Europe's Health Data Shift: Regulation, Anonymisation, and Security

T. Chib, R. van Kempen and A.I. Hakkers[1]

The 2021 ransomware attack on Ireland's Health Service Executive[2], where attackers threatened to publish patient data, presaged a new era of healthcare vulnerability. As Europe implements ambitious data-sharing frameworks in 2025, this incident reminds us of the central challenge facing modern healthcare: how do we make data useful without making it dangerous?

To understand this fundamental tension, we examine it through three interconnected lenses that together form what we call the '*Privacy-Security-Utility Triangle*', or the '*Golden Triangle*':

The regulatory analysis reveals how Europe's new frameworks attempt to legislate the balance between usefulness and danger. In this article, we analyse where these regulations create genuine progress, such as patient control mechanisms and cross-border portability and where they multiply complexity without resolving core tensions. The convergence of four major frameworks over twenty-four months forces healthcare organisations to navigate conflicting requirements, and where making data useful for research may make it dangerous for privacy.

The anonymisation analysis examines why the regulatory promise of "anonymous" health data sharing fails against technical reality. We suggest how medical data's inherent uniqueness makes true anonymisation impossible while retaining utility. Every technique that preserves usefulness leaves fingerprints that enable re-identification, while every method that ensures anonymity destroys the very characteristics that make data valuable for research.

The cybersecurity analysis shows how every step toward better, faster care also creates new risks. As hospitals connect more systems and share more data, they become easier targets. Old machines, rushed innovation, and complex networks open the door to attacks that can shut down care or leak sensitive data. In healthcare's digital evolution, the question isn't whether we can make data useful without increasing risks. It's how dynamically we can balance creating systems resilient enough to bend without breaking, open enough to enable care without enabling attacks, and sophisticated enough to know when each matters most, and continue doing so as technology, threats, and care needs keep evolving.

These three perspectives converge to reveal that healthcare's digital transformation is not a technical challenge with a definitive solution, but rather a perpetual balancing act between competing imperatives that cannot be permanently resolved. The regulatory frameworks promise control through compliance, anonymisation promises safety through transformation, and cybersecurity promises protection through barriers, yet each solution creates new vulnerabilities even as it addresses existing ones.

1. Privacy in Healthcare: how regulations are reshaping rights and responsibilities

The European health data landscape is undergoing its most significant transformation since GDPR, driven by a deceptively simple question: how do we make data useful without making it dangerous? Four interconnected frameworks each attempt their own answer, yet their interaction reveals that every mechanism designed to unlock data's utility, such as standardised access, transparency, and interoperability, also simultaneously amplify risks.

1.1. European Health Data Space ("EHDS")

The European Health Data Space[3], which entered into force in January 2025 after three years of debate, represents Europe's most ambitious attempt to create a unified health data ecosystem. Its dual infrastructure: MyHealth@EU for patient data portability and HealthData@EU for research, promises to revolutionize both individual empowerment and collective medical advancement. Yet the decade-long implementation timeline stretching to 2035 itself acknowledges the enormity of harmonising 27 national health systems while introducing unprecedented patient controls.

1. Tanya Chib is a regulatory and privacy lawyer, Dr. Anna Hakkers is a cybersecurity expert specializing in Data Security and Renate van Kempen is a data anonymisation consultant focused on health data protection and re-identification risk assessment.
2. https://www2.hse.ie/services/cyber-attack/what-happened/
3. https://health.ec.europa.eu/ehealth-digital-health-and-care/european-health-data-space-regulation-ehds_en

EHDS Implementation Timeline
2027: Member States designate Digital Health Authorities and Health Data Access Bodies.
2029: Basic health data (patient summaries, e-prescriptions) available cross-border via MyHealth@EU
2031: Expanded to include medical imaging, lab results, and discharge summaries
2033: European Commission evaluation and potential adjustments
2035: Full implementation including HealthData@EU infrastructure for secondary use

The EHDS does introduce genuinely transformative patient rights that, if properly implemented, could fundamentally alter the power dynamic in healthcare. Patients gain control through opt-out rights for secondary data use without needing to justify their decision. However, the exceptions for public interest research create ambiguity about when individual preferences can be overridden[4]. Nonetheless, EDHS enables granular access control, allowing patients to shield sensitive information about mental health or sexual health from certain providers while sharing other medical data. Real-time notifications promise to alert patients whenever their data is accessed, creating an audit trail of who sees what and why. Patients can demand immediate error corrections at no cost, addressing long-standing frustrations with inaccurate medical records that follow individuals across providers.

Yet these empowering provisions contain subtle contradictions that may undermine their effectiveness. The emergency override provision, while necessary for life-threatening situations, lacks clear boundaries on what constitutes sufficient emergency to bypass patient restrictions and remains undefined, potentially creating a loophole that normalises circumventing preferences. The technical complexity that makes data useful for granular control, also makes it dangerously inaccessible for vulnerable populations who need it most. Further, the framework's data sharing mechanisms could enable large technology companies to access aggregated health datasets for AI development, potentially prioritising commercial innovation over individual privacy despite regulatory safeguards.

The digital divide across Europe further threatens to transform EHDS from an equalising force into another source of healthcare inequality[5]. While Denmark's advanced digital infrastructure can readily support real-time notifications when patient data is accessed, healthcare facilities in less advanced countries may struggle. Vulnerable populations such as the elderly, disabled, and migrant communities will require extensive support to navigate granular controls, further adding costs that cash-strapped systems cannot afford. Without addressing these disparities, EHDS risks creating a two-tier system where digital literacy determines healthcare quality. The 2027 milestone will reveal whether Member States pursue genuine transformation or mere compliance. If stakeholders embrace both the letter and spirit of the regulation, Europe could pioneer a model balancing individual autonomy with collective benefit. If not, the EHDS risks becoming another hollow framework - well-intentioned but ultimately ineffective.

As exemplified above, the EDHS embodies the 'useful-dangerous' paradox at many levels. For example, while standardising access allows a Spanish tourist to get his prescriptions in Sweden, this does create new attack vectors. The framework's patient control features (opt-out rights, granular access, real-time notifications) make data useful for individual empowerment but dangerous through complexity, in cases of vulnerable populations where elderly patients might accidentally restrict critical information from emergency providers.

1.2. The EU-US Data Privacy Framework ("DPF")[6]

The implementation challenges highlighted above multiply when EHDS intersects with the EU-US Data Privacy Framework, established in July 2023 as the third attempt to enable transatlantic data flows. If EHDS shows how making data useful for patients makes it dangerous for privacy, the DPF reveals how making it useful for research makes it dangerous for sovereignty.

The DPF aims to establish "adequacy" for personal data transfers under GDPR, creating a mechanism for legal health data sharing between European and American institutions and organisations. The DPF's predecessor, the Safe Harbor and Privacy Shield, failed due to fundamental conflicts between EU privacy rights and US surveillance practices[7], so it remains to be seen if this third attempt can survive legal challenges.

Complexity arises because the DPF must coexist with multiple regulatory frameworks. Healthcare

4. Marelli L, Stevens M, Sharon T, Van Hoyweghen I, Boeckhout M, Colussi I, Degelsegger-Márquez A, El-Sayed S, Hoeyer K, van Kessel R, Zając DK, Matei M, Roda S, Prainsack B, Schlünder I, Shabani M, Southerington T. The European health data space: Too big to succeed? Health Policy. 2023 Sep;135:104861. doi: 10.1016/j.healthpol.2023.104861. Epub 2023 Jun 26. PMID: 37399677; PMCID: PMC10448378
5. https://www.eu-patient.eu/globalassets/ehds-analysis---final.pdf
6. https://www.dataprivacyframework.gov/Program-Overview
7. https://curia.europa.eu/jcms/upload/docs/application/pdf/2020-07/cp200091en.pdf

organisations transferring data between the EU and US must satisfy not only DPF principles but also Health Insurance Portability and Accountability Act ("HIPAA")[8] requirements, while navigating the patchwork of country-specific health data rules that EU Member States impose beyond GDPR[9]. Without clear guidance on how these frameworks interact, organisations are forced to layer Standard Contractual Clauses atop DPF certification, essentially creating their own reconciliation between competing regulatory demands. The absence of healthcare-specific provisions means medical data transfers operate in a regulatory grey zone where generic privacy principles inadequately address sector-specific realities.

The DPF represents progress but falls short of providing comprehensive solutions for healthcare data transfers. Its generic framework, combined with Member States' retained flexibility to impose additional health data restrictions, forces organisations into complex multi-layered compliance strategies. Healthcare institutions engaged in transatlantic collaboration must combine DPF certification with tailored contractual arrangements, conduct exhaustive jurisdictional analyses for each EU Member State involved, and essentially create their own sector-specific protections within the generic framework. This patchwork approach reflects a deeper failure: the absence of healthcare-specific provisions that acknowledge how medical data differs from commercial information in sensitivity, use patterns, and ethical considerations. Until regulators develop frameworks that genuinely address health data's unique nature, rather than treating it as merely another data category with higher risks, organisations must compensate through costly and inefficient workarounds that may still face legal challenges.

1.3. The European Data Act ("Data Act")

The Data Act[10], entering into force in January 2024 with general applicability from September 2025, introduces another layer of complexity that fundamentally challenges medical device innovation. While promising to create a thriving data sharing economy by granting users rights to access and share data from their pacemakers to fitness trackers, the regulation's "access by design" mandate demands architectural changes that many manufacturers may not have anticipated.

The Data Act's nuanced approach to access rights reveals tensions between promise and practice. Article 3.1 establishes that data must be "by default, easily, securely and directly accessible", yet manufacturers retain control over initial contract terms, creating immediate ambiguity about what constitutes genuine access. Users gain rights only to data generated by their use, not all data processed by devices, a distinction that becomes critical for AI-enabled medical devices that process far more information than they store.

Article 4's requirement for "continuous and real-time" data access poses particular challenges for implanted devices designed with closed architectures for safety reasons. Medical device manufacturers must somehow fit user-accessible data ports or wireless interfaces without compromising therapeutic functions, all while maintaining compliance with the EU Medical Device Regulation ("MDR") and In Vitro Diagnostic Regulation ("IVDR").

The framework's third-party sharing provisions under Article 5 create unexpected limitations. While patients can theoretically share their device data freely, the explicit exclusion of platforms designated as "gatekeepers" under the Digital Markets Act means patients cannot integrate pacemaker data with major tech companies' health platforms that they already might be using. This restriction, intended to prevent Big Tech dominance, may impair patients seeking consolidated health management tools.

More concerning, the Act's allowance for "reasonable compensation" for data access could transform the promise of free patient access into a cost barrier, particularly for complex medical devices requiring substantial infrastructure investments to enable sharing capabilities.

The Data Act's relationship with existing regulations adds layers to an already complex compliance puzzle[11]. The regulation explicitly preserves GDPR requirements, intellectual property rights, and Member States' prerogatives in public health and security, creating a multi-dimensional challenge where each regulatory framework pulls in different directions. A solution compliant with the Data Act might violate GDPR's data minimisation principles or compromise trade secrets protected under intellectual property law.

This ambitious scope reflects the EU's commitment to digital sovereignty but risks undermining the very innovation it seeks to democratise. Success requires more than technical compliance; it demands fundamental rethinking of product design, business models, and data governance strategies. Medical device companies must balance the Act's democratising vision with practical realities of device safety, inno-

8. https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html
9. Tschider, Charlotte and Corrales Compagnucci, Marcelo and Minssen, Timo, The New EU–US Data Protection Framework's Implications for Healthcare (September 27, 2024). Journal of Law and the Biosciences, volume 11, issue 2, 2024[10.1093/jlb/lsae022], Available at SSRN: https://ssrn.com/abstract=4983419
10. https://digital-strategy.ec.europa.eu/en/policies/data-act
11. Casolari, Federico & Buttaboni, Carlotta & Floridi, Luciano. (2023). The EU Data Act in Context: A Legal Assessment. SSRN Electronic Journal. 10.2139/ssrn.4584781.

vation cycles, and global competition. Those engaging early with the evolving regulatory framework and actively participating in guidance development may find opportunities within the constraints, and many may redirect innovation efforts to less regulated markets.

1.4. The EU AI Act[12]

While the Data Act struggles with making device data useful without compromising safety, the AI Act faces an even sharper dilemma. The EU AI Act, coming into force in August 2024 with full applicability by August 2027, adds a final layer of regulatory complexity as the world's first comprehensive legal framework for artificial intelligence systems. Its risk-based approach creates immediate categorisation challenges for medical AI, where the same technology might be prohibited in general use but permitted in healthcare contexts. Emotion recognition systems, normally banned, receive exemptions for medical purposes like psychological treatment, this flexibility that acknowledges healthcare's unique needs while creating interpretation challenges about what constitutes legitimate medical use.

The AI Act's risk categories translate into different compliance burdens. For example, high-risk AI systems, such as AI-assisted X-ray diagnosis, emergency triage systems, and medical training assessment tools face stringent requirements encompassing risk management protocols and human oversight. Low-risk systems like administrative AI for structured radiology reporting (unless integrated into medical devices) have minimal obligations. Approximately 75% of current AI medical devices will be classified as high-risk, requiring compliance with 16 distinct requirements within 36 months of the Act coming into force[13]. These requirements span continuous risk management, data governance protocols, extensive technical documentation, automatic event logging, fundamental rights assessments, and post-market surveillance systems, each adding layers of complexity and cost.

AI-enabled medical devices must satisfy both the AI Act's horizontal requirements and MDR/IVDR's vertical sector-specific regulations. Both frameworks demand risk management systems, technical documentation, and conformity assessments, but with different specifications, timelines, and interpretations[14]. Manufacturers of medical devices must make adjustments in a number of areas in order to comply with the requirements of the AI Act in addition to the requirements of the above-mentioned regulations.

The EU AI Act represents both a necessary step toward responsible AI governance and a potential impediment to medical innovation. While its comprehensive requirements may protect patients from harmful AI applications, the regulatory complexity, such as the overlap with existing frameworks, threatens to stifle the very innovation it seeks to regulate, or worse: create a compliance theatre with checkbox compliance.

These regulations are not merely coexisting but also compounding the useful-dangerous paradox. The EHDS makes data useful for cross-border care, DPF makes it useful for transatlantic research, Data Act makes it useful for device transparency, and AI Act makes it useful for algorithmic accountability. But each layer of usefulness multiplies dangers: more access points, more complexity, more attack surfaces, more ways for well-intentioned utility to become unintended vulnerability. This creates an urgent question: how can healthcare organisations enable all this data flow while maintaining privacy? The answer many are turning to is anonymisation, which appears to be deceptively simple until you examine what it actually takes to anonymise health data in a way that satisfies all regulatory frameworks simultaneously.

2. Anonymisation in healthcare datasets, an analysis to maintain utility

Anonymisation within the European healthcare sector has long been treated as a compliance afterthought, something that can be just "done" by removing direct identifiers only. The reality however is far more advanced. It is essential to distinguish between **pseudonymisation** and **anonymisation**, as the two are often confused or incorrectly used interchangeably. According to the European Data Protection Board, pseudonymised data remains personal data under the GDPR and does not meet the criteria for anonymisation[15]. The UK Information Commissioner's Office emphasises that anonymisation must render data incapable of identifying individuals, even when cross-referenced with other data[16].

In the European legal context, for data to be considered truly anonymised under GDPR, the risk of re-identification must be negligible when assessed against all means reasonably likely to be used[17]. This

12. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng
13. Emmanouil P Vardas, Maria Marketou, Panos E Vardas, Medicine, healthcare and the AI act: gaps, challenges and future implications, *European Heart Journal - Digital Health*, 2025;, ztaf041, https://doi.org/10.1093/ehjdh/ztaf04
14. Busch, F., Kather, J.N., Johner, C. *et al.* Navigating the European Union Artificial Intelligence Act for Healthcare. *npj Digit. Med.* 7, 210 (2024). https://doi.org/10.1038/s41746-024-01213-6
15. EDPB, guidelines 01/2025 on Pseudonymisation https://www.edpb.europa.eu/system/files/2025-01/edpb_guidelines_202501_pseudonymisation_en.pdf
16. ICO, *Anonymisation.* https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/anonymisation/
17. GDPR, recital 26: https://gdpr-info.eu/recitals/no-26/

bar is high, and many healthcare organisations are realising that anonymisation is no longer a fixed outcome, but a discipline requiring judgement, technical skills and constant adaptation.

This shift is driven by both regulatory expectations and operational needs. On the one hand, GDPR excludes anonymised data from its scope, creating strong incentives for sharing data in an anonymous format, particularly for secondary use, such as research, benchmarking or training (AI) algorithms. On the other hand, achieving this legal status is difficult in practice, especially when working with rich and complex health datasets. The technical challenges, institutional misunderstandings and a lack of consistent standards create real barriers to progress.

2.1. Key challenges in anonymisation

While awareness of anonymisation has grown and regulatory expectations have become clearer, many healthcare organisations, including those in pharma, healthtech and clinical research, still face significant challenges in practice. These range from technical constraints to persistent misconceptions, and they continue to complicate the safe and effective reuse of health data. The main challenges observed across these sectors are:

1. Risk of singling out through rare combinations:
Health data often includes detailed combinations of attributes such as age, treatment date, diagnosis, and geographic location. These combinations, while not directly identifying, can be unique. For instance, a 42-year-old patient with a rare diagnosis admitted to a specific regional hospital in March may be the only person matching that profile. This kind of uniqueness can enable re-identification even in the absence of names or IDs[18].

2. Ambiguity about what constitutes "sufficient" anonymisation:
Although the GDPR sets a high bar, it does not define measurable thresholds. Phrases like "reasonably likely" and "negligible risk" leave room for interpretation. In practice, this legal ambiguity makes it difficult for organisations to assess whether a dataset is anonymous enough to no longer fall in scope of the GDPR. This is particularly problematic when data is reused or shared across borders where interpretations may differ.

3. Limitations of current tools:
While anonymisation tools are improving, most are optimised for structured data and still require expert configuration. Automated settings often fail to capture the entire complexity of clinical or longitudinal datasets. Also, there are no tools built yet that cover all anonymisation challenges and techniques, and many tools struggle with the fact that no dataset is the same. For example, applying a uniform generalisation rule on identifiers might protect identity but can destroy important useful information in treatment timing or outcomes. Hence, at the moment, most datasets require a tailored approach that tools alone cannot provide.

4. Misunderstanding anonymisation as a one-time action:
A common misconception is that once data is anonymized, it stays that way forever. However, as new external data becomes available, new techniques emerge that adversaries can use or new regulations are implemented, a dataset that was previously considered safe may no longer meet the anonymisation standards. It is important to bear in mind that anonymisation is not a one-and-done process. As highlighted by the Spanish Data Protection Agency in their publication on common misunderstandings[19], it requires periodic re-evaluation, especially if data is re-used or shared in a new context, out of scope of the original consent.

5. Inconsistent adoption of advanced techniques:
While some organisations have integrated advanced techniques like hashing, differential privacy and k-anonymity[20] to mitigate risk, implementation is often fragmented. One research team may apply a secure hashing method, such as SHA-512, while another uses a weaker approach, like simple randomisation to generate pseudonyms. The adoption of most advanced techniques is often hindered by a lack of technical expertise and proper tooling or other resources. This inconsistency undermines trust and increases the likelihood of either over-anonymisation and under-protection, so a higher risk of leaving room for reversing the techniques used and allowing the adversary to re-identify subjects.

2.2. Synthetic data as a solution to maintain utility?

Synthetic data is often presented as a privacy-preserving solution, but its limitations are frequently overlooked. Multiple studies show that disputes remain over whether synthetic datasets leak identifiable patterns or maintain analytic utility. A narrative review found concerns that synthetic data can "not replicate precisely the content and properties of the original dataset," leading to risk of data leakage[21]. Another recent study described a "data-sharing paradox" in healthcare: synthetic data is designed for sharing yet often remains overly re-

18. ICO, *Anonymisation*. ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/anonymisation/
19. AEPD, *10 Misunderstandings related to Anonymisation*. https://www.aepd.es/guides/10-anonymisation-misunderstandings.pdf
20. AEPD, *K-anonymity as a privacy measure https://www.aepd.es/guides/k-anonymity-as-a-privacy-measure.pdf*
21. Gonzales et al., *Synthetic data in health care: A narrative review*, PLOS Digital Health (2023) https://pmc.ncbi.nlm.nih.gov/articles/PMC9931305/

stricted due to ambiguous re-identification risk metrics[22].

Based on our experiences within many health-related organisations, synthetic data is currently seen as not yet mature enough for widespread healthcare use. For it to be safe and effective, it must be demonstrably free from identifiable patterns and retain sufficient clinical relevance, a delicate balance that has not been consistently achieved.

An example might be if, if synthetic patient data were generated for a rare neuromuscular disease cohort and researchers relied upon this data to develop diagnostic models. If the synthetic dataset closely mirrored a unique, real patient (e.g. combining rare genetic markers with clinical outcomes), there is a re-identification risk. Conversely, if too much distortion is applied, critical genotype–phenotype links may break, rendering the data scientifically useless. Without rigorous privacy and utility evaluations, the data may expose patients' identities or invalidate research outcomes.

A 2024 study in *Nature* confirms these risks, warning that synthetic data may allow membership inference or re-identification attacks, particularly when auxiliary datasets are available[23]. Until standard metrics and rigorous governance frameworks exist, synthetic data remains a useful tool, but not yet a plug-and-play alternative for anonymisation.

2.3. How to maintain as much utility as possible?

To make data useful without making it dangerous, healthcare organisations must begin treating anonymisation as a strategic capability, not just another technical step. Leading organisations are already moving in this direction by embedding anonymisation into their broader data governance programs. This includes incorporating risk assessments during project design, adopting context-based transformation strategies, and assigning clear responsibilities for reviewing anonymisation outcomes.

Crucially, anonymisation must be integrated early into the data life cycle, well before data is shared or published (anonymisation-by-design). This requires the involvement of qualified experts who understand both the technical risks and the regulatory landscape, as well as the clinical context of the data. Without this, organisations risk either over-transforming data and losing its utility value, or under-protecting it and exposing patients to privacy harms, therefore making it too dangerous.

Several sector-specific initiatives are helping to close this gap. MedTech Europe, for example, has developed a practical anonymisation framework to support its members in applying structured, risk-based approaches. While not yet formally published as standalone guidance, the framework was outlined in a recent MedTech Europe article and aims to offer clear starting points for organisations that need to operationalise anonymisation within compliant data sharing practices[24].

Ultimately, anonymisation done well enables the secondary use of health data, so maintaining as much utility as possible while safeguarding individual privacy. It allows organisations to meet GDPR expectations, build public trust, and support responsible innovation. But this balance cannot be achieved through a single action or via generic tools only. Proper anonymisation must be seen for what it really is: a discipline that requires ongoing investment, expert judgement and continuous attention. This ensures a maximum of utility without the data becoming dangerously harmful.

Even the most sophisticated anonymisation strategies require protection, especially when data is shared publicly or across borders. Protecting information in motion and at rest becomes just as critical as transforming it at the source. This is where cybersecurity steps in, not as a final barrier, but as a dynamic additional safeguard for possible and potential outside threats.

3. Cybersecurity in healthcare: balancing protection, privacy and progress

The Ireland Health Service Executive attack mentioned in the introduction exemplifies a pattern that has only intensified since 2021. When Germany's Düsseldorf University Hospital suffered a ransomware attack in 2020, the consequences went beyond data breaches when a patient died after emergency services had to be rerouted to a more distant hospital[25]. Such incidents reveal healthcare's fundamental challenge: every connection that enables better care also creates new vulnerabilities.

This question becomes more urgent as Europe implements ambitious data-sharing frameworks in 2025. Every technological advance that promises better care through connected systems simultaneously creates new vulnerabilities. The challenge isn't simply

22. Achterberg et al., *The Data Sharing Paradox of Synthetic Data in Healthcare* (2025) https://arxiv.org/html/2503.20847v1
23. Recent *Nature Scientific Reports* study on membership inference and re-identification risk in healthcare synthetic data https://www.nature.com/articles/s41598-024-72894-y
24. van Kempen, R. *Can Europe unlock the power of data while protecting privacy?*, MedTech Europe (2024). https://www.medtecheurope.org/medtech-views/policy-views/can-europe-unlock-the-power-of-data-while-protecting-privacy/
25. Ransomware's impact on patient care, including the first reported death: https://arstechnica.com/information-technology/2020/09/patient-dies-after-ransomware-attack-reroutes-her-to-remote-hospital/_

technical. It reflects deeper tensions between competing imperatives that cannot be permanently resolved, only dynamically balanced.

3.1. The data paradox: connected care, multiplied risk

Healthcare data possesses unique characteristics that intensify such tensions. Clinical records combine personally identifiable information, protected health information, payment details, and increasingly, biometric markers within single datasets. For pharmaceutical and life sciences organisations, patient data coexists with research findings and manufacturing processes representing millions in intellectual property value[26].

The Northeast Radiology breach illustrates how data's inherent value makes it vulnerable. Unauthorized access persisted from April 2019 through January 2020, with discovery only occurring in March 2020. By then, 298,532 patient records had been exposed[27]. The breach's extended timeline reveals a crucial insight: the very characteristics that make health data valuable for longitudinal care make it attractive for theft. Rich, comprehensive records enable better treatment decisions but also command premium prices in criminal markets.
Recent statistics underscore this reality. Ransomware accounts for 54% of reported cyber incidents in EU healthcare, with 43% including confirmed data leaks according to ENISA findings[28, 29]. These aren't merely IT disruptions. They represent fundamental breakdowns in healthcare's ability to simultaneously share and protect information.

The Medefer case in early 2025 demonstrates how modern interoperability solutions embody these contradictions[30]. An authentication flaw in APIs used by this NHS contractor exposed patient referral data for years. The system was functioning exactly as designed, efficiently sharing information between providers. Its vulnerability stemmed from utility. Making data flow seamlessly between organisations meant reducing friction, and reduced friction meant fewer barriers for both authorised and unauthorised access.

This paradox extends beyond individual breaches. Healthcare's push toward integrated care requires breaking down information silos that, however inefficient, provided natural segmentation against attacks. Modern initiatives like integrated data lakes and FHIR APIs promise transformative benefits: population health insights, precision medicine breakthroughs, rapid research capabilities. Yet aggregating scattered data into centralised, accessible platforms creates what security professionals recognise as high-value targets.

Consider how security context typically disappears as data moves through integration layers. Access controls, classifications, and audit trails that protect information in its original system often get stripped away when data flows through APIs into analytics platforms. Once decoupled from native protections, this information gets reused in dashboards, research datasets, and decision support tools with minimal oversight. The very processes that unlock data's potential simultaneously erode its protection[31].

3.2. The innovation trap: racing ahead of protection

Healthcare's digital transformation accelerated dramatically during the COVID-19 pandemic, with telehealth adoption jumping from peripheral service to core delivery mechanism virtually overnight. This transformation saved lives by maintaining care continuity during lockdowns. It also opened vast new attack surfaces faster than organisations could secure them[32].

The pandemic revealed how utility demands can overwhelm security capabilities. Healthcare organisations deployed collaboration platforms, remote access systems, and patient engagement tools at unprecedented speed. Each new capability addressed urgent care needs but also expanded the perimeter that security teams needed to defend. Traditional security models built on controlling defined network boundaries became obsolete when care delivery itself became boundaryless.

This acceleration continues with artificial intelligence adoption. Healthcare organisations rush to implement AI for diagnostics, treatment planning, and operational efficiency. Yet research indicates most lack AI-specific governance frameworks[33]. The

26. Pharmaceutical companies and life sciences organisations face unique challenges in protecting both patient data and intellectual property. See: https://www.techtarget.com/healthtechsecurity/news/366594393/Inadequate-Healthcare-Cybersecurity-Maturity-Jeopardizes-Patient-Privacy
27. Northeast Radiology settlement details: https://www.hhs.gov/press-room/hhs-ocr-hipaa-settlement-nerad.html
28. ENISA Health Threat Landscape Report, p. 3, p. 13: https://www.enisa.europa.eu/sites/default/files/publications/Health%20Threat%20Landscape.pdf
29. European Commission on healthcare cybersecurity: https://commission.europa.eu/news/bolstering-cybersecurity-healthcare-sector-2025-01-15_en
30. NHS API vulnerability investigation: https://www.computerweekly.com/news/366620174/NHS-investigating-how-API-flaw-exposed-patient-data
31. API security incidents in healthcare have risen sharply, with 79% of organisations experiencing incidents: https://www.hipaajournal.com/79-of-healthcare-organisations-experienced-an-api-security-incident-in-the-past-12-months/
32. COVID-19's impact on healthcare cybersecurity: https://pmc.ncbi.nlm.nih.gov/articles/PMC8059789/
33. HIMSS report on AI governance gaps: https://www.himss.org/news/report-health-system-cybersecurity-budgets-increasing-lack-ai-governance-threatens-security

pattern repeats: transformative utility drives adoption before protective measures mature. Each wave of innovation creates new imbalances that organisations struggle to address while already managing previous ones.

The innovation trap isn't simply about moving too fast. It reflects how healthcare's mission creates different risk calculations than other sectors. When a new technology might improve patient outcomes, the ethical imperative to adopt it can override security concerns. This calculation makes moral sense in individual cases but creates systematic vulnerabilities when repeated across thousands of decisions.

3.3. The architecture of vulnerability

Healthcare's technical landscape reveals how historical decisions about balancing access and protection accumulate into current vulnerabilities. Organisations operate environments mixing 20-year-old medical devices with cloud-native applications, multivendor systems with proprietary protocols, and safety-critical applications running on unsupported operating systems. Each element represents a past decision where immediate utility took precedence, creating what security teams now experience as a nearly unmanageable attack surface[34].

Medical devices are a great example of these accumulated trade-offs. A ventilator or MRI machine designed for a 20-year service life cannot be patched like consumer electronics. The device critical for care simultaneously endangers the network through unpatched vulnerabilities[35].
Legacy clinical systems present similar dilemmas. Picture Archiving and Communication System ("PACS") imaging archives, laboratory information systems, and specialised departmental solutions often run on outdated platforms because replacement would disrupt care delivery. These systems grew in isolation, use unique or outdated data formats and protocols that are not supported by modern security tools. A Digital Imaging and Communications in Medicine ("DICOM") imaging file that enables critical diagnostics may be invisible to security scanners designed for conventional documents.

The architectural complexity multiplies through third-party dependencies. Modern healthcare relies on intricate webs of vendors, from electronic health record ("EHR") providers to device manufacturers, imaging centres to laboratory networks. Each connection enables essential services but also creates potential entry points. Supply chain attacks exploit these relationships, compromising trusted partners to reach ultimate targets. The same interconnections that enable coordinated care, also enable coordinated attacks.

3.4. The maturity spectrum in the sector

Cybersecurity maturity in healthcare remains uneven across subsectors. Providers, particularly hospitals and clinics, tend to lag behind pharma and life sciences, where protecting intellectual property and complying with regulations have driven earlier investments in security. Medical device manufacturers are only recently shifting from a focus on safety to one that includes robust digital protection, responding to both rising connectivity and regulatory scrutiny.

The result is a patchwork with islands of resilience within a broader environment that is still catching up. But this unevenness also reflects deeper questions of strategic priorities. Organisations more advanced in protecting security tend to be those that have explicitly recognised the tension between utility and risk and invested accordingly.

By 2025, a significant shift is underway. Healthcare technologies are increasingly being built with embedded security from the start. Whether it's a medical device, mobile health application, or an EHR system, incorporating protective measures early in development is now the new norm. This evolution is being shaped by regulations such as the U.S. Food and Drug Administration's cybersecurity guidance and the EU's Cyber Resilience Act, as well as the cumulative impact of previous high-profile security lapses.

3.5. Toward adaptive resilience

Healthcare can accelerate progress by learning from sectors that faced similar challenges in securing complex and distributed systems, such as finance and telecommunications. Over the past decade, these sectors have built mature practices around API governance, zero-trust architecture, and real-time auditing across cloud environments. While healthcare's risks are unique in their clinical consequences, the underlying technical problems are often shared.

The key insight from mature sectors is that balance points shift constantly. What works during normal operations fails during crisis. What protects adequately today becomes vulnerable tomorrow. Organisations need capabilities to sense these shifts and adjust accordingly rather than seeking permanent solutions. Static security models that try to lock down systems inevitably fail because they conflict with care delivery needs.

This adaptive approach extends to governance mod-

34. Medical device security evolution: https://galendata.com/updating-cybersecurity-for-advanced-medical-devices-2024-insights-and-best-practices/

35. Government investigation into medical device vulnerabilities: https://www.reuters.com/article/technology/us-government-probes-medical-devices-for-possible-cyber-flaws-idUSKCN0IB0DQ/

els. Instead of centralised control that stifles innovation or distributed chaos that enables breaches, leading organisations create flexible frameworks. These establish baseline requirements while allowing justified variations. They emphasise outcomes over compliance, measuring whether data remains protected rather than whether specific controls exist.

Healthcare's digital future requires accepting that making data useful without making it dangerous isn't a problem to solve, but a tension to manage. Every advance in care delivery through connected systems creates new vulnerabilities. Every security measure that adds friction potentially delays treatment. Every confidentiality protection that segments data might prevent crucial insights.

4. Conclusion

The Privacy-Security-Utility Triangle we've explored reveals that Europe's health data transformation is not heading toward a destination but embarking on a continuous journey. Each vertex of this triangle, regulation, anonymisation, and security pulls in its own direction, creating tensions that cannot be resolved but must be dynamically balanced.

For healthcare leaders navigating this landscape in 2025 and beyond, the key insight is to stop seeking perfect solutions and start building adaptive capabilities. Success will belong to organisations that can sense when balance points shift, adjust protections without paralysing innovation, and maintain resilience while enabling transformation. In healthcare, both the music and the stakes keep changing. In this dance, standing still is not an option; only those who keep moving, learning, and adapting will thrive in Europe's bold new health data ecosystem.