



Ernst & Young LLP
1775 Tysons Blvd
Tysons, VA 22102

Tel: +1 703 747 1000
Fax: +1 703 747 0100
ey.com

Mr. Alan Davidson
Assistant Secretary of Commerce for Communications and Information
National Telecommunications and Information Administration
1401 Constitution Avenue, NW
Washington, DC 20230

12 June 2023

National Telecommunications and Information Administration AI Accountability Policy Request for Comment [Docket No. 230407-0093]

Dear Mr. Davidson

Ernst & Young LLP is pleased to submit comments in response to the *AI Accountability Policy Request for Comment* (RFC) issued by the National Telecommunications and Information Administration (NTIA) ", which seeks input on what policies can support the artificial intelligence (AI) accountability ecosystem.

The opportunities presented by the widespread deployment of AI technologies could potentially lead to scientific breakthroughs, increase the efficiency of business operations, conserve resources for complex and creative thinking and improve livelihoods. However, the deployment of such powerful technologies must be done responsibly to promote fairness, reduce risks and instill confidence in the processes and outcomes fueled by AI.

Ernst & Young LLP is the US member firm of the global network of EY member firms that provides advisory, assurance, tax and transaction services to entities worldwide. We comprise approximately 54,000 of the more than 365,000 EY professionals worldwide and serve many companies across a wide range of industries in the US.

Ernst & Young LLP and artificial intelligence

We have a unique perspective in the field of AI and machine learning given our network's global, cross-sector footprint. Our team of dedicated data scientists advise entities across every sector of the economy on adopting a holistic view of automation, process and service improvement. We help organizations craft, deploy and evaluate AI tools to responsibly accomplish their digital transformation goals.

As an early leader in identifying the need to build trust in AI, we created a Trusted AI Framework¹ to assist organizations in understanding and mitigating the risks that arise throughout the AI lifecycle. This framework may be a useful resource for NTIA to gain a business perspective on AI risk and measures to manage it, focusing on:

- Design risk – Is the AI application designed to meet the business objective?
- Data risk – Is the right data available to make this work? Is the solution sufficiently unbiased and resilient?
- Algorithmic risk – Is the AI application explainable, and does it leave an appropriate audit trail?
- Performance risk – Does the AI application meet performance standards, and is its use transparent to the business owners? Is the output of the AI application appropriately monitored?

¹ https://assets.ey.com/content/dam/ey-sites/ey-com/en_gl/topics/digital/ey-how-do-you-teach-ai-the-value-of-trust.pdf

Through integrated teams involving members from our data and analytics practice and our risk practice, we have helped organizations evaluate whether the AI systems they deploy meet their specific goals, adhere to their predetermined ethical principles, comport with their risk tolerance level, meet their stakeholders' expectations regarding using AI responsibly and comply with applicable laws, standards and regulations. In addition, we have developed our own AI principles to guide the appropriate implementation and use of AI in our firm, based on the EY Trusted AI Framework.

We are also a leader in performing high-quality audits and other attest services that build trust in financial markets and business. Our audit professionals lead the profession in providing assurance services to the largest share of Fortune 10, Fortune 500, Fortune 1000 and Russell 3000 companies. In providing audit services, we are actively engaged with entities to understand the risks that AI systems may represent to the financial statements of organizations. In addition, our assurance services require us to evaluate the risk resulting from the use of emerging technologies and design procedures to evaluate and conclude on those risks. To determine that such services meet the expectations of stakeholders, our assurance professionals also actively lead and support the efforts of standard setters as they develop guidance for the public accounting profession that directs efforts to mitigate the risks introduced through the organization's use of new technologies and provide assurance over them.

This combination of experience obtained through developing and deploying AI systems and addressing the risks that arise from the organization's use of AI systems makes our perspective on AI accountability particularly applicable to the NTIA RFC.

Our responses to questions raised in the NTIA's RFC

Our views below are informed by our experience advising entities on trusted AI approaches and performing financial statement audits, program audits, governance and sustainability assurance reporting, and examinations of internal controls and other subject matters. They are also based on our involvement in the public accounting standard setters' efforts to develop assurance reporting guidance for new technologies.

Q1: What is the purpose of AI accountability mechanisms such as certifications, audits, and assessments (subparts b, c, d and e)?

When the risks associated with a matter, such as an AI system, may have a significant negative effect on the wellbeing of society or individuals or may prevent beneficial activity from occurring, there is a benefit to identifying and mitigating those risks and preparing information about the risks and their mitigation. Such information may include data related to the matter or a description of the process and controls over the matter. This risk identification and mitigation may be initiated and addressed by the party responsible for the matter, the parties affected by the matter or, when the matter is in the public interest, a government body.

When the identification and mitigation of risks are sufficiently important, accountability mechanisms over the matter and the associated information may be desirable or necessary. Such information may cover system functionality and compliance with legal requirements, industry standards or the commitments made by the persons responsible. In some instances, the information and its verification may be used by the party responsible for the matter to implement improvements. In other cases, the information and its verification may be provided voluntarily to the parties affected by the matter. In other cases, public reporting may be required.

AI systems hold great promise for improving individual and societal wellbeing but also present significant risks. Furthermore, because AI is a complex and fast-evolving technology, the benefits and the risks are, as

yet, not completely known. Consequently, there is likely to be a significant benefit to making sure AI systems are “legal, effective, ethical, safe, and otherwise trustworthy.”² While there is currently no generally accepted, standardized accountability mechanism or supporting information disclosure regime for AI systems, it may be appropriate to establish such accountability programs in the future.

In evaluating the need and appropriateness of AI system accountability mechanisms, consideration needs to be given to the fact that one single accountability mechanism will not likely meet the needs of all stakeholders in a cost-effective manner. Furthermore, due to the rapidly changing AI ecosystem, any mechanisms developed today will need to be modified or replaced over time. Regardless of whether an accountability mechanism is established voluntarily by an organization, an industry body or a government, key components of accountability mechanisms typically include but are not limited to:

- Identification of parties affected by the AI systems (stakeholders) and their needs
- Identification of persons or bodies able to represent both internal and external stakeholders in establishing and operating accountability mechanisms
- Identification of the effect of the AI system on the stakeholders
- Identification of the appropriate objectives and boundaries of the AI systems
- Determination of how to measure whether AI systems have exceeded the boundaries or failed to achieve their objectives
- Identification of the information needed from those responsible for the AI system and from other sources to permit evaluation of whether an AI system is operating within the boundaries and is achieving its objectives
- Provision of information to measure operation within the boundaries and the achievement of objectives by those responsible for the AI system and collection of relevant information from other sources
- Verification (such as certifications, audits and assessments) of the information provided (discussed further in the answer to question 8 below)
- Evaluation of the information
- Accountability of those responsible for operating the AI system within the boundaries and achieving its objectives

Each of these elements requires significant consultation to determine a workable mechanism for the assessment of AI systems. However, we agree with the NTIA that verification is crucial in:

“[h]elping to hold entities accountable for developing, using, and continuously improving the quality of AI products, thereby realizing the benefits of AI and reducing harms. These mechanisms can also incentivize organizations to invest in AI system governance and responsible AI products. Assurance that AI systems are trustworthy can assist with compliance efforts and help create marks of quality in the marketplace.”³

² [National Telecommunications and Information Administration AI Accountability Policy Request for Comment \[Docket No. 230407-0093\]](#)

³ [National Telecommunications and Information Administration AI Accountability Policy Request for Comment \[Docket No. 230407-0093\]](#)

Q2: 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?

Verification schemes are one component of accountability mechanisms. The value of verification schemes in the context of AI accountability can have both external and internal benefits for an organization. While they can contribute to promoting trust among external stakeholders such as customers, users and the public, they also play a role in identifying potential weaknesses in internal processes in organizations and strengthening those internal processes.

Internal verifications are performed to obtain input to help organizations change and improve their internal processes and are designed for internal use. Certifications and audits performed by objective third parties usually focus on providing assurance to external stakeholders.

External stakeholders, including investors, consumers, regulators and the general public, often rely on certifications and audits as indicators of an organization's commitment to responsible practices and governance. These types of verification can also help establish trust and confidence in AI systems, demonstrating that appropriate measures have been taken to address issues like fairness, transparency, privacy and security by providing assurance to external stakeholders that the AI systems have been evaluated against established standards or measures.

Internally, AI system verification schemes serve to evaluate and improve an organization's processes, policies and practices. They can identify gaps, vulnerabilities and areas of improvement in the organization's AI development, deployment and governance processes. Through verification against appropriate criteria, organizations can gain insights into potential risks, biases or ethical concerns associated with their AI systems. This allows them to make necessary adjustments, enhance their internal processes and make sure they are in compliance with relevant guidelines and regulations.

Q6: The application of accountability measures (whether voluntary or regulatory) is more straightforward for some trustworthy AI goals than for others. With respect to which trustworthy AI goals are there existing requirements or standards? Are there any trustworthy AI goals that are not amenable to requirements or standards? How should accountability policies, whether governmental or non-governmental, treat these differences?

The measurement and evaluation of the functionality of AI systems is a relatively new concept. Ongoing academic research regularly identifies new methods for evaluating AI systems and new challenges in their operating integrity. The National Institute of Standards and Technology (NIST) continues to research and drive the development of mechanisms for measuring the behavior of AI systems and the International Standards Organization (ISO) and International Electrotechnical Commission (IEC) are jointly working to address accountability mechanisms.

Currently, AI governance is the most amenable to measurement and verification. Consequently, the alignment of AI governance with the goals for AI systems is the scenario most likely to be ready for accountability mechanisms. Accountability mechanisms intended to address aspects of AI systems for which measurement is nascent will be challenging to implement.

Q8: What are the best definitions of and relationships between AI accountability, assurance, assessments, audits, and other relevant terms?

Due to the great diversity of stakeholders, the differing effects of AI systems on them, and the different risks and benefits that they will experience from an AI system, multiple accountability mechanisms and different levels of verification should be considered. Furthermore, the diversity of the current stakeholders leads to different definitions, potentially creating confusion. At the current stage, it may be more helpful to focus the concepts that underly specific terms rather than on the specific definitions.

As discussed above, verification schemes are one component of accountability mechanisms. Broadly speaking, AI accountability refers to the responsibility and answerability of individuals and organizations for the achievement of established objectives during the development, deployment and use of AI technologies. It may be most helpful to divide the concept of AI accountability into two categories:

- Internal accountability refers to the accountability within the organization for operating the AI system to achieve the entity's objectives related to operations, compliance, and effective reporting, and is overseen by those within the organization charged with governance. AI accountability starts with an organization establishing policies, procedures and controls over the entire AI lifecycle. These controls should include governance controls, monitoring controls, IT process controls and other relevant internal controls.
- External accountability refers to the accountability of those responsible for the AI system to stakeholders outside the organization. It will likely include designing AI systems to be transparent, fair, secure and reliable, and in a way that the impacts and outcomes are understood and monitored, and exceptions are addressed appropriately.

Verification refers to the evaluation of whether an AI system is operating within its established boundaries and whether the AI system and the processes supporting or supported by the AI system are achieving their objectives. Verification is typically intended for either internal accountability or external accountability and can be performed by internal or external parties with varying degrees of objectivity, formality and rigor.

Key factors to consider in establishing a verification scheme include:

- Whether it is intended for internal accountability (e.g., measuring operational efficiency), external accountability (e.g., measuring compliance with a regulation) or internal accountability that precedes external accountability (e.g., management's evaluation of the effectiveness of their system of internal control in meeting its objectives)
- What objective achievement is to be measured (it is usually not cost-effective to measure the achievement of all objectives, therefore, verification schemes usually focus on the objectives associated with higher-risk topics.)
- The criteria to be used to perform the verification (verification schemes that use relevant, neutral/objective, reliable/measurable, complete and understandable criteria are likely to be more valuable to users of the verification. Criteria may be established by those responsible, by an objective body using due process or by a governmental body.)
- The form of the verifier's conclusion, whether a schedule of results or a formal statement of conclusions
- Other information to be provided by the verifier
- The level of formality with which the verification is to be performed (e.g., is the verification to be performed using a process established by the organization being assessed, formal standards established by a body using due process (such as the AICPA *Attestation Standards*) or standards established by a governmental body?)
- The amount of evidence to be collected to support the verifier's conclusions

- The qualifications of the verifier and their personnel, including competence, capability, and objectivity and independence
- Any certification or authorization requirements issued by recognized bodies or licensure requirements established by government bodies

For example:

- An organization that has concerns about whether its AI system development processes are being followed may establish an ad hoc team of its best personnel to investigate the matter and report to the head of the AI development function.
- An organization that has identified an internal concern about whether an AI system is using an appropriate model may engage an academic researcher in AI models to assess the appropriateness of an AI system and report on findings to management.
- An organization that wishes to provide information to its business partners regarding the controls over its AI governance may wish to engage a certified public accountant to perform an evaluation of its controls and express an examination opinion in accordance with the AICPA *Attestation Standards*.

Q9: 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?

AI accountability mechanisms are still evolving, and various frameworks and practices are being developed and implemented across different sectors. Here are a few examples:

- *Technology sector:* Companies in the technology sector, particularly those developing and deploying AI systems, have undertaken initiatives to establish AI accountability mechanisms. Many large technology companies have published AI ethical principles or guidelines outlining their commitment to responsible AI development. These frameworks often cover areas such as fairness, transparency, privacy and robustness. There are also a number of industry collaborations focused on producing research, guidelines and recommendations on addressing the ethical and social implications of AI.
- *Financial services sector:* The financial services industry has been actively working on AI accountability due to the significant impact AI can have on financial markets and consumer experiences. Some international bodies have issued guidelines addressing ethical concerns, algorithmic transparency and risk management. Some examples include:
 - AI Governance Principles: The Financial Stability Board has released a set of AI governance principles that provide guidance on accountability, fairness, transparency and explainability in the use of AI and machine learning.
 - Guidelines on the Prudential Use of AI: The Bank for International Settlements' guidelines on the prudential use of AI in the financial sector provide recommendations on model management practices, robustness and governance.
 - Principles for the Ethical Use of Artificial Intelligence in Insurance: The International Association of Insurance Supervisors has developed a set of principles for the ethical use of

AI in insurance that focus on fairness, transparency, accountability and compliance with legal and regulatory requirements.

- *Healthcare sector:* The healthcare industry focuses on the importance of AI accountability in safeguarding patient safety, privacy and ethical use of data. Organizations like the World Health Organization have developed guidelines for AI in healthcare, focusing on issues like clinical validation, explainability and data governance. In addition, the US Food and Drug Administration (FDA) has released guidance on AI accountability in the context of medical devices.
- *Governmental bodies:* Governments are increasingly involved in developing AI accountability frameworks and regulations. For example, the European Union and Canada have released strategies and guidelines for responsible AI development and deployment. Singapore's AI Verify has developed certification schemes to allow developers and users of AI systems to have their procedures and systems assessed and certified by a third party.
- *Civil society organizations:* Various civil society organizations, advocacy groups and research institutions are contributing to the development of AI accountability frameworks. These organizations often focus on the ethical and societal implications of AI and advocate for transparency, accountability and human rights. Examples include the AI Now Institute and Future of Life Institute, as well as the Institute of Electrical and Electronics Engineers Standards Association (IEEE SA) Ethics Certification Program for Autonomous and Intelligent Systems (ECPAIS) and CertifAIEd program.

People currently working on AI accountability come from various backgrounds and include researchers, policymakers, industry professionals and civil society activists. We are actively collaborating with organizations such as research institutions, technology companies, government agencies, non-profit organizations and regulatory bodies to assist in the development of AI accountability frameworks, guidelines, standards and regulations. Additionally, our interdisciplinary collaborations involving experts in law, ethics, computer science and social sciences make important contributions to the development of AI accountability mechanisms.

Q10: What are the best definitions of terms frequently used in accountability policies, such as fair, safe, effective, transparent, and trustworthy? Where can terms have the same meanings across sectors and jurisdictions? Where do terms necessarily have different meanings depending on the jurisdiction, sector, or use case?

Transparency, safety, fairness and robustness are typically used to describe the goals of AI accountability mechanisms. While such terms are appropriate in providing overall direction in identifying risks associated with AI systems and establishing the types of boundaries to be established, they (alone) are too subjective to permit objective assessment. As part of the development of the accountability mechanism, such terms should be translated into requirements to provide criteria that meet the attributes of relevance, objectivity, measurability and completeness.

Q20: 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?

Developers and deployers of AI systems should maintain certain records and documentation to support AI accountability. While the specific requirements will vary depending on the context and regulations, the following are some examples of the types of records and documentation that can be important to support trust in AI:

- An organization's AI governance framework, which should include policies and procedure requirements for creating ethical AI
- Application design, including the purpose of the application and expected functionality
- Data selection and preprocessing, including documentation on the data sources, sampling methods, preprocessing steps, and data cleaning processes applied
- Model selection, including the criteria, processes, and considerations for selecting or designing the AI system
- Model training, including the training process, hyperparameters, convergence criteria and any performance evaluation metrics used
- Testing and validation, recording the testing methodologies, evaluation metrics and results obtained during the validation and performance assessment of the AI system
- System architecture and configuration
- System updates and patches, including complete logs of all changes
- System security design
- Monitoring mechanisms, such as the organization's controls to make sure the AI system continues to operate as expected, including how the AI is changed, how access to AI is managed and what monitoring is in place to make sure the AI application continues to operate as expected
- Communications to stakeholders, including user instructions, system processing specifications, descriptions and advertising, regarding the AI system's intended use methods of use, and intended functionality to support proper use and avoid unintended misuse.

The retention period for this documentation could range from a couple of years to over the life of the AI system, depending on the goals of the accountability system, legal requirements, regulatory frameworks, organizational policies and other factors. Consideration should also be given to the evidentiary requirements of assessments and audits. For example, when there are independent assessors, the documentation retention must be sufficient for an assessor to evidence that controls were in place over a defined period of time.

Given the complex nature of AI, humans play an important role in the accountability framework by guiding AI evolution in a way that incorporates accountability and governance. To foster accountability by design, AI systems can follow certain design principles and technical considerations, including:

- *Explainability and interpretability*: Design AI systems that can provide understandable explanations for their decisions and outputs, allowing users and stakeholders to comprehend the reasoning behind them
- *Data governance and transparency*: Establish clear data governance policies, including data provenance, documentation of data sources and usage, and transparency regarding data collection and handling practices
- *Robustness and resilience*: Build AI systems that are robust to adversarial attacks, noise, and varying operating conditions and implement safeguards to make sure the systems perform reliably even in challenging circumstances

- *Bias detection and mitigation:* Develop mechanisms to detect and mitigate biases in data, algorithms and decision-making processes, and regularly evaluate and monitor for fairness and bias issues throughout the system's lifecycle
- *Privacy and security:* Prioritize privacy protection and implement security measures to safeguard sensitive data and prevent unauthorized access or misuse
- *Model monitoring and continuous evaluation:* Establish monitoring mechanisms to track the performance, accuracy and behavior of AI models in real-world deployments, and continuously evaluate and assess the system's performance to identify potential risks or issues

By incorporating these design principles and technical considerations into the development process, organizations can proactively embed accountability measures into AI systems, promoting responsible and transparent practices from the outset.

Q23: How should AI accountability "products" (e.g., audit results) be communicated to different stakeholders? Should there be standardized reporting within a sector and/or across sectors? How should the translational work of communicating AI accountability results to affected people and communities be done and supported?

Communicating AI accountability "products," such as audit results, internal and external certifications or accountability assessment reports, to different stakeholders requires careful consideration to make sure there is transparency, clarity and effective understanding about what each "product" provides and does not provide to stakeholders. Here are some considerations for AI accountability communication:

- *Stakeholder engagement:* Identify the key stakeholders who would benefit from understanding AI accountability and involve them in the process. This can include industry peers, affected communities, advocacy groups, academia and researchers, and the general public.
- *Clear and accessible reporting:* Develop standardized reporting formats that present the AI accountability information in a clear, concise and accessible manner. Use plain language and avoid technical jargon to enhance understanding by different stakeholders.
- *Tailored communication:* Customize the communication approach to the specific needs and interests of different stakeholder groups. Consider their level of technical expertise, cultural background and preferred channels of communication.
- *Sector-specific and cross-sector standardization:* Explore the potential for sector-specific reporting standards that align with the unique characteristics and challenges of each industry. Additionally, consider cross-sector collaborations to establish harmonized reporting frameworks and facilitate comparisons and benchmarking across industries.
- *Multi-modal communication:* Use a variety of communication channels and formats to reach a wide range of stakeholders effectively. This can include written reports, infographics, visualizations, public consultations, public hearings, online portals and public presentations.
- *Community engagement and translational work:* Engage with affected communities and involve them in the communication process. Implement community outreach programs, public consultations and focus groups to make sure that AI accountability results are effectively translated and understood by those who may be impacted by AI systems.
- *Training and education:* Support training and education initiatives to enhance AI literacy among stakeholders. This can include workshops, webinars and educational resources that explain the concepts of AI accountability and enable stakeholders to meaningfully engage in the discourse.

- *Independent oversight and verification:* Consider involving independent third-party entities to provide additional credibility and verification of AI accountability reports. Independent oversight can enhance trust and confidence in the communicated results.
- *Feedback mechanisms:* Establish mechanisms for stakeholders to provide feedback, ask questions and seek clarification on AI accountability reports. This can include dedicated contact points, public forums or online platforms for engagement.

The form of communication may differ depending on the AI accountability product and the impact to stakeholders. For example, if the AI accountability product is reported on a webpage, assurance may be provided with a “seal” that communicates that the product has been evaluated against a standard set of expectations that has been evaluated by an independent third party. If the AI accountability product is part of generating input or performing controls in a process generating financial statement information, the communication may take on a more formal attestation report.

Standardized reporting should be considered where practical and in the context of the considerations listed above because it provides consistency and comparability and, if required, it should be based on formalized requirements and standards.

There should be clear communication on what AI accountability means to the organization and the different roles in the process. The organization that develops and maintains the AI system is responsible for the AI accountability product it creates and for the controls it puts in place. The attestation providers give reasonable assurance that the organization developed and maintain the AI system based on a standard framework. The users of the AI system have a responsibility to understand the AI system being used and its limitations.

Q34: 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?

The importance of uniformity in AI accountability requirements and practices can vary depending on the context. However, in general, there are several arguments in favor of uniformity across the US and global jurisdictions.

In the US, having uniform AI accountability requirements and practices can provide clarity and consistency for both developers and users of AI systems that operate in multiple states, as well as consumers. Uniformity establishes a level playing field and provides that that all entities operating in the country adhere to a common set of standards. This can be particularly crucial when it comes to addressing ethical concerns, privacy issues and potential biases in AI systems. Uniformity can also facilitate effective enforcement and regulation, making it easier for authorities to monitor compliance and for stakeholders to have a consistent assessment system in which to build trust in AI systems.

On a global scale, harmonization of AI accountability requirements and practices also can be beneficial for various reasons. First, it can facilitate international cooperation and collaboration in AI research, development and use. When countries share common standards, it becomes easier to exchange information, share best practices and jointly tackle challenges associated with AI technologies. This would allow for focused investments in the development of effective governance and control mechanisms for AI systems, an area that is currently struggling to keep pace with the design of AI functionality. Harmonization also can

help prevent regulatory fragmentation, where different countries or regions have conflicting or contradictory requirements, leading to inefficiencies and barriers to innovation.

Achieving uniformity or harmonization can be challenging due to the diversity of legal systems and regulatory approaches across jurisdictions. However, several approaches would help promote consistency and collaboration:

- *International cooperation and agreements:* Encouraging countries to collaborate and negotiate international agreements or frameworks on AI governance can foster uniformity. This could involve discussions on ethical principles, data protection, transparency and accountability.
- *Cross-sector collaboration:* Promoting collaboration among different sectors, such as academia, industry and government, could help identify common challenges and develop shared practices. This could be facilitated through the establishment of multidisciplinary working groups or expert panels.
- *Standardization efforts:* Encouraging the development of international standards for AI accountability could help establish common benchmarks. Standardization bodies and organizations can play a significant role in creating guidelines and best practices.
- *Knowledge sharing and capacity building:* Facilitating the exchange of knowledge, research findings and case studies could help raise awareness and promote a shared understanding of AI accountability. Capacity-building initiatives, including training programs and workshops, could enhance the expertise of policymakers, regulators and practitioners.

Although there are many benefits to promoting harmonization and interoperability of accountability requirements, complete uniformity across all sectors may not be practical since the risks of AI systems will differ depending on the use case, data type and impacts. Even within sectors, the risks of using AI will be heavily dependent on the use case. A risk-based approach may be most suitable to evaluate the AI risks and impact. This approach is then used to determine the governance and control mechanisms that should be put in place. These can be consistent across the same use case and AI risk profile. An important focus is establishing core principles and frameworks that can be adapted to specific contexts while making sure there is compatibility and cooperation among different jurisdictions.

* * * * *

Conclusion

We commend the NTIA for seeking feedback on the important topic of AI and accountability. As we have described, there are various levels of accountability measures that should be explored, and there are many stakeholders in this ecosystem. Similar to certain other emerging technologies, AI has incredible potential and risk, both of which must be managed effectively to promote successful implementation and public trust. We would welcome the opportunity for further discussion with the NTIA on AI accountability measures that are both operationally effective and cost efficient.

If you have further questions, please contact John Hallmark in the EY Office of Public Policy at john.hallmark@ey.com.

Yours sincerely,

Ernst & Young LLP

Ernst & Young LLP