Atlantic Council
CYBER STATECRAFT INITIATIVE
ai x ac
April 2026

Issue brief Navigating the European Union's AI and health data framework

Written by Mark Scott

The EU's overlapping regulatory frameworks — the GDPR, AI Act, and European Health Data Space Regulation — are reshaping how companies can access and use health data for AI development. This policy brief outlines the growing tension between the EU's rights-based, data-localization approach and the US preference for open, globally pooled datasets, while offering practical recommendations for pharmaceutical and technology companies navigating this shifting landscape.

Executive summary

- The European Union's regulatory structures of the General Data Protection Regulation (GDPR), Artificial Intelligence Act, and European Health Data Space Regulation represent a step change in the bloc's approach to AI governance, cross-border data flows, and the use of health data. Collectively, these rules impose significant governance obligations on AI systems—from training data provenance and AI model validation to post-deployment oversight—which mirror existing regulations for the pharmaceutical and medical device industries.
- The ability to transfer European data outside the EU for pooled global datasets is now a legal liability. The EU is shifting toward a federated data infrastructure model, which requires high-risk health data to be either segmented from non-EU data or, increasingly, stored locally within secure infrastructure environments. This change, which will come into force over the next five years, will reshape pharmaceutical research pipelines and create friction between diverging European and American approaches to data access.
- The EU's evolving governance structures will require companies to adopt "regulation-by-design" principles, which may offer a competitive advantage for well-resourced pharmaceutical companies. Firms which develop research pipelines that comply with EU rules may be able to access EU health data more reliably, while also engaging with regulators more proactively in ways that can reduce legal uncertainty compared to less well-resourced or smaller competitors.

Introduction

The EU is at an inflection point when it comes to data usage, digital regulation and its embrace of AI-powered economic growth.

There are two primary drivers for this change.

Ursula von der Leyen's return in 2024 for a second five-year term as president of the European Commission coincided with the rise of her center-right European People's Party as the leader of digital and economic policymaking within the 27-country bloc. These politicians have prioritized pa-

red-back regulation and an emphasis on industrial development over the singular promotion of EU fundamental rights.

This dynamic coincides with the publication of an EU competitiveness report in late 2024 by Mario Draghi, the former Italian prime minister. The analysis highlighted, in part, the perception of overly complex and burdensome regulatory practices as a cause of the bloc's reduced economic competitiveness compared to other parts of the world, notably the United States and China.

Taken together, these dual policymaking forces have led to a fundamental reassessment of the EU's industrial strategy, which has created both opportunities and friction for industries seeking to access data within the bloc.

On the one hand, the EU's approach to digital regulation, underpinned by its GDPR, Artificial Intelligence Act, and European Health Data Space Regulation, has not changed. Together, this legislative package defines AI use cases within the health sector as a high-risk, systemically-significant activity that require additional regulation. That includes mandatory safeguards for data protection, transparency and accountability in how the bloc's data may be utilized by AI-powered pharmaceutical development.

On the other hand, senior European leaders, both within the European Commission and at the member country level, have encouraged industries operating within the EU to make better use of the Continent's treasure trove of personal and non-personal data for AI-related economic growth. The European Commission published its so-called Digital Omnibus Regulation Proposal in late 2025, which aims to reduce regulatory burdens associated with the EU's data protection and AI legislation to enable economic development, as outlined in the Draghi Report. These proposals must still be enacted by both the European Parliament and Council of the European Union, and negotiations are likely to drag on into late summer 2026.

Adding to the current complexity, geopolitical tensions—particularly the transatlantic relationship with the United States—have increased scrutiny these industries operating on both sides of the Atlantic.

Doubts over the sustainability of future EU-US data flows, as well as the rise of digital sovereignty as a policymaking priority within the 27-country bloc, make it difficult to navigate around artificial intelligence and health data. The US and EU approaches to these topics is fundamentally different, and it is unclear whether, in the short-term, there is a path forward to reconcile these diverging regulatory and policymaking landscapes.

Roots and drivers of the European Union's cross-border data policy

Laws, regulations, and policies

The EU's cross-border data policy is based on the notion that the protection of personal data is a fundamental right, as outlined in Article 7 and 8 of the EU Charter of Fundamental Rights. When assessing third-party jurisdictions' ability to access such data—including so-called special category data like for health—the bloc must determine if such cross-border data transfers are lawful through "essential equivalence" protection.

In practice, this means third-party countries must guarantee that EU citizens' data will be treated in their jurisdictions with substantially the same level of protection as outlined in the EU Charter of Fundamental Rights and its protection rules.

The foundation of the EU's data privacy regime is the GDPR. These rules impose strict conditions on the processing and international transfer of data. It has made the EU's data protection rulebook the *de facto* global standard for online privacy rules.

Within this regulatory regime, there are additional safeguards for sensitive data, including information about individuals genetic, biometric, and other health data, as outlined in Article 9 of GDPR. Such data is believed to be intrinsically higher-risk, and therefore additional oversight is required to reduce the likelihood of harm and reflect the fundamental rights protection within EU law.

The bloc's AI Act extends this protection approach to AI systems—both in terms of the importance of special category data and requirements for third-party jurisdictions to have EU-equivalent protections.

Under the legislation, which will progressively come into force by late 2027, specific AI use can been deemed to be "high-risk," as outlined in Article 6 of the Act, if it's believed to significantly impact EU citizens' fundamental rights, safety, or health. That definition means that almost all health- and pharmaceutical-related applications, as defined within the Act, fall within the high-risk category and, therefore, require regulation for not only specific AI outputs but also the training data quality, provenance, and representation, as well as requirements for bias mitigation and documentation that underpin such high-risk applications.

The European Health Data Space Regulation (EHDS) completes the EU's regulatory approach by institutionalizing how national health data is controlled and accessed. It requires all secondary use (by industry or research groups which do not own the data) to take place through public or supervised infrastructure. This builds on both the GDPR and the AI Act by basing all such data access on the promotion of privacy as a

fundamental right and the need to demonstrate how AI systems uphold such values via accountability and data minimization procedures.

In practice, the EHDS shifts the balance from the transfer of data, often outside of the bloc, to the access of such data via secure processing environments within the EU. The regulation prioritizes secure data processing environments where AI systems can demonstrate their ability to comply with both the GDPR and AI Act before they can access special category data.

Political system

The trifecta of the GDPR, AI Act, and EHDS is shaped by the EU's history and regulatory traditions, namely the friction between fundamental rights protection and member state market integration.

Within the data protection policymaking space, EU courts have shown a significant willingness to upend long-standing cross-border data flow agreements with third-party countries (primarily the United States) over concerns that such jurisdictions do not uphold equivalent protections as outlined within the EU Charter of Fundamental Rights.

The most notable cases are colloquially referred to as Schrems 1 and Schrems 2, in which an Austrian privacy campaigner successfully overturned transatlantic agreements on cross-border data flows. Judges from the EU's highest court agreed that the US did not offer so-called "essentially equivalent" data protection safeguards to EU citizens when their data was transferred to the US.

In response, the US offered new redress mechanisms to EU citizens who believed their information had been mishandled by the federal government. That agreement, known as the EU-US Privacy Framework, is now being tested within EU courts over renewed claims of equal protection.

Unlike in the US, where sectoral regulators have broad discretion in how they enforce their mandate, their EU counterparts operate primarily under judicially-enforceable duties. In practice, this means that supervisory authorities are legally obligated to suspend cross-border data transfers if adequate protection cannot be ensured.

These baked-in legal obligations lead to a precautionary, *ex ante* enforcement approach where supervisory authorities often taken a conservative view when implementing specific legislation. This is particularly true when dealing with special category data, like that associated with the health sector. This regulatory regime, combined with the bloc's focus on the protection of fundamental rights, results in strict government oversight for health and pharmaceutical use cases.

Evolution of the tech industry

Despite its large consumer market and economic dynamism, the EU does not have a world-class technology industry compared to those of the US and China. With no Silicon Valley-style tech giants based in the continent, EU economic policymakers—until recently—have focused on tech for industrial applications, and regulated sectors like health and energy as part of wider industrial strategy priorities.

The lack of a well-functioning digital single market across EU—where, for example, a Portuguese startup could raise funds from a Polish venture capitalist to produce digital services for Swedish consumers—is a primary driver for its lack of consumer-tech companies. Ongoing regulatory bottlenecks, including in nation-specific labor markets and domestic-focused capital markets, have only added to the EU's disconnect with the US and China—arguably the world's two most important tech industries.

The absence of a well-functioning EU consumer technology sector has had two major policy implications.

First, large international firms have greater capacity to comply with EU's digital regulations compared to smaller European competitors. That has led, particularly in relation to the GDPR and AI Act, to greater compliance by non-EU companies, many of which have used this fact to maintain their market position within the bloc.

Second, EU digital oversight policy has focused on making technology compatible with existing regulations in high-performing European sectors. This is evident in the AI Act, whose risk-based approach, *ex ante* controls, and clear responsibility and accountability requirements more resemble existing pharmaceutical and medical device regulation than those for the software industry.

The role of AI

EU policymakers have sought to create governance structures around artificial intelligence that meet existing societal and political norms. It is notable that the bloc was the first democratic jurisdiction to pass comprehensive AI legislation, hoping to use the AI Act to shape emerging global policy—an example of the so-called "Brussels Effect."

In general, EU officials view AI systems as potential risk amplifiers, and not merely neutral tools. Those risks relate to both the undermining of EU fundamental rights and also include scaling bias (when an AI model's accuracy drops at increasing scales of complexity), errors, and discrimination via opaque models that are not accountable to the public.

For data, there is a similar concern that AI systems may undermine EU fundamental rights if sensitive information, collected

across entire populations, is transferred outside of the bloc in ways that does not uphold the privacy equivalence.

In practice, the EU's AI governance is based on the belief that any potential risks must be addressed before deployment, especially for high-risk uses cases like those involving health data. Unlike other jurisdictions, where AI policy is treated as a relatively new and separate domain, the EU's approach embeds AI into other existing policy areas and regulatory regimes, like data protection, product safety, and market competition.

The state of health data in the AI sector

Health data has increasingly become a geopolitical and economic asset. Within this context, the US and EU have embraced opposing models in how to leverage such society-wide assets.

The US promotes a relatively open strategy towards data use, though it has recently imposed how such data is exported to "foreign adversaries," such as China and Russia. The country relies on these cross-border data flows to power AI growth. The EU, in contrast, is moving toward data localization practices and intra-European AI development where training algorithms are tested and utilized within closed infrastructure.

Where the US favors the global pooling of health data, which enables significant cost efficiencies and faster iteration for innovations, the EU prefers a more targeted approach with federated, localized infrastructure.

The result has been the legal fragmentation of AI-enabled use cases for health data between the democratic world's two most important players. Ongoing concerns from the EU around cross-border data flows to the US, where there is legal uncertainty over domestic surveillance laws, has only exacerbated this divide.

Current successes: cross-border data collection, use, and transfer

While the ability for companies to share health data outside the bloc has become more difficult, there are also extensive cross-border data modalities within the EU fostered by the GDPR and, increasingly, the EHDS. The EU's data protection framework has created a shared baseline across all 27 member countries, equating to a population of 450 million. This harmonization enables data pooling, at scale, with a stable regulatory stable environment that is unavailable in most other parts of the world.

The EHDS builds on this regulatory alignment by permitting entities that meet specific data protection and security requirements to request access to national health data. Use cases permitted within the regulation include scientific research associated with AI development.

The legislation reinvents how health data is typically handled, based on the EU's collective promotion of fundamental rights and its *ex-ante* regulatory oversight. Researchers do not receive raw global datasets. Instead, data is accessed through secure processing environments and external AI systems are used to maintain adequate data protection controls.

To be EU compliant, companies must invest in federated learning infrastructure across national registries provided via EHDS. This requires technical expertise to leverage compliant secure processing environments, which can be a barrier to entry for smaller firms and a potential competitive advantage for larger rivals with existing compliance capacity.

Given the size of the EU health data market, the combination of the GDPR, AI Act, and EHDS offers two mutually reinforcing benefits. First, it allows companies access to a relatively genetically diverse and extensive health data market with built-in data protection and security protocols. Second, it theoretically incentivizes engagement with the European Medicines Agency and aligned national Health Technology Assessment Bodies, which are actively seeking to enable AI-driven research and innovation from compliant researchers and companies.

As of early 2026, however, a few examples of such engagement between the EU and such companies exist. According to the DARWIN EU Coordination Center, which facilitates access to health data across the bloc, most health data pooling has been used for academic endeavors.

Current Challenges: Cross-Border Data Collection, Use, and Transfer

The EU's broader regulatory regime and oversight of health data and AI compliance structures stand at odds with global pooled datasets, which remain the cornerstone of current AI foundation models. Mixing EU and non-EU health data will inevitably breach one, if not all three, of the GDPR, AI Act, and EHDS legislative packages.

These governance and regulatory mechanisms —and their disconnect with traditional global pooled data models favored by othe jurisdictions—will increase costs and drive compliance complexity. Federated learning models, in which data is held in local secure processing environments, are slower and more intensive compared to global datasets. For most companies, including those with global teams and extensive resources, the further development of parallel EU and US AI application pipelines will be cost-prohibitive.

From a corporate culture perspective, the EU's regulatory structures will also require significant changes within individual companies. So-called "data versioning," or the systemic tracking, management, and identification of changes to datasets over time—now obligatory within the EU—requires signifi-

cant capacity-building for teams accustomed to rapid innovation timelines.

Similarly, the AI Act's audit demands, in which companies and researches with models using high-risk datasets must provide information on their development and bias mitigation approaches and demonstrate how their data is representative, will pose a challenge to many smaller businesses and research initiatives.

Future growth areas and tensions

The EU will continue to push for geographically segmented AI models, particularly around high-risk and sensitive datasets. This includes EU-specific training environments, limited datasets, and modular innovation pipelines that separate AI training from deployment.

As the bloc's EHDS develops, the level of sophistication and breadth of federated and privacy-enhancing infrastructure will mature and become more sophisticated. This may eventually open significant amounts of health data for intra-EU cross-border collaboration.

Yet the EU does not operate in a vacuum. The bloc's regulatory approach stands apart from more permissive jurisdictions, including those available within the US. Without significant investment, US-based pharmaceutical AI models and other health-related technologies will fail EU regulatory tests barring significant redesign.

Future policy directions

Near-term (2026 – 2027)

The EU's AI Act is still a work in progress, and compliance for high-risk AI in regulated products has been given an extended transition until August 2027. Most other obligations under the regulation, however, come into force in August 2026. To support industries in their regulatory planning, the European Commission published its voluntary General-Purpose AI Code of Practice in 2025.

Collectively, this requires companies, including those using health data, to build sufficient internal governance structures to meet the upcoming regulatory deadlines. That requires greater documentation, traceability, and oversight of AI models before deployment. For health data in particular, the European Medicines Agency and Federal Drug Administration published joint guidelines in January 2026 that further outlined principles for "good AI practice in drug development."

International cross-border data transfers remain reliant on the European Commission's adequacy decisions within individual third-party countries. This includes the current EU-US Data Privacy Framework, although its legality has yet to be fully addressed within EU courts.

This country-by-country determination does not bar all transfers, as companies can still share data outside of the EU via standard contractual clauses and binding corporate rules. However, for health data, the EHDS includes a "permission layer" that sits above GDPR which requires data exports to be authorized, analyzed in secure processing environments, and be purpose limited.

The result of these standards is that EU health data to the US is currently permissible, though has moved toward a risk-managed flow model, including the segmentation of EU data from non-EU data and the addition of security protocols to meet EU regulations.

Medium-term (2027-2031)

The implementation of the AI Act's requirements for high-risk AI in regulated products in August 2027 will be a step change in regulatory compliance.

EU officials are positioning this upcoming deadline as complementary to existing medical device rules, including the Medical Device Regulation and In Vitro Diagnostic Medical Device Regulation. For companies, this will require an integrated compliance approach that for all regulatory functions.

Cross-border data transfers— particularly sensitive data— are likely to become more difficult. Ongoing transfer will require significant auditability and be segmented between EU and non-EU data, which will drive up costs. Companies' reliance on standard contractual clauses, or pre-approved legal provisions to facilitate cross-border data exchange, will become more difficult as EU data protection regulators exert greater control of their high-risk datasets.

The EHDS will similarly evolve from an untested policy instrument to the cross-border infrastructure that facilitates the secondary use of national health data across the bloc. Policymakers hope that this will reshape the pharmaceutical industry's research approach by nudging it toward EU-based analysis within secure digital environments and promoting federated research models compliant with EHDS governance structures.

It is unclear, however, if such policy priorities align with industry objectives.

Long-term (2031 – 2035)

One of the AI Act's stated objectives is to treat high-risk AI models as a regulated entity akin to current standard practices in both the pharmaceutical and medical device industries. Documented and auditable inputs, demonstratable pre-deployment risk- and harm-mitigation, and structured compliance management in alignment with EU standards are expected to become the norm.

The long-term sustainability of the EU's adequacy agreement structures, in which currently a few predominantly Western

countries are deemed to have equivalent data protection standards to those within the bloc, is uncertain. A shift toward multistakeholder approaches where nation states negotiate cross-border data agreements based on mutual assurance is gaining traction with non-European countries, but these agreements have not been supported by the EU.

Access to global data sets will remain, although there will be increased pressure to legally segment EU and non-EU datasets to ensure the bloc's data protection standards are upheld. In this world of siloed datasets, interoperability may continue via the deployment of AI models within secure data environments, which is a marked shifted away from the common industry standard of global pooled datasets.

The future of EHDS is also unknown. Currently, this regulation is more expected standards than defined infrastructure, with significant variation in national level implementation plans. EU officials hope this infrastructure—and the chance for compliant companies to be granted secondary access to such large pools of health data—will drive pharmaceutical innovation within Europe.

This new approach assumes that a combination of competitive advantages of scaled federated and privacy-preserving infrastructure, the maintenance of sought-after long-term datasets, and closer alignment between regulators and researches will spark an economic engine. Such results have yet to materialize, however, as this new regime will take three to five years to operationalize.

Policy recommendations

The EU's evolving approach to cross-border data flows and AI models connected to health data are at an inflection point.

The regulatory trifecta of the GDPR, AI Act, and EHDS will drive up costs, create internal and external regulatory challenges, and further divide industry between opposing EU and US governance approaches.

Given these complexities, below are three policy recommendations for how companies operating in the EU can navigate the current environment:

1. Adopt "regulation by design" principles

Instead of retrofitting compliance onto technologies and workflows, companies should instead design their AI systems to meet EU regulatory requirements from day one. They should treat these untested governance models as a core feature of product design, not a *post-hoc* legal obligation.

Under the EU's AI Act, health data is considered high-risk and requires significant regulation for its training, governance, and post-deployment management. Applying compliance after-the-fact risks insecure data provenance and AI models that cannot be fully audited, and ongoing iteration of these systems breaches "regulation by design" principles. The result will be compliance failure.

By redesigning production workflows based on existing or future EU requirements, companies can better meet criteria while also reducing potential regulatory roadblocks that may stall the deployment or force wholesale redesign.

2. Shift toward federated data access

The ongoing use of pooled global datasets at scale is no longer compliant with EU regulation. To maintain access to European health data, companies should invest in federated and jurisdictional-specific data infrastructure that complies with GDPR and EHDS.

In practice, that necessitates distributed analytics and localized datasets for EU data that maintains European legal control over sensitive and high-risk health datasets. This means that AI models are ported into such secure environments, rather than EU data being externally transferred within pooled global datasets.

The potential benefits of such an approach are threefold: 1) it allows companies to access EU health data at scale; 2) it reduces the regulatory risk associated with continued cross-border data flows; and 3) it allows for corporate innovation pipelines with future federated infrastructure associated with the EHDS.

3. Reframe regulatory compliance as a strategic asset

Given the similarities between the EU's regulatory model toward AI and existing governance structures for the pharmaceutical and medical device industry, companies should position themselves as enablers of EU AI governance—and not merely rule-takers which have limited engagement with policymakers.

The EU's revamped interest in AI-driven economic growth opens possibilities for companies and researchers which can navigate future regulatory structures and strategically access EU health data.

The specifics of many of these rules have not yet been decided. This places those who proactively engage with authorities at an advantage over those who view these regulations as a mere compliance cost. This type of relationship with the European Medicines Agency, national data protection authorities, and EHDS health data access bodies can position the health and pharmaceutical industries as enablers of innovation-friendly policy structures.

Such policymaking engagement would allow companies to meet their regulatory requirements while also supporting agencies and policymakers in the development of pragmatic and flexible regulatory oversight and governance.

About the author

Mark Scott is a senior resident fellow at the Digital Forensic Research Lab's (DFRLab) Democracy + Tech Initiative within the Atlantic Council Technology Programs. In this role, he is engaged in expanding the Initiative's ongoing work around comparative digital policy, regulation, and governance, as well as efforts linked to the European Union's Digital Services Act and Digital Markets Act. He currently sits on the international advisory board of RegulAite, a project at the University of Amsterdam dedicated to artificial intelligence policymaking. He is also a research fellow at the Centre for Digital Governance at the Hertie School in Berlin.

Acknowledgments

The author would like to thank Ken Propp, Celine Lee, and Kenton Thibaut for their comments on earlier drafts of this report, Nitansha Bansal for critical help in getting the report to final form, as well as all the individuals who participated in background and Chatham House Rule discussions about EU's regulatory approach AI and health data.

About the center

The **Cyber Statecraft Initiative** works at the nexus of geopolitics, technology, and security to craft strategies to help shape the conduct of statecraft and to better inform and secure users. This work extends through the competition of state and non-state actors, the security of the internet and computing systems, the safety of operational technology and physical systems, and the communities of cyberspace. The Initiative convenes a diverse network of passionate and knowledgeable contributors, bridging the gap among technical, policy, and user communities.

The Atlantic Council is a nonpartisan organization that promotes constructive US leadership and engagement in international affairs based on the central role of the Atlantic community in meeting today's global challenges.

Atlantic Council
1400 L Street NW, 11th Floor
Washington, DC 20005
(202) 778-4952
www.AtlanticCouncil.org