TANGIBLE SECURITY

Preparing for AI Compliance

www.tangiblesecurity.com

Table of Contents

This eBook gives an overview of the current compliance environment for AI, and how this will lead to new compliance obligations for organizations deploying AI solutions. Important new standards include the NIST AI Risk Management Framework and ISO/IEC 38507, 24028, 23894, and 22989. Discover the common themes in these new standards, and the practical steps organizations should take to be prepared. Steps include governance, risk and security assessments, and how Tangible Security's services can help.

Coming compliance requirements for AI

The rapid rise of powerful, breakthrough AI technologies is generating excitement in the technology sector, but at the same time there are concerns about potential harms to safety, privacy, human rights, and other spheres of human society. Leaders in government as well as civil society worry that companies developing powerful AI systems may not be adequately assessing the risks new technologies could pose and want formal governance and transparency requirements to help assess and mitigate risks.

Traditional information technologies have long been the subject of both regulation and standards. Prominent examples include the GDPR law for privacy and the PCI DSS standard for ecommerce security. But new regulations and standards for AI are developing much faster, because of the unique aspects of AI technologies, such as the possibility that automated decision-making could have significant consequences for individuals and society. The large amounts of data used by AI applications raise special privacy concerns, and generative AI applications could be used to create deceptive or malicious content.

More than 800 measures are under consideration in 60-plus countries to regulate artificial intelligence.

According to Boston Consulting Group, there are more than 800 measures to regulate AI under consideration in over 60 countries. China, the EU, the UK, and the U.S. are all considering significant AI legislation. The International Organization for Standardization (ISO) committee on Artificial Intelligence currently has either published or is considering 55 separate standards addressing AI. Wherever these efforts lead, it now seems likely they will have a significant impact on the development and deployment of many AI applications in the coming years.

Emerging AI standards

The hundreds of proposed regulations around the world addressing AI include the EU AI Act, which would impose significant governance requirements and restrictions on certain AI technologies. The scope of this eBook is limited to the most significant compliance standards, specifically the U.S. National Institute of Standards and Technology (NIST) AI Risk Management Framework, and a subset of the most consequential standards of the International Organization for Standardization/International Electrotechnical Commission (ISO/IEC).

These standards will likely be influential for the governance of AI technologies, and so should be of particular concern to businesses deploying or considering deploying AI. The NIST Framework defines AI systems as “engineered or machine-based systems that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy.”

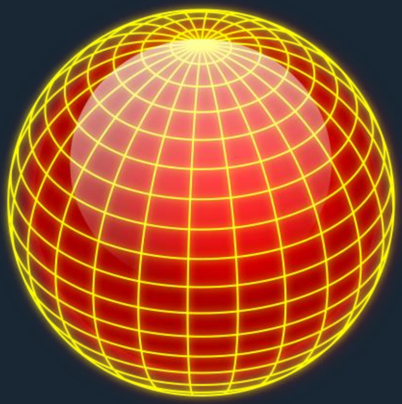

The ISO has either published or is considering 55 separate standards addressing AI.

NIST AI Risk Management Framework 1.0

The NIST is an agency of the U.S. Department of Commerce dedicated to advancing measurement science, standards, and technology. The NIST AI Risk Management Framework 1.0 is a set of guidelines released on January 26, 2023. The Framework is intended to help organizations design, develop, use, and evaluate artificial intelligence (AI) systems in a trustworthy and responsible way. The Framework aims to address the various risks to privacy, security, and safety that AI systems may pose to individuals, organizations, and society.

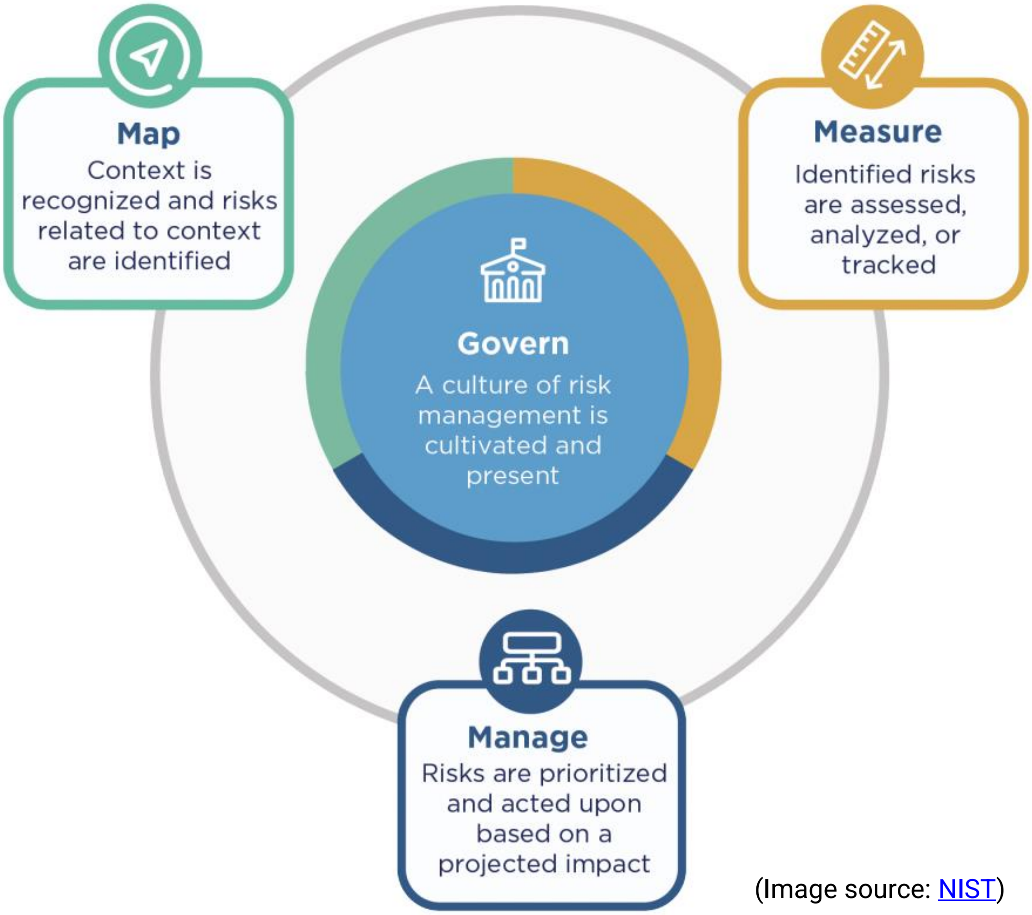
Map
Context is recognized and risks related to context are identified
Measure
Identified risks are assessed, analyzed, or tracked
Govern
A culture of risk management is cultivated and present
Manage
Risks are prioritized and acted upon based on a projected impact

(Image source: NIST)

The Framework is divided into two parts. Part 1 discusses how organizations can frame the risks related to AI. Part 2, the core of the Framework, consists of four functions: Govern, Map, Measure, and Manage. Govern applies across all AI risk management processes and procedures, and the Map, Measure, and Manage functions can be applied at specific stages of the AI development process. Each function is further broken down into specific sub-functions.

NIST Framework: Core functions

Organizations can leverage the Framework to apply the functions best suited to their environments, deploying them at various points during the AI lifecycle and selecting the sub-categories appropriate to their needs.

Govern

The Govern function cuts across the entire risk management process through a series of processes, documents, and organizational schemes to help identify and manage risk across the entire system lifecycle. Govern also provides a structure that can align risk management and technical functions with organizational principles, and thereby cultivate a culture of risk management in organizations developing or deploying AI systems.

Map

The Map function establishes the context of the AI system by identifying the purpose, use, laws, norms and expectations, and deployment settings. The Map function includes the categorization of the system, system capabilities, usage and goals, risks and benefits, and impacts to individuals, groups, communities, organizations, and societies. Outcomes in Map form the basis for the Measure and Manage functions.

Measure

The Measure function employs quantitative, qualitative, or mixed-method tools to assess, benchmark, and monitor AI risks. Measure leverages information gathered in the Map function. Risks measured include trustworthy characteristics, social impact, and human-AI configurations. Processes used in Measure can include software testing, performance benchmarks, formalized reporting, and documentation of results.

Manage

The Manage function prioritizes risk and allocates resources to address risk as defined by the Govern function. Manage categories include the prioritization of risks as derived from Map and Measure, strategies to maximize AI benefits and minimize negative impacts, management of risks and benefits from third-party entities, and risk treatments, and include response, recovery, and communication plans.

Read the full NIST AI Risk Management Framework.

ISO/IEC Standards

ISO (International Organization for Standardization) is an independent, non-governmental international organization with a membership of 169 national standards bodies. The organization develops and publishes international standards, including many foundational standards familiar to IT professionals.

In 2017, ISO created the ISO/IEC JTC 1/SC 42 committee to develop international standards for AI. Since then, the committee has published 20 standards with 35 under development.

The standards developed by the committee are intended to have broad application across the entire AI ecosystem and across industries. They address such foundational topics as governance (ISO/IEC 38507), concepts and terminology (ISO/IEC 22989), and system life cycle (ISO/IEC 5338).

Because of the significance of AI technology, the standards developed by the committee also consider non-technical factors, including regulation, business models, societal impacts, and ethical issues.

Like existing ISO standards for IT, the emerging ISO standards for AI will provide a framework for managing information security, data privacy, service quality, and customer satisfaction. Additionally, they will likely be a key part of compliance with legal and regulatory requirements, as well as to show competence to customers, suppliers, and stakeholders.

The committee maintains a continuously updated list of ISO/IEC standards.

Key Published Standards

ISO/IEC 38507:2022 Governance of the use of AI

Provides guidance on the role of a governing body deploying AI within its organization and encourages it to use appropriate standards to underpin governance of the use of AI.
https://www.iso.org/standard/56641.html

ISO/IEC TR 24028:2020 Trustworthiness in AI

Surveys existing approaches that can improve trustworthiness and their potential application to AI systems, as well as approaches to mitigating AI system vulnerabilities related to trustworthiness.
https://www.iso.org/standard/77608.html

ISO/IEC 23894:2023 Guidance risk management

Intended to be used in connection with ISO 31000:2018 – Risk Management Guidelines. Contains three main parts: principles, a framework, and risk management processes applied to AI.
https://www.iso.org/standard/77304.html

ISO/IEC 22989:2022 AI concepts and terminology

Provides standardized concepts and terminology to improve the understanding of AI technology and to promote use by a broader set of stakeholders.
https://www.iso.org/standard/88145.html

Standards in Development

ISO/IEC 5338 AI system life cycle processes

Will define a set of processes and associated concepts for describing the life cycle of AI systems based on machine learning and heuristic systems.
https://www.iso.org/standard/81118.html

ISO/IEC 42001 Management system

Will provide a certifiable AI management system framework within which AI products can be developed as part of an AI assurance ecosystem.
https://www.iso.org/standard/81230.html

Common themes in AI compliance

The emerging compliance standards for AI in many ways leverage and build on existing standards for information technologies. The ISO/IEC 38507 standard for governance of AI derives from the earlier ISO 37000 standard of Governance of organizations, and the NIST AI Risk Management Framework builds upon the NIST Cybersecurity Framework and Risk Management Framework. Like earlier standards, emerging AI standards focus on risk management practices, security, reliability, and interoperability.

The unique nature of AI systems leads to some important differences from earlier standards. Because AI systems can have a larger societal impact, compliance standards for AI involve a broader and more inclusive number of stakeholders, and are more focused on the ethical, social, and legal implications of AI systems, such as fairness, transparency, accountability, privacy, and security. Because AI systems change rapidly, compliance standards for AI are designed to be more dynamic.

Trustworthy

Organizations need to ensure systems are trustworthy by accounting for risks to safety, security, privacy, and accountability.

Dynamic

AI systems are dynamic and involve continuous learning and updating, so processes must be regularly reviewed and revised.

Inclusive

AI systems should be developed with the input of a broad array of stakeholders across the entire AI lifecycle.

Governance

Governance and oversight must be robust, and humans must ultimately be accountable for decisions and outcomes

Responsible

Organizations should consider potential harms to individuals and specific groups including civil liberties and economic opportunity.

How to prepare for AI compliance

The unique aspects of AI compliance, such as the need for greater care, broader stakeholder involvement, and continual review have been emphasized throughout this eBook. Organizations should prepare for AI compliance requirements, such as the need for greater care, broader stakeholder involvement, and continual review, by implementing practices that map to the main themes of emerging AI standards.

Create a V-team

Governance programs for AI should involve an array of stakeholders that includes business units, IT, legal, HR, marketing, and PR to ensure the broadest set of perspectives to analyze the risk of harms to individuals, society, and to the reputation of the business.

Initiate governance

Develop a responsible AI initiative that follows a set of principles that stress accountability, transparency, privacy, security, fairness, and inclusiveness in the development and deployment of AI systems.

Document compliance

Create technical documentation, processes, and policies, that employees can follow. Document these compliance initiatives in preparation for third-party audits.

Regularly assess risk

Build regular impact assessments into your AI systems to identify and mitigate any potential risks, such as bias, error, or harm. Establish a risk management system, including regular testing, for the life cycle of the system.

How Tangible Security can help

For over 25 years, security-minded organizations have trusted Tangible Security with protecting their sensitive assets. We offer a full range of services from penetration testing and risk assessments to staff training, compliance assessments, and staff augmentation such as fractional CISOs that will ensure that security in your organization becomes tangible.

Governance, risk, and compliance consulting

Tangible Security's GRC consulting services provide organizations with expert guidance and support in managing your governance, risk, and compliance initiatives. Our consultants work closely with organizations to develop and implement effective GRC frameworks, policies, and processes tailored to your specific needs and industry requirements. We offer expertise in risk assessment, regulatory compliance, policy development, and monitoring, helping organizations streamline their operations, mitigate risks, and ensure adherence to legal and industry standards. With our GRC consulting services, organizations can enhance governance practices, strengthen risk management capabilities, and achieve comprehensive compliance across operations.

AI Technical Security Assessment

Tangible Security's AI Technical Security Assessment service provides a comprehensive evaluation of your AI-powered tools, adhering to the OWASP Top 10 Framework for Large Language Model applications and other critical areas. Our team of AI and cybersecurity experts delivers detailed reports on vulnerabilities and actionable recommendations, ensuring the security of your AI tools and safeguarding internal and confidential data. By leveraging our expertise, you can streamline the process of creating a secure environment, maintain compliance, and align business needs within budget constraints, all while maximizing the potential of cutting-edge productivity tools. Our service grants immediate access to a professional team of cybersecurity leaders, backed by a published AI researcher with extensive software engineering experience.

TANGIBLE SECURITY

We can help

To learn more about how we can help you prepare for AI compliance, contact us:

info@tangiblesecurity.com
1 800 913-9901
7408 Knightdale Blvd Ste 220B
Knightdale, NC 27545
www.tangiblesecurity.com

About Tangible Security

Tangible Security is a full-service cybersecurity services firm providing advanced protection to customers' sensitive data and infrastructure. Applying an attacker's mindset and innovative methodologies, our team of experts implement tailored security solutions that make security tangible to each customer. Founded in 1998 and headquartered in North Carolina, Tangible Security provides services to industries throughout North America.