Adopt 5 Key Privacy Lessons When Preparing for AI Regulations

2 August 2024 - ID G00817005 - 17 min read

By Bernard Woo, Bart Willemsen

Major economies have passed or are actively debating AI regulations to prevent harm to people and/or society. Security and risk management leaders can use this research to help ease AI governance efforts and enhance AI project success while facilitating compliance.

Overview

Key Findings

- The number of AI regulations worldwide is increasing, with the EU and U.S. states enacting laws in 2024, and multiple jurisdictions actively debating draft bills.
- Laws regulating the use of AI required that organizations track how the technology is adopted and used, both in current deployments and future efforts.
- AI regulations require organizations to focus first on the potential harm to individuals, thus building on top of principles found in modern privacy regulations and human rights. Other common themes like the need for organizations to provide transparency with respect to the use of the technology are included as well.

- Elements of the privacy program are fundamental precursors and can be extended to AI governance efforts, allowing organizations to gain efficiencies and agility in their programs.

Recommendations

As organizations prepare for the regulated use of AI, security and risk management (SRM) leaders should guide AI governance efforts through:

- Building an inventory of AI uses and projects by connecting with sponsoring business functions to identify all current and future deployments.
- Determining which AI adoption efforts affect individuals by identifying the purpose(s) for each use case, with records of processing activities (RoPAs) representing a potential short cut.
- Identifying the potential harm to individuals by expanding the privacy impact assessment process to include an assessment of fundamental rights and the AI system(s) to be used.
- Reducing the risks associated with the processing of personal data by proactively applying privacy-enhancing technologies (PETs) and other elements of AI trust, risk and security management (AI TRiSM).
- Providing transparency to individuals by giving notice about the use of AI technology and allowing avenues for people to exercise their rights.

Strategic Planning Assumption

By 2026, 50% of governments worldwide will enforce use of responsible AI through regulations, policies and the need for data privacy.

Introduction

Interest in the potential of AI technology has exploded since OpenAI's ChatGPT arrived on the scene in late 2022.

According to the 2024 Gartner Growth Agenda Survey, 47% of executive leaders involved in their company's growth initiatives identified generative AI (GenAI) as the most critical technology for achieving high revenue growth. Simultaneously, the second-ranked most important technology was privacy and security. [1]

In parallel, governments around the world are moving to regulate the use of AI technologies. The EU passed the EU AI Act in 2024, joining China in having a far-reaching law on the books. [2,3] Meanwhile, in the U.S., a Presidential Executive Order demanding the use of trustworthy AI was announced in late 2023 with some sector-specific laws in place and other draft bills actively debated in Congress. [4] Individual states are also taking action; for example, the state of Colorado enacted a "comprehensive" law of its own in 2024. [5] Other countries, including Canada, have announced similar developments of AI-focused regulations. [6]

One common driving force for the AI regulations, whether passed or under consideration, is a need to protect the safety and well-being of individuals and society at large. They also help to establish a common set of ethical standards and practices that benefit all parties. Such concerns are captured and reflected in the AI principles that the Organisation for Economic Co-operation and Development (OECD) have adopted. [7] This ought to come as no surprise. After all, AI works based on data, and any lack of control over data given a specific purpose or intent will be amplified by AI usage to extreme levels.

It is noteworthy that at the heart of AI principles is a declaration that individual rights, including privacy, be respected.

This overlap between AI principles and privacy, and the emphasis of having to be in absolute granular control of both, means that organizations use their privacy and data protection programs as a starting point when embarking on AI governance efforts. Gartner has observed, via client interactions, that some organizations have already tasked the privacy office with leading AI governance efforts. Furthermore, this overlap seems to have been recognized at the highest levels of the organization.

Data from the 2024 Gartner Board of Directors Survey on Driving Business Success in an Uncertain World indicates that "data privacy or security risks" was identified by nonexecutive directors as the second-highest concern (after "information integrity") for 2024 through 2025 with respect to GenAI risks. [8]

SRM leaders, particularly those with AI or privacy responsibilities, must recognize the overlap between AI principles and privacy. In doing so, they lead their organizations to apply the knowledge gained from five key privacy lessons to enhance AI governance efforts and facilitate success for AI projects (see Figure 1).

Figure 1: Five Privacy Lessons When Preparing for AI Regulations

Five Privacy Lessons When Preparing for AI Regulations

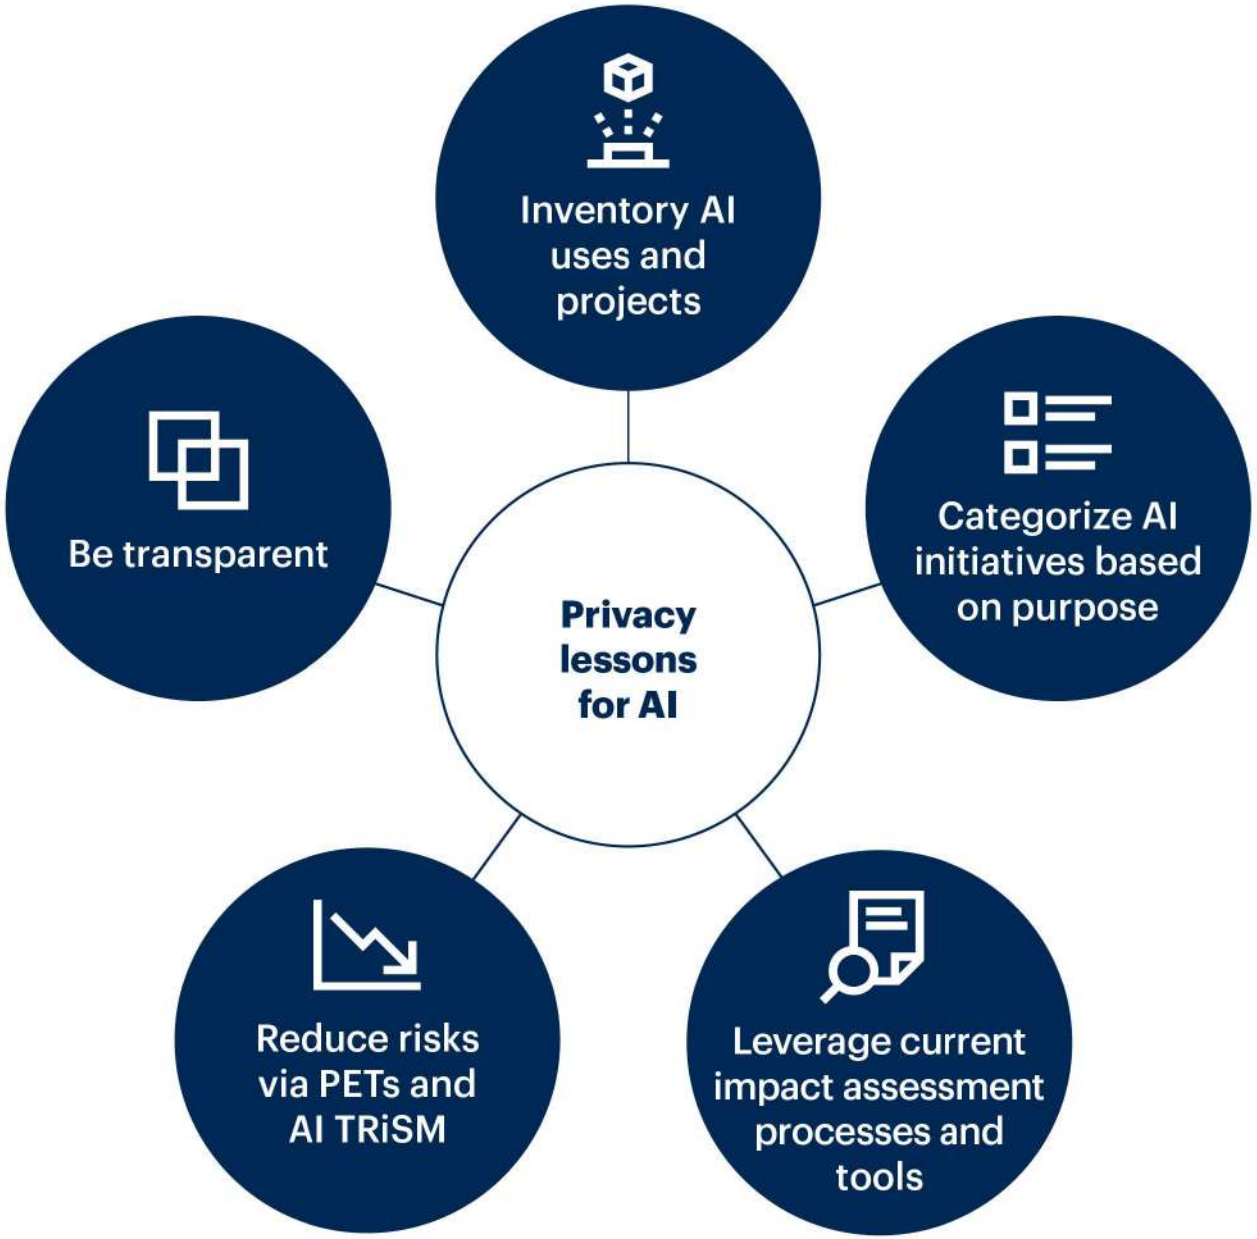

PETs = privacy-enhancing technologies

TRiSM = trust, risk and security management

Source: Gartner
817005_C

Gartner

Analysis

Lesson 1: Know What You Have (and Maintain It!)

The starting point for both the privacy program and AI governance efforts requires an organization to have detailed knowledge over where personal data is processed and which AI projects are in progress or being planned.

For your privacy program, the knowledge of where and which personal data is processed is critical. At minimum, this enables your organization to respond to subject rights requests (SRRs), such as access to data, correction and deletion. Beyond this, having and maintaining a data inventory allows the organization to perform myriad tasks, including:

- Classifying the sensitivity of the data processed.
- Identifying when data protection/privacy impact assessments (PIAs) are needed.
- Determining purpose-based access controls and retention periods, and lineage for data governance efforts.

Data classification, though not a new topic, is turning out to be a crucial factor in AI projects. Building on top of data discovery efforts, the proper labeling of data enables organizations to implement controls over the movement of data. This turn provides an avenue for limiting which data elements, as well as ensuring their proper labeling, are passed into large language models (LLMs).

For more, see How to Succeed With Data Classification Using Modern Approaches.

Similarly, building an inventory of AI projects provides your organization with the foundational knowledge needed to then identify where risk exists, including situations where regulations may apply (such as those that involve processing personal data).

To build an inventory of AI uses and projects, you need help. Start by reaching out to leaders of business functions to explain the changes coming from the regulatory landscape and the need for partnerships now to prepare for the oversight by regulators. The effect is to be bidirectional — to both have insight into existing usage that might have started unnoticed, and to proactively gain control over models that are intended to be used at a later stage.

Build a process or system for maintaining this inventory, as different business functions continue to experiment with AI technologies. This is where automated data discovery tools can provide much needed scalability and efficiency. The use of such tools can help identify not only where sensitive data (including personal data) is processed, but trace its movement as well.

Other tools that can help the discovery efforts include security service edge (SSE) and security information and event management (SIEM), which track the usage of applications and provide hints as to where sensitive data may be processed. Once the repositories are identified, putting together a map of sensitive data flows provides SRM leaders with greater efficiencies in identifying AI-related workloads (see **7 Ways Privacy Can Drive Data Mapping Initiatives** for an example). Then, you can initiate the workflows needed to determine potential risks and regulatory obligations. This is particularly important when it comes to identifying situations involving third parties that may be using your organization's data for their foundation models.

As an added bonus, such a map can be used for other purposes, such as helping security teams home in on high-value data assets that require additional controls or the building of data catalogs to support data- or information-governance-related initiatives.

Lesson 2: Identify the "Why" — What Is the Purpose?

A key tenet within the privacy world revolves around the purpose(s) for processing personal data. For example, data about an individual's health condition is among the most sensitive of information. Most times, access to such information is strictly limited. But, if the person involved is about to undergo surgery, medical professionals who would normally be prevented from accessing the data may be provided access in order to ensure the procedure is successfully carried out.

Similarly, AI regulations are focused on identifying the uses of technology and placing limits around those that can harm individuals. The EU AI Act, for example, groups use cases into four categories based on risk of harm to individuals (see **Getting Ready for the EU AI Act, Phase 1: Discover & Catalog**). Those use cases with the highest level of potential harm to individuals are considered undesired, and organizations are generally prohibited from using them.

For existing AI usage, once an inventory of AI use cases has been compiled, the next step is to identify the intended purposes. For newly intended ones, this obviously needs to be switched around, and establishing purpose comes first. A potential short cut exists for organizations that have put together RoPAs as part of their privacy programs (a sample template is available from the U.K. **Information Commissioner's Office**). A RoPA is intended to be a catalog of all the processes that handle personal data. Details recorded include the process owners and purposes by which personal data is handled, and the data elements involved, retention periods and parties with whom information has been shared. Such information can prove invaluable when a particular AI use case is using or seeking to use the personal data.

Even when a RoPA does not exist, looking into the elements one would need can serve as guidance for the information you want to gather about AI use cases.

Be deliberate with this step. There will likely be one or more instances where the processing purpose is simply to identify previously unknown patterns. Given that AI regulations focus on potential harm to individuals, purposes that are not clearly defined or represent "fishing for the unknown" exercises carry significant risk — for both the organization and the people involved. Your organization will need to take steps to address the potential harms to individuals, which means a risk assessment is necessary.

The good news is, you likely have existing practices (see Lesson 3 section).

Lesson 3: Use What You Have — Extend the PIA Process

While these are early days in terms of AI regulations, with additional guidance to be developed by regulators, a pattern is emerging where such laws demand that users and developers of AI technologies focus on the potential harms to individuals. The EU AI Act, for example, requires that a Fundamental Rights Impact Assessment (FRIA) be completed for "high-risk" AI use cases. [9] As its name implies, the FRIA focuses on AI use cases and determines if they negatively affect individuals at large.

Before your organization rushes to create a new process, examine the PIA process from your privacy program (see **Use a Privacy Impact Assessment to Ensure Baseline Privacy Criteria**). A PIA analyzes a given process that handles personal data, determines the potential risks to individuals and identifies mitigating measures. Many of the questions asked in the PIA are likely to be easily extendible to examine AI use cases. In fact, the EU AI Act has seemingly recognized this by indicating that a FRIA can exist as a complement to the assessments demanded by the General Data Protection Regulation.

As your legal team examines the details in AI regulations, take note of the templates that are being developed for FRIAs, examples of which include the OECD's Catalogue of Tools & Metrics for Trustworthy AI and the Dutch government's Impact Assessment Fundamental Rights and Algorithms. Then, determine how your current PIA process (including templates, such as **Toolkit: Assess Your Personal Data Processing Activities**) can be extended so you continue to have one tool and process.

Lesson 4: Reduce Risk by Implementing PETs and Security Measures

As you carry out FRIAs/PIAs, you will need to start examining how privacy risks can be reduced. The first question in this regard is always: Is it necessary to use personal data? After all, the need for AI models to train using vast amounts of data puts them at a diametric opposite to the principle of data minimization where personal data is involved. Therefore, the first risk reduction question should always be whether the models can be trained without using identifiable personal data. Contrary to popular belief, models can often be architected to function on altered or transformed data with the same precision in operation.

And, if it is necessary to train on data with some level of identifiability (for example, where one needs to separate data belonging to different individuals within the model), then consider the use of PETs (see **Three Critical Use Cases for Privacy-Enhancing Computation Techniques**). Opportunities for PET usage occur in training, architecting and operating a model.

PETs represent a family of techniques that allows for the reduction of privacy risks while supporting different workloads that involve processing sensitive data. Gartner has observed, via conversations with clients, techniques such as synthetic data, federated machine learning (ML) and differential privacy being adopted for AI model training. Sometimes, these are combined — for example, using federated ML to generate synthetic data for training a model. However, it is important to note the focus cannot be on the ingestion of data only. The power of AI is such that even when you have taken precautions and transform the data prior to its ingestion into models/algorithms, reidentification can occur. Therefore, ensure that you have a process to examine all outputs to observe that anonymity of individuals is preserved or that such results are discarded to avoid bringing unintended consequences to individuals.

One emerging technique for managing AI risks, particularly those related to the use of GenAI and the potential of inaccurate conclusions ("hallucinations"), is to deploy augments to control what goes into the prompts and correct/restrict the output that is returned. For more, see How CISOs Are Supercharging Their Teams With Generative AI Augments.

In addition to the use of PETs, it is necessary to implement security measures at other parts of AI workloads. AI technologies, particularly projects involving GenAI, bring new risks in three categories: content anomalies, data protection and AI application security (see **Innovation Guide for Generative AI in Trust, Risk and Security Management**). These risks manifest themselves in attack surfaces as shown in Figure 2.

Figure 2: Generative AI Attack Surfaces Across the AI Life Cycle

Generative AI Attack Surfaces Across the AI Life Cycle

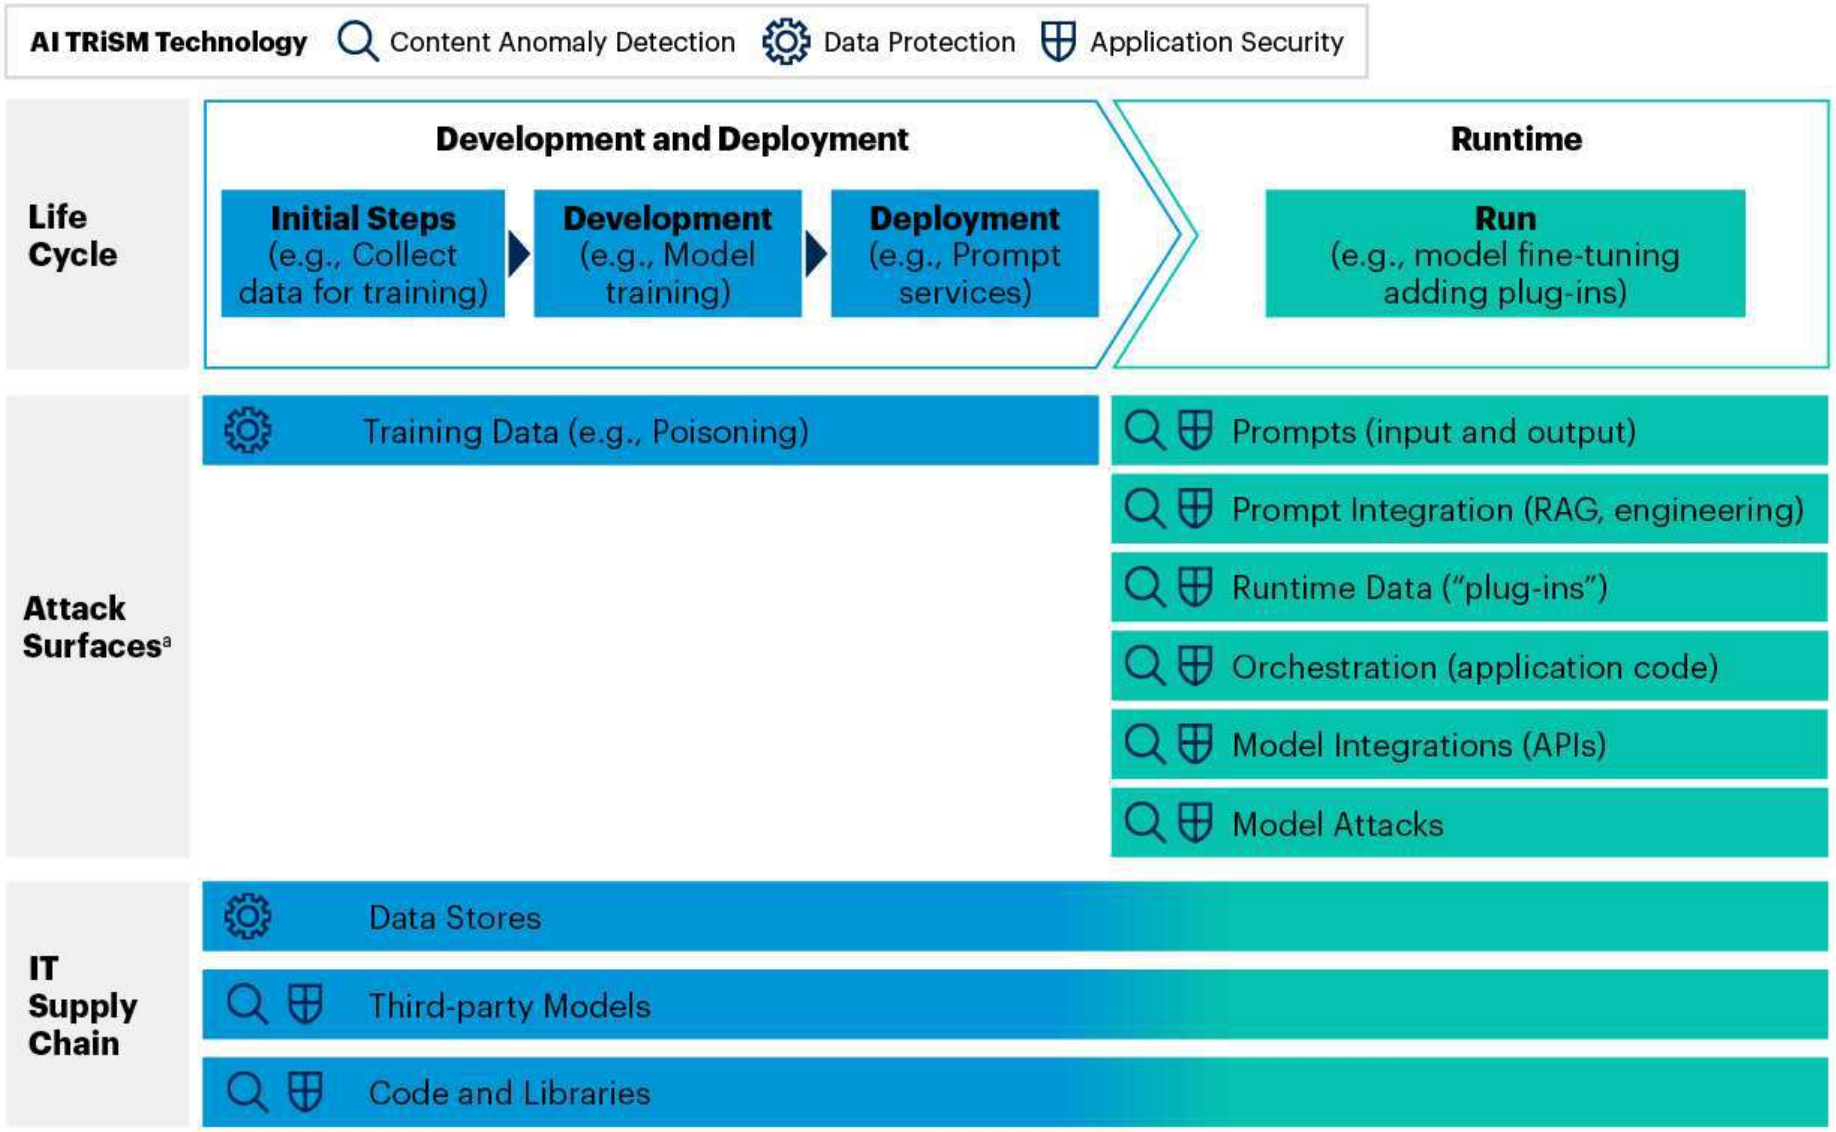

Source: Gartner

[a] Main sample attack surfaces only; others not shown

796422_C

Gartner

It is important to note that all elements covering AI and security risks have yet to be identified — perhaps even by potential attackers themselves! Organizations need to find a balance — be proactive and anticipate when/where attacks may occur, but do not become so preoccupied that it severely restricts attempts to test AI technologies.

While the use of PETs (and augments) can lower the risks of associated content anomalies and data protection, they cannot be assumed to be sufficient when facing unknown attack methods. Furthermore, they do not protect the AI applications/models. SRM leaders should lead their organizations in the adoption of all elements of Gartner's AI TRiSM framework, which encompasses a set of governance processes and emerging technology components (such as PETs) for managing the risks associated with AI technologies. For more, see **Top Strategic Technology Trends for 2024: AI Trust, Risk and Security Management**.

Lesson 5: Be Transparent — Just Like the Privacy UX

Finally, in addition to internal transparency regarding what is used, where, how and why, it is also elemental to provide external transparency — such as to consumers, patients, citizens and staff — regarding AI usage. External transparency allows them to decide if and how they want to engage with you, let you engage with them, and/or let you potentially expose their data to AI models.

This is yet another similarity to privacy regulations, where a core tenet is that organizations are open with individuals about how they (organizations) intend to process personal data and the practices they have in place to protect people's privacy. This permits individuals to then make an informed choice about whether they wish to engage with the organization. On this front, Gartner has been advising clients on the necessity for a privacy user experience (UX) (see **State of Privacy: The Privacy Tech Driving a New Age of Data Wealth**).

Here again, you and your organization can apply the lessons learned from helping the privacy UX to bring transparency to the use of AI technologies. For example, when it comes to presenting information about your organization's use of AI, take care to use simple, easy-to-understand language. Furthermore, this should be "just-in-time" messaging, to the point of interaction only, and as brief yet complete as possible. There is a place for more legal-oriented language (this cannot be avoided) but where possible, respect your (potential) customers and employees by using language they can easily digest. Strengthen your processes, especially around responding to requests to delete personal data ("right to be forgotten") or limiting the processing of personal data that lead to automatic decision making about individuals (see **Market Guide for Subject Rights Request Automation**).

Evidence

[1] **2024 Gartner Growth Agenda Survey.** This study was conducted to better understand which tactics and approaches differentiate high-growth companies (those who achieved revenue growth of 10% or more) from others. The research was conducted online from 4 October through 27 November 2023 among 288 executive leaders from North America (n = 105), Latin America (n = 35), Europe (n = 104) and Asia/Pacific (n = 44), across all commercial industries, excluding information technology. Respondents were from companies with at least $250 million or equivalent in annual revenue. Qualified respondents had personal knowledge of their company's financial performance and either led or participated in their company's growth initiatives. Disclaimer: Results of this survey do not represent global findings or the market as a whole but reflect the sentiments of the respondents and companies surveyed. There are no respondents from China in compliance with the Personal Information Protection Law (PIPL).

[2] **Artificial Intelligence Act: MEPs Adopt Landmark Law**, European Parliament.

[3] **China's New Rules For Generative AI: An Emerging Regulatory Framework**, Fasken.

[4] **Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence**, The White House.

[5] **From Brussels to Boulder: Colorado Enacts Comprehensive AI Law on the Heels of EU's AI Act with Significant Obligations for Business and Employers**, Connect On Tech.

[6] **An Act to Enact the Consumer Privacy Protection Act, the Personal Information and Data Protection Tribunal Act and the Artificial Intelligence and Data Act and to Make Consequential and Related Amendments to Other Acts**, Parliament of Canada.

[7] **OECD AI Principles Overview**, Organisation for Economic Co-operation and Development.

[8] **2024 Gartner Board of Directors Survey on Driving Business Success in an Uncertain World.** This survey was conducted to understand the new approaches adopted by nonexecutive boards of directors (BoDs) to drive growth in a rapidly changing business environment. The survey also sought to understand the BoDs' focus on investments in digital acceleration; sustainability; and diversity, equity and inclusion. The survey was conducted online from June through August 2023 among 285 respondents from North America, Latin America, Europe and Asia/Pacific. Respondents came from organizations with $50 million or more in annual revenue in industries

except governments, nonprofits, charities and nongovernmental organizations (NGOs). Respondents were required to be nonexecutive members of corporate boards of directors. If respondents served on multiple boards, they answered questions about the largest company, defined by its annual revenue, for which they are a board member. Disclaimer: The results of this survey do not represent global findings or the market as a whole, but reflect the sentiments of the respondents and companies surveyed.

[9] **Fundamental Rights Impact Assessment Under EU AI Act**, Van Bael & Bellis.