI audits, whether cooperative or adversarial? What are the best strategies and interventions to overcome these barriers?*

The lack of sufficient voluntary consensus-based standards impedes the ability to conduct effective third-party audits. Without common standards, the quality of third-party audits will vary, undermining the utility of attempting to evaluate systems based on an objective benchmark. Further, this variation allows organizations to select auditors based on the most favorable criteria, methods, and scope, making them less reliable. In addition, the lack of professional bodies governing AI auditors also limits the perceived assurances provided by third-party audits.

To overcome these barriers, NIST and industry should continue to work in international standards development organizations to develop appropriate standards. We note that the International Organization for Standardization has issued some AI-related standards, but there are several more under development. Impact assessments are increasingly important accountability tools that do not have the same challenges as third-party audits and are therefore more readily implemented for high-risk AI systems.

***Question 25**. Is the lack of a general federal data protection or privacy law a barrier to effective AI accountability?*

Establishing a strong, comprehensive federal privacy law is a top priority for BSA. A federal privacy law would provide important protections for consumers' personal information and limit how that data is collected, used, and shared. Implementing measures to safeguard personal information is crucial to the development of AI systems that may use this data to train models, and to the deployment of AI systems that may collect personal information during their use. The lack of a federal privacy law means that companies can adopt varied approaches to how they collect, use, and protect data, leading to different levels of protection for how personal information is treated. A federal privacy law would create clear and uniform guardrails around these practices, which would support the further development of AI accountability mechanisms.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

***Question 26**. Is the lack of a federal law focused on AI systems a barrier to effective AI accountability?*

Laws and regulations already govern AI today, because an action that is already unlawful is not somehow made lawful through the use of AI. For example, the Consumer Financial Protection Bureau has confirmed that adverse action notification requirements still apply when creditors use AI to make credit decisions.[14] The Federal Trade Commission has also highlighted its authority over issues implicating AI, including addressing false and unsubstantiated claims and unfair practices.[15] Agencies should review their existing legal authority and assess whether there are any gaps that warrant further regulation on AI. Some agencies have already announced their intention to use their existing enforcement authority to address AI-related issues.[16]

A federal law focused on high-risk AI systems would create an important mechanism to further promote effective AI accountability. BSA has for several years called for a federal law that establishes an accountability framework for companies developing and deploying AI systems for high-risk uses. Such a law should require developers and deployers to conduct impact assessments for high-risk uses, tailored to their respective roles in the AI ecosystem. It should also require developers and deployers to implement risk management programs that establish the policies, processes, and personnel that will be used to identify, mitigate, and document AI risks. Without a federal law, responsible companies will continue to examine and address AI risks, but this will not be the case for all AI actors. In addition, a patchwork of state and local laws could emerge, creating differing obligations for the same AI systems.

***Question 27**. What is the role of intellectual property rights, terms of service, contractual obligations, or other legal entitlements in fostering or impeding a robust AI accountability ecosystem? For example, do nondisclosure agreements or trade secret protections impede the assessment or audit of AI systems and processes? If so, what legal or policy developments are needed to ensure an effective accountability framework?*

AI accountability policies should account for intellectual property rights and contractual obligations. Companies invest valuable time and resources into developing innovative AI systems and appropriately seek to protect the confidential business proprietary information related to these efforts. Third-party audits often require disclosure of this confidential information, which may impede the ability of companies to hire outside auditors, even in an ecosystem where there is an organizing body to accredit third-party auditors and existing standards for auditors to apply. In addition, companies may have contractual obligations that require certain privacy and security protections or otherwise limit their ability to share

---

[14] *See* Consumer Financial Protection Bureau, Consumer Financial Protection Circular 2022-03, *available at* https://www.consumerfinance.gov/compliance/circulars/circular-2022-03-adverse-action-notification-requirements-in-connection-with-credit-decisions-based-on-complex-algorithms/.

[15] *See* Federal Trade Comm'n, Keep Your AI Claims in Check, Feb. 27, 2023, *available at* https://www.ftc.gov/business-guidance/blog/2023/02/keep-your-ai-claims-check; Federal Trade Comm'n, Chatbots, Deepfakes, and Voice Clones: AI Deception for Sale, Mar. 20, 2023, *available at* https://www.ftc.gov/business-guidance/blog/2023/03/chatbots-deepfakes-voice-clones-ai-deception-sale; Federal Trade Comm'n, The Luring Test: AI and the Engineering of Consumer Trust, May 1, 2023, *available at* https://www.ftc.gov/business-guidance/blog/2023/05/luring-test-ai-engineering-consumer-trust.

[16] *See* Consumer Financial Protection Bureau, Dept. of Justice, Equal Employment Opportunity Comm'n, Federal Trade Comm'n, Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems, *available at* https://www.eeoc.gov/joint-statement-enforcement-efforts-against-discrimination-and-bias-automated-systems.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

relevant data — including limits that are imposed on them by privacy, cybersecurity, and other laws. Impact assessments of high-risk AI systems do not create the same interference with intellectual property rights or the same privacy and security concerns, which enhances the ability of companies to more readily implement impact assessments to enhance accountability.

***Question 29**. How does the dearth of measurable standards or benchmarks impact the uptake of audits and assessments?*

As discussed in response to Question 24, the lack of voluntary consensus-based standards impedes the ability to conduct effective third-party audits. The AI auditing field is nascent, and there are ongoing efforts to develop international standards. Currently, this means that third-party auditors would apply varied standards, undermining efforts to assess AI systems against an objective and shared benchmark. Further, this variation can allow organizations to select auditors based on the most favorable criteria, methods, and scope, making third-party audits less reliable.

***Question 30**. What role should government policy have, if any, in the AI accountability ecosystem? For example:*
b. Should AI accountability regulation, if any, focus on inputs to audits or assessments (e.g., documentation, data management, testing and validation), on increasing access to AI systems for auditors and researchers, on mandating accountability measures, and/or on some other aspect of the accountability ecosystem?
c. If a federal law focused on AI systems is desirable, what provisions would be particularly important to include? Which agency or agencies should be responsible for enforcing such a law, and what resources would they need to be successful?

As discussed above, agencies should review their existing legal authority and assess whether there are any gaps that warrant further regulation on AI.

For several years, BSA has called for legislation to require companies that develop and deploy high-risk AI systems to conduct impact assessments. We believe that approach will help companies identify and address potential issues across the lifecycle of an AI system. At the same time, we believe any legislation on AI accountability should accomplish several objectives.

First, it should focus on high-risk uses of AI. Because AI is integrated into an incredibly broad array of products and services, many of which provide great value at low risk and low cost, policymakers should focus legislative requirements on AI systems likely to create high risks. These are the systems where accountability mechanisms are most needed — and should be used to identify and address risks across those systems.

Second, it should distinguish between AI developers and deployers and tailor any obligations according to their respective roles.

Third, legislation should require developers and deployers of high-risk systems to conduct an impact assessment, tailored to their respective roles in the AI ecosystem. The assessments should require documentation of key aspects examined, which are outlined in response to question 20.

Fourth, legislation should also require these accountability mechanisms to be grounded in company-wide policies and procedures for addressing AI risks. BSA supports requiring companies that develop or deploy high-risk AI systems to implement a risk management

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org

program that establishes the policies, processes, and personnel that will be used to identify, mitigate, and document AI risks.

***Question 34**. Is it important that there be uniformity of AI accountability requirements and/or practices across the United States? Across global jurisdictions? If so, is it important only within a sector or across sectors? What is the best way to achieve it? Alternatively, is harmonization or interoperability sufficient and what is the best way to achieve that?*

Uniform legal requirements for high-risk AI systems will be important to foster responsible development in these technologies. In the United States, any national law should include federal preemption that applies consistent standards throughout the country. As more countries seek to regulate AI, global interoperability across legal requirements applied to AI systems will become increasingly important so that companies that develop and deploy AI systems can create strong compliance programs that meet legal obligations across the countries in which they operate. As a result, there is a need for consistent rules on AI accountability to ensure that companies identify and mitigate potential risks involved in that AI system across its lifecycle.

* * *

Thank you for the opportunity to provide comments on AI accountability. We look forward to serving as a resource as you continue to engage in policy discussions on this issue.

200 Massachusetts Avenue, NW
Suite 310
Washington, DC 20001

P 202-872-5500
W bsa.org