


WHITE PAPER

THE EU AI ACT

GDPR for the 2020s, or the
Limits of the Brussels Effect?

A SANS Institute Discussion
Paper on What the EU AI Act
Means for Your Company



Introduction

For companies based in the European Union (EU), or doing business there:

01. Is the new EU Artificial Intelligence Act the same kind of seismic moment for the 2020s that the General Data Protection Regulation (GDPR) was for the 2010s?

02. More practically, what must companies do to comply with it?

03. How can they do so in a way that optimises their global effectiveness?

This exploratory discussion paper from SANS Institute sets out some of the key questions arising from the EU AI Act as its multiyear implementation timetable takes effect.

Alongside this paper, we are launching a survey to learn more about how companies are approaching this crucial strategic issue. The findings will form a crucial part of a major white paper on EU technology regulation, to be published in early 2027.

Overview

In recent decades, Europe has not been the epicentre of technological innovation and production. But for more than a decade, it has been the world's foremost power in setting and enforcing rules for modern technology and its use. The high point of this so-called 'Brussels Effect' was the introduction of the GDPR in 2018.

The GDPR was neither the world's first data protection law nor its toughest. But it became the embodiment of the Brussels Effect. Multinational corporations found it economically impractical to maintain separate standards for different regions. To gain access to the EU's massive, affluent market of 450 million consumers, many companies chose to adopt the EU's strict regulations across their global operations. Whether or not a company was European, it often complied with the GDPR.

In the 2020s, the focus is Artificial Intelligence (AI). As companies rapidly adopt AI, Europe's structural innovation challenges, including weaker and smaller capital markets and barriers to creating hyperscalers, have become more pronounced than ever. These challenges were set out in detail in former Italian Prime Minister Mario [Draghi's report on EU competitiveness](#).

At the same time, EU institutions have sought to replicate their achievements of the previous decade by providing the first comprehensive AI rulebook. When provisional agreement on the EU AI Act was reached at the end of 2023, European Commission President Ursula von der Leyen [called](#) it 'an historic moment' and described it as the first ever comprehensive framework on artificial intelligence *worldwide*.

In this paper:

- First, we look at the GDPR story and what businesses learnt from it about the implications of EU technology regulation.
- Second, we consider how the EU AI Act is taking shape.
- Third, we identify the key questions facing businesses as they grapple with AI adoption in the years ahead. These questions also form the basis of our survey.

Europe has not been the epicentre of technological innovation and production. But for more than a decade, it has been the world's foremost power in setting and enforcing rules for modern technology and its use.

Part 1: The GDPR Story

GDPR was, and remains, a principles-based regulation that required significant up-front work to establish the structures needed for compliance. Once those structures are in place, ongoing compliance must still be managed, but it is generally straightforward as long as the organisation's business and data-processing activities remain stable.

The GDPR sets out seven principles: lawfulness, fairness, and transparency; purpose limitation; data minimisation; accuracy; storage limitation; integrity and confidentiality; and accountability.

In practice, the regulation required every organisation handling the personal data of people in the EU to do five things:

01. Know its data.

Organisations had to map what personal data they held, where it was stored, how it flowed between systems and third parties, and the legal basis on which it was processed. For many organisations, doing this for the

02. Build structures for ongoing compliance.

This meant establishing a lawful basis for data processing activities and creating processes to support new individual rights, including access to personal data, the erasure of certain data and data portability.

03. Develop documentation and governance capabilities.

Organisations had to record processing activities and risk assessments. Many were also required to appoint a data protection officer.

04. Manage supply chains and vendors.

The GDPR required third parties to provide 'sufficient guarantees' that suitable arrangements were in place for broader compliance.

05. Establish an incident response capability.

Organisations needed a tested process for notifying regulators within 72 hours, where required, and managing the consequences of a personal data breach.

The Impact on Businesses

Around all this was the visible infrastructure of GDPR compliance, including privacy notices, cookie consent mechanisms and organisation-wide staff training.

It wasn't cheap. The most widely cited estimate, from a 2018 survey by the International Association of Privacy Professionals and EY, put the combined preparatory spending of the world's 500 largest companies at approximately \$7.8 billion, or an average of roughly \$15.8 million each.

Large British firms in the FTSE 350 were estimated to have spent approximately \$1.1 billion between them. Costs for individual enterprises ranged from under \$2 million to \$70 million, depending on the complexity of their legacy systems and the extent of their data processing activities.

American and other non-European companies tended to spend more than their European counterparts because much of the GDPR's substance already existed in earlier EU law, giving European companies a head start. Across both EU and non-EU companies, however, compliance presented a material, board-level cost.

Did it work? Assessing that depends on the outcomes being measured. The GDPR certainly didn't stop data breaches. They still happen, and according to some analysis, the trends are moving in the wrong direction, although this may partly reflect improvements in reporting and enforcement.

However, the GDPR has certainly made privacy a board room issue and created permanent data protection and security functions within organisations.

Enforcement, done at the national level by designated regulators, has resulted in fines of approximately €7 billion. However, this total is heavily concentrated in a small number of very large penalties against a handful of major technology platforms. Several of those penalties are under appeal, and the sums collected are a fraction of the sums announced.

The fact that most organisations haven't been fined is not evidence that the regulation has been ineffective. There is abundant evidence that the GDPR forced major organisational change. What it has shown is that headlines about punitive, ruinous turnover-based fines are not the best way to think about the impact of EU regulation, as such fines rarely materialise.

Headlines about punitive, ruinous turnover-based fines are not the best way to think about the impact of EU regulation, as such fines rarely materialise.

\$15.8M

average prep spend across the world's 500 largest companies

€7B


in total fines

\$1.1B

estimated spend by large British firms in the FTSE 350

As we look ahead to the era of the EU AI Act, three general conclusions can be drawn from the GDPR experience:

- EU technology regulation attached to access to the single market is not optional. Companies doing business in the EU must comply. For EU companies, it means there is no choice. For non-EU companies, the choice is between adopting EU regulations in their global business model, establishing separate arrangements for EU business, or withdrawing from the EU market. In the case of the GDPR, most big global companies tended towards the first option, making the regulation the epitome of the Brussels Effect.
- The lasting impact of EU technology, security, and privacy regulation is not found primarily in landmark court cases or fines. It lies in the extensive, unglamorous changes to organisational functions and processes that become permanent.
- There is some evidence those organisations that took GDPR implementation seriously improved their cyber security posture and outcomes.



Part 2: The EU AI Act – The Emerging Story

The EU AI Act, Regulation (EU) 2024/1689, is the first comprehensive law governing AI in a jurisdiction of this scale. It entered into force on 1 August 2024.

The Act's foundations lie in classifying AI systems according to their level of risk, rather than regulating 'AI' as a single category. There are four categories of risk:

Prohibited Risk

This category is designed to root out potential abuses like social scoring or excessive surveillance.

High Risk

This category imposes heavy compliance obligations on uses of AI that have anything to do with safety, critical services, and potentially controversial issues of law enforcement and security.

Limited Risk

These systems are subject primarily to transparency obligations. This includes everyday AI use already popular in areas like large language models (LLMs) and chatbots.

Minimal Risk

Systems in this category are not subject to mandatory obligations under the Act. Examples include basic AI productivity tools, games, and other uncontentious uses.

Alongside this risk-based framework is a separate set of rules for providers of general-purpose AI models, with additional obligations for models designated as posing systemic risk.

The obligations arrived in stages:

2025

2 FEBRUARY 2025

Prohibitions and AI literacy duties began to apply on 2 February 2025.

2 FEBRUARY 2025

General-purpose AI model obligations began to apply on 2 August 2025.

2026

AUGUST 2026

Transparency obligations generally began to apply in August 2026.

2 DECEMBER 2026

Key provisions on watermarking duties to label AI-created material will not apply

2027 - 28

2 DECEMBER 2027

Implementation of the high-risk regime was postponed because the necessary enforcement architecture, including national competent authorities and harmonised technical standards, was not ready. Following the EU's AI Omnibus, high-risk obligations for standalone AI systems will apply from 2 December 2027.

2 AUGUST 2028

The rules for high-risk AI systems embedded in regulated products will apply from 2 August 2028.

A deployer that fine-tunes, rebrands, or repurposes a third-party system can be reclassified as its provider, inheriting a far heavier set of obligations.

Enforcement is shared between the Commission's AI Office and national competent authorities. Penalties can reach €35 million or 7% of global annual turnover for prohibited practices, and €15 million or 3% of global annual turnover for most other infringements.

There are four obvious similarities between the GDPR and the EU AI Act:

- ✓ Both have extraterritorial reach. Their requirements can apply outside the EU when those organisations place systems or services on the EU market or when the outputs of those systems are used in the EU.
- ✓ Both are 'horizontal' laws. In other words, they apply across all sectors rather than one specific industry.
- ✓ Although both laws are characterised by stringent, turnover-based fines, compliance is about painstaking internal work.
- ✓ Regulatory enforcement is the responsibility of member states through the by now well-established architecture of competent authorities operating within EU-wide standards set centrally. A similar model applies to the Network Information Systems Directives, NIS 1 and NIS 2, which address the cyber security and resilience of critical infrastructure.

However, there are also crucial conceptual and practical differences:

01.

The extent of an organisation's regulatory obligations varies *far more* under the EU AI Act than under the GDPR.

The GDPR requires organisations to build a compliance architecture for processing personal data. The EU AI Act requires triaging business operations to work out which AI systems and activities are covered. That assessment determines what compliance architecture needs to be built. Most companies will find that most of their actions are unregulated, while some will find themselves with heavy compliance duties. The outcome depends on the organisation's business model. This differs from the GDPR because virtually all organisations had stringent obligations under that regulation.

02.

Compliance is a far more *dynamic process* under the EU AI Act than it is under the GDPR.

Given the fast-changing nature of AI and its uses, complying with the Act will be a more dynamic process for most organisations than compliance with the GDPR. The concept of personal data doesn't change much. Under the AI Act, however, organisations will need continuously to evaluate which systems they use and what those changes mean for compliance with the Act.

03.

Designations and supply chains are *much bigger issues* under the EU AI Act than under GDPR.

Under the GDPR, an organisation is a controller or processor, and that distinction is usually stable. Under the AI Act, most companies are deployers that buy and use AI systems. However, a deployer that fine-tunes, rebrands, or repurposes a third-party system can be reclassified as its provider, inheriting a far heavier set of obligations. Similarly, the GDPR's controller-processor relationship is essentially bilateral. Under the EU AI Act, responsibilities may extend across a chain involving a general-purpose model provider, an AI system provider or integrator, and a deployer. A deployer's ability to comply may depend on documentation and compliance work done upstream by parties it does not control.

04.

Key parts of the legal framework are *fundamentally different*.

The GDPR is primarily an administrative compliance regime built around controls for protecting personal data. The EU AI Act more closely resembles product safety law, drawing on concepts used to regulate machinery and medical devices. All organisations are used to administrative legal obligations like the GDPR. However, many organisations subject to the EU AI Act, including some financial institutions, may be less familiar with the product safety model. Responsibility under the GDPR is structured principally around controllers and processors, with a Data Protection Officer required in specified circumstances. The EU AI Act distributes responsibilities across several roles, including providers, deployers, importers, and distributors. This makes it more difficult to identify a single organisational home for compliance.



Part 3: The Key Questions for Businesses

These similarities and differences offer important lessons from the GDPR experience for organisations inside and outside the EU as they prepare for the EU AI Act.

The geopolitical environment, which is not the subject of this paper, is different from 2018. Questions remain about whether the Brussels Effect is as strong as it once was. That said, it seems clear that many organisations selling into the EU single market realise they must comply with the Act.

To that extent, we are experiencing something similar to the GDPR moment. However, the two moments are not the same.

SANS Institute is committed to helping organisations manage these changes as effectively as possible and supporting the safe and secure adoption of new AI capabilities. We are therefore publishing a series of two papers that seek to:

- Explain the changes arising from the new legislation and identify lessons from the GDPR experience.
- Consult on organisations' experiences of preparing for those changes.
- Use those responses to publish further free guidance in early 2027.

The survey accompanying this paper is designed to examine how companies are responding in practice. It is based on the questions below, which we believe are particularly important to organisations and to the strategic direction of EU AI implementation. We present them here to encourage respondents to provide additional comments in the survey.

The questions fall into three categories:

01.

Readiness

02.

Organisational Implications

03.

Strategic Implications

The work that compliance requires, including understanding which systems an organisation operates, why it operates them, and who is accountable for them, is worth doing in its own right.

01. Readiness

Under the general heading of readiness for the EU AI Act, we are particularly interested in the following questions:

- Are organisations ready and able to compile an inventory of the AI systems they use, including systems embedded in software they already license and capabilities introduced through vendor updates?
- Can organisations classify those systems by the risk tiers established by the new law, and have they done so?
- Do organisations understand when fine-tuning AI systems could cause them to be reclassified from a deployer to a provider?
- How have organisations implemented the new transparency obligations?

02. Organisational Implications

To understand what the Act means for organisations and how important the change is, we are interested in the following questions:

- To what extent do existing EU laws, including the GDPR, the Cyber Resilience Act, and the NIS Directives, provide governance structures and expertise that can support implementation of the EU AI Act?
- Where does responsibility for EU AI Act compliance sit within organisations, given that the Act does not specify an organisational home equivalent to the GDPR's data protection officer role?
- How much is implementation costing, and how do those costs compare with the cost of implementing the GDPR?

03. Strategic Implications

No organisation exists in a vacuum. For better or worse, the EU AI Act will influence how AI is implemented and developed, both within the EU and beyond it. We want to understand what organisations think about this strategic direction. In particular:

- Are organisations applying the EU AI Act globally, as many did with the GDPR, or are they partitioning and restricting AI capabilities by region? In other words, is the Brussels Effect still as effective?
- Do organisations directly affected by the Act believe that it will ultimately produce positive strategic outcomes?

We want to hear how your organisation is approaching these questions. The survey accompanying this paper will remain open throughout the autumn, and its findings will provide the foundation for our 2027 white paper on EU technology and regulatory trends.

Even at this stage, we can offer one conclusion drawn from the lessons of the GDPR: the work that compliance requires, including understanding which systems an organisation operates, why it operates them, and who is accountable for them, is worth doing in its own right.