White Label Consultancy

Health Data Beyond Healthcare: Managing Inferred Health Risk in Digital Ecosystems

November 2026

Table of Contents

1. **Introduction** 3

2. **From Consumer Data to Health Insight: Structural Drivers and Key Risks** 4
   - *Inferred Health Data and the Special Category Threshold* 5
   - *Continuous Monitoring and Ambient Inference* 5
   - *Connected Ecosystems and Diffused Accountability* 6
   - *AI Profiling, Health Prediction, and High-Risk Processing* 6

3. **The Regulatory Landscape for Inferred Health Data** 7
   - *Layer 1 - General Data Protection Law (the Baseline)* 7
   - *Layer 2 - Cross-Border Transfer Restrictions: Adequacy versus Appropriate Safeguards* 9
   - *Layer 3 - Health-Sector Localisation* 10
   - *Layer 4 - IoT-Specific Law* 11
   - *Comparative Overview* 11
   - *Why the Patchwork Leaves a Gap* 13

4. **What Regulators Expect: Core Compliance Mechanisms** 14
   - *Recognising Inferred Special Category Data* 14
   - *Privacy by Design and by Default in Connected Products* 14
   - *Consent Engineering Done Properly* 15
   - *Impact Assessments Calibrated to Inferred Data* 15
   - *Transparency About Inferences* 16
   - *AI and Profiling Governance* 17

5. **The Governance Gap: Why Inferred Health Risk Escapes Existing Controls** 17

6. **Practical Governance Framework for Organisations** 18
   - *Step 1: Map Where Ordinary Data Can Yield Health Inferences* 18
   - *Step 2: Assess Special Category Thresholds* 18
   - *Step 3: Conduct Inference-Calibrated Impact Assessments* 18
   - *Step 4: Embed Privacy by Design and Clean Consent Engineering* 19
   - *Step 5: Govern AI and Profiling Systems* 19
   - *Step 6: Coordinate Privacy, AI, Cybersecurity, Product, and Risk Functions* 19
   - *Step 7: Provide Transparency and Data-Subject Controls Over Inferred Data* 19
   - *The Framework at a Glance* 19

7. **Conclusion: Governing Health Data Beyond Healthcare** 20

1. Introduction

What qualifies as health data under privacy legislation is often nuanced and not always obvious. Personal data does not need to be collected in a healthcare context to constitute health data. Behavioural patterns, wearable-device outputs, voice recordings, location data, purchasing habits, and other seemingly ordinary data points may, individually or in combination, reveal information about an individual's physical or mental health, disease risks, or physiological condition. As organisations increasingly rely on advanced analytics and artificial intelligence to generate such insights, they may find themselves processing health data without ever collecting traditional medical information.

Consider an organisation that develops a calorie-counting application, a wellness platform that assesses stress levels from typing cadence and voice recordings, or an insurer that analyses behavioural data to refine pricing decisions. None of these organisations may view itself as handling medical records. Yet the inferences they generate may reveal information about a person's health and therefore attract the heightened protections associated with health data under privacy law. The challenge is not merely that health information can be inferred, but that organisations do not always recognise when such inferences cross the threshold into regulated health data.

This distinction matters because health data occupies a special position under modern privacy frameworks. Under the GDPR and the many privacy frameworks modelled on it,[1] data concerning health forms part of the special or sensitive categories of personal data and is subject to additional legal requirements and safeguards. Misclassifying health-related information as ordinary personal data can therefore have significant compliance consequences, affecting lawful basis assessments, the additional conditions for processing, data protection impact assessments, retention practices, vendor management, and international transfers.

This paper argues that because health data can be inferred as readily as it is collected, the centre of gravity of health-data governance is shifting from what organisations collect to what they can infer. It maps the structural drivers and key risks of health inference (**Section 2**); sets out the overlapping layers of regulation that can apply at once and what regulators expect organisations to do (**Sections 3 and 4**); and explains why inferred health data so often escapes existing controls (**Section 5**). It then offers a practical, implementation-ordered governance framework to ensure that inferred health data is properly identified, classified, and governed (**Section 6**).

The analysis is anchored in the GDPR, which offers one of the most developed frameworks for health and inferred data, but it examines the same problem across the United Kingdom, the UAE (including the DIFC and ADGM free zones), Saudi Arabia, the United States, Canada, and China, because an organisation operating internationally must satisfy several of these regimes at once.

[1] Regulation (EU) 2016/679 (GDPR), Arts. 4(15) and 9(1); Recital 35. Comparable provisions appear in, among others, the UK GDPR, Art. 9; the DIFC Data Protection Law 2020, Art. 11; the ADGM Data Protection Regulations 2021, s. 7; the UAE PDPL (Federal Decree-Law No. 45 of 2021), Art. 1; and the Saudi PDPL, Art. 1. See Section 3.

2. From Consumer Data to Health Insight: Structural Drivers and Key Risks

Regulatory guidance and existing research have highlighted several ways in which seemingly ordinary personal data can reveal health-related information.[2] Three interconnected factors are particularly relevant in this context:

1. **Sensor-based collection:** wearables, connected products, and the broader Internet of Things generate dense, uninterrupted streams of behavioural and physiological data.

2. **Inference:** analytics derive health-related conclusions from data that is not, on its face, health data.

3. **Invisibility:** these inferences emerge indirectly and often go unrecognised by the user and the organisation alike.

Together, these factors blur the line between ordinary personal data and health data. As a result, digital-economy companies that never set out to process health data may do so without realising it.

Inference deserves particular emphasis. Step counts, scrolling patterns, purchase histories, sleep data, and typing cadence can all become proxies for conditions, states, and risks; as the Swiss Data Protection Authority notes, data measured directly on the body is "highly personal", and the patterns it yields enable "conclusions to be drawn about users' health, lifestyle habits and private lives".[3]

This matters because health data is defined broadly. Under Article 4(15) GDPR, "data concerning health" means personal data relating to a natural person's physical or mental health, including the provision of health care services, which reveal information about their health status. Recital 35 confirms that this covers information about past, current or future health status regardless of its source, and is not confined to data collected in a clinical setting.[4]

The Court of Justice has taken the same broad approach in *Lindenapotheke*, holding that the data a customer enters when ordering pharmacy-only medicines online constitutes data concerning health, because it is capable of revealing information about the customer's health status "by means of an intellectual operation involving collation or deduction", and that this is so even where the medicine is only probably, not certainly, intended for the customer.[5] In *Meta Platforms v Bundeskartellamt*, it held that collecting a user's visits to websites and apps relating to health (among other Article 9 categories) and linking them to the user's account is processing of special category data where it

[2] See: EDPB-EDPS Joint Opinion 03/2022 on the European Health Data Space (12 July 2022); the ICO's Guidance for consumer Internet of Things products and services (11 June 2026); Article 29 Working Party, Opinion 8/2014 on Recent Developments on the Internet of Things (WP 223, 16 September 2014); and Article 29 Working Party, Annex on health data in apps and devices (letter to the European Commission, 5 February 2015). For academic analysis, see S. Wachter and B. Mittelstadt, 'A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI' (2019) Columbia Business Law Review 494; D. J. Solove, 'Data Is What Data Does: Regulating Based on Harm and Risk Instead of Sensitive Data' (2024) 118 Northwestern University Law Review 1081.

[3] Swiss FDPIC, "Wearables: Smartwatches, fitness trackers, smart glasses" (edoeb.admin.ch).

[4] GDPR Recital 35.

[5] Case C-21/23 *Lindenapotheke* (CJEU, Grand Chamber, 4 October 2024), paras 83-84 and 90-91.

allows such information to be revealed, whether or not the information revealed is correct.[6] Neither case involved a medical record: in each, ordinary transactional or browsing data crossed the Article 9 threshold because of what it revealed.

Inferred Health Data and the Special Category Threshold

The threshold question for any service is whether it processes health data, including data it never set out to collect. The Article 29 Working Party's guidance on apps and devices gives the assessment its practical shape: data is health data where it is inherently medical; where it is raw sensor data that can be used, alone or in combination with other data, to draw a conclusion about a person's actual health status or health risk; or where conclusions about health status or risk are in fact drawn, whether or not they are accurate. A single weight or step-count reading may fall outside the category, but the same measure tracked over time, combined with age and sex, or transferred to a third party can move inside it.[7] Sleep, nutrition, mood, and skin-scan data will often cross the threshold, particularly once tracked over time or combined, so a "wellness" or "lifestyle" positioning offers no shelter in itself.

Under the GDPR, once that threshold is crossed, processing is prohibited pursuant to Article 9(1) unless one of the conditions in Article 9(2) applies. For most consumer services, that will mean the explicit consent of the data subject. An Article 9(2) condition is required in addition to, not instead of, a lawful basis under Article 6.[8] The Article 6 basis chosen for the original, non-health purpose may not carry over either: where deriving health information constitutes a new purpose, its compatibility with the original purpose must be assessed under Article 6(4), which expressly identifies the processing of special category data as a factor in that assessment.[9]

Further obligations may also arise, including a data protection impact assessment where special category data is processed on a large scale and, where such processing is a core activity, the designation of a data protection officer; transparency about the inference itself; a prohibition on solely automated decisions producing legal or similarly significant effects based on the inferred data unless explicit consent or substantial public interest applies; and closer scrutiny of sharing and international transfers, since each disclosure requires its own Article 9 condition and any transfer risk assessment must reflect the sensitivity of the data.[10] These obligations can arise without the organisation ever deciding to handle health data, and that is precisely what makes them easy to miss.

Continuous Monitoring and Ambient Inference

Continuous monitoring devices, such as smartwatches, fitness bands, and smart rings, are the clearest illustration of how data tracked over time can cross the threshold. They reveal health status, behaviour, and lifestyle at a granularity earlier services could not approach, and the picture sharpens as data accumulates: a single reading may fall short of health data, but a year of readings gathered quietly in the background can describe a person's health with uncomfortable precision. The sensitivity is emergent, often

[6] Case C-252/21 *Meta Platforms and Others v Bundeskartellamt* (CJEU, Grand Chamber, 4 July 2023), paras 68-73 and operative part, point 2.

[7] Article 29 Working Party, Annex - health data in apps and devices (5 February 2015), pp. 3-5.

[8] GDPR, Recital 51; ICO, "What are the rules on special category data?"

[9] GDPR, Arts. 5(1)(b) and 6(4).

[10] GDPR, Arts. 13-14, 22(4), 35(3)(b) and Chapter V.

crystallising through aggregation long after collection, which is one reason it escapes the controls applied at the point of data capture. Smart glasses and always-on sensors add a further layer, ambient inference: health signals derived from a person's surroundings, and from the people in them, without any deliberate act of data entry. The problem therefore extends beyond the wearer to bystanders who have no relationship with the provider and no practical way to object.

Connected Ecosystems and Diffused Accountability

The Internet of Things refers to networks of physical objects, from wearables and smart speakers to home hubs and connected appliances, that are embedded with sensors and connectivity. These products rarely operate alone, and the resulting chain of devices, apps, analytics partners, and cloud layers diffuses accountability. A sensitive inference may be generated at a layer of the stack that the device maker neither sees nor controls, and controller and processor roles across the ecosystem are frequently unclear. The assumption that an organisation knows what data it holds breaks down exactly where the sensitivity is highest.

The UK regulator has addressed this directly. In its guidance on consumer Internet of Things products and services, the Information Commissioner's Office (ICO) recognises that such products may process special category data either directly or by inference, giving the example of a device that intentionally infers health data in order to give the user a health score. It also cautions that most processing involving IoT products is likely to result in a high risk, so that a data protection impact assessment will usually be required.[11] Although the guidance applies the UK GDPR, the relevant provisions mirror the EU GDPR and the reasoning carries across.

AI Profiling, Health Prediction, and High-Risk Processing

Predictive systems can assign people to inferred health categories, such as likely pregnant, likely depressed, or likely managing a chronic condition, on the basis of behaviour that is not medical on its face, raising acute questions of fairness, discrimination, and lawful basis. The distinguishing feature is opacity: these are often black-box models, continuously updated and imperfectly understood even by the organisations that deploy them, so that neither the user nor, frequently, the deployer can explain how a particular conclusion was reached.[12] This makes the transparency obligations described above difficult to meet in practice.

Risk also escalates with scale, and this is where inference becomes hardest to manage. Where an organisation knows from the outset that it is processing health data at scale, the obligations are at least visible and can be planned for. The difficulty with inferred health data is that the organisation frequently does not know: it can be genuinely hard to tell whether a system is generating incidental, minimal health inferences or whether it has tipped into large-scale processing of special category data, which brings in the full set of obligations attached to such processing, including a mandatory DPIA and potentially a

[11] ICO, "New guidance to help smart product manufacturers get data protection right" (16 June 2025); ICO, "Setting out our expectations for the smart device industry" (11 June 2026).

[12] S. Wachter and B. Mittelstadt, "A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI" (2019) 2 Columbia Business Law Review 494; Article 29 Working Party Guidelines on Automated individual decision-making and Profiling (WP251rev.01), endorsed by the EDPB.

DPO. Separately, an AI system that infers health in a sensitive context, such as risk assessment and pricing in life and health insurance, recruiting or evaluating workers, or recognising emotion (the last of these prohibited outright in the workplace and in education, except for medical or safety reasons), may be classified as high-risk under the EU AI Act, triggering conformity assessment and ongoing obligations for both providers and deployers.[13] The same data that looks low-stakes in a single app can therefore, at scale or inside an AI system, carry obligations of an entirely different order.

3. The Regulatory Landscape for Inferred Health Data

Even once an organisation knows it holds health data, the law that applies to a consumer health product in the tech space is not a single regime that can be satisfied with a single control. It is a stack of overlapping layers, and one product can trigger several of them at once. To make this concrete, consider four applications that do not regard themselves as "medical": a sleep tracker; a skin-scan app that assesses moles, scars, and skin conditions; a calorie and nutrition tracker; and a mood-and-stress app of the kind used for habit-building or meditation. Each processes data the law may treat as health data, and each can engage up to four distinct layers of regulation depending on how it is built and where it operates. The four layers are set out below, followed by a comparative table and jurisdiction-level notes.

Layer 1 - General Data Protection Law (the Baseline)

Comprehensive privacy laws form the baseline, and across them health data is consistently singled out for treatment that ordinary personal data does not attract. Under the GDPR, the UK GDPR, China's PIPL, and the DIFC and ADGM data protection laws, it is a special or sensitive category that requires an additional condition, over and above an ordinary lawful basis, before it may be processed.[14] These regimes also layer heightened accountability duties on top of the rules on lawful processing: for health data a data protection impact assessment is mandatory where it is processed on a large scale (and, in some regimes, whenever sensitive data is processed at all), and a Data Protection Officer must be appointed where such processing is a core activity or reaches a prescribed scale.

Saudi Arabia's PDPL classifies health data as sensitive and, although it has no free-standing processing condition for sensitive data, it excludes legitimate interest as a basis for processing it, requires explicit consent where consent is relied upon, imposes specific handling controls on health data, and requires an impact assessment whenever sensitive data is processed and a Data Protection Officer where processing sensitive data is a core activity;[15]

[13] EU AI Act (Reg. (EU) 2024/1689), Art. 5(1)(f) and Annex III, points 1(c), 4(a) and 5(c); Annex III obligations apply from 2 December 2027 following Regulation (EU) 2026/1744 (Digital Omnibus on AI, in force 27 July 2026); European Commission, draft Guidelines on the classification of high-risk AI systems under Art. 6 (19 May 2026; consultation extended to 23 July 2026; not yet adopted).

[14] GDPR, Art. 9(1)-(2); UK GDPR, Art. 9 and Data Protection Act 2018, Sch. 1; PIPL, Arts. 28-29; DIFC Data Protection Law, DIFC Law No. 5 of 2020, Art. 11; ADGM Data Protection Regulations 2021, s. 7(2).

[15] Saudi PDPL (Royal Decree M/19 of 2021, as amended by M/148 of 2023), Arts. 1, 6(4), 22, 23 and 30(2); Implementing Regulation (September 2023),

Canada's PIPEDA, which applies to data handled in the course of commercial activity, treats health information as almost always sensitive, which raises the bar for consent and safeguards.[16] The Canadian regulator has applied that approach to inferences directly, finding that health-related browsing used to infer an advertising interest was sensitive personal information for which express consent was required.[17] Federal reform is also in train: Bill C-36, introduced in June 2026, would define personal information as including information inferred about an individual, list health information as sensitive, and require a privacy impact assessment before personal information is transferred outside Canada.[18]

In the UAE, three data protection regimes operate in parallel. At federal level the Personal Data Protection Law applies generally; within the financial free zones, the DIFC Data Protection Law 2020 and the ADGM Data Protection Regulations 2021 apply in its place. Each of the three defines a category of special or sensitive personal data that includes health data, in terms broadly similar to one another and to the GDPR. The federal PDPL, however, does not impose an additional processing condition for sensitive data, although processing it can trigger additional obligations, such as a data protection impact assessment under Article 21 or the appointment of a Data Protection Officer under Article 10. The PDPL Executive Regulations had not been issued at the time of writing.[19]Health data is then addressed more specifically by the federal ICT Health Law, which governs health information generated in connection with health services in the UAE and sets out its own definition of health data together with its own handling and localisation requirements.[20]

The United States offers the clearest contrast, because it has no comprehensive federal privacy law. HIPAA is the principal federal health-privacy statute, but it binds only defined "covered entities" such as providers, plans, and clearinghouses, and their business associates, so consumer and wellness apps that are not provided by or on behalf of a

Arts. 11(2), 16(1)(c), 25(1)(a), 26 and 32(1)(c); Regulation on Personal Data Transfer outside the Kingdom (Version 2.0, August 2024), Arts. 4(2)(b), (d) and (e) and 7(1)(b).

[16] PIPEDA, s. 4(1)(a) and Schedule 1, cl. 4.1.3, 4.3.4, 4.3.6 and 4.7.2; OPC, Guidelines for obtaining meaningful consent (2018). Quebec: Act respecting the protection of personal information in the private sector, CQLR c. P-39.1, s. 17. Nova Scotia: Personal Information International Disclosure Protection Act, s. 5 (public bodies).

[17] OPC, PIPEDA Report of Findings 2014-001 (Google Inc., 14 January 2014): health-related browsing used to infer an advertising interest category was sensitive personal information for which express consent was required (Principles 4.3 and 4.3.6).

[18] Bill C-36, An Act to enact the Protecting Privacy and Consumer Data Act, to amend the Personal Information Protection and Electronic Documents Act and to make amendments to other Acts, 45th Parliament, 1st Session (first reading 15 June 2026), ss. 2(1) ("personal information [...] including information that is inferred about the individual"; "sensitive" including health information), 15(5) and 57.

[19] UAE PDPL (Federal Decree-Law No. 45 of 2021), Arts. 1 ("Sensitive Personal Data"), 2(2)(e), 4, 10(1) and 21; the Executive Regulations provided for by the Decree-Law, which are to specify the thresholds for both obligations, had not been issued at the time of writing (September 2026).

[20] UAE PDPL (Federal Decree-Law No. 45 of 2021), Art. 2(2)(e) (exclusion of personal health data governed by its own legislation) and Arts. 22-23; DIFC Data Protection Law, DIFC Law No. 5 of 2020, Arts. 11 and 26-27; ADGM Data Protection Regulations 2021, ss. 7 and 41-42; UAE ICT Health Law (Federal Law No. 2 of 2019), Arts. 1, 2, 13 and 24 (official English translation; the Arabic text prevails); Cabinet Resolution No. 32 of 2020 (Executive Regulation; not reviewed); Ministerial Decision No. 51 of 2021 (permitted cases for storage and transfer abroad; content taken from published summaries).

covered entity generally fall outside it. The gap is filled by state law. Comprehensive state privacy statutes, led by the California Consumer Privacy Act as amended by the California Privacy Rights Act, classify health information, alongside precise geolocation and the contents of communications, as "sensitive personal information", A newer wave of consumer-health-data rules, beginning with Washington's My Health My Data Act and followed by Nevada, Connecticut and Maryland, define consumer health data broadly enough to capture inferences (Washington and Nevada expressly reach data "derived or extrapolated" from non-health information) and reaches well beyond HIPAA. The Federal Trade Commission adds enforcement under its consumer-protection powers and the Health Breach Notification Rule. For an organisation operating in the United States, the practical question is therefore rarely whether HIPAA applies, but which combination of federal and state regimes does.[21]

Layer 2 - Cross-Border Transfer Restrictions: Adequacy versus Appropriate Safeguards

Once an application processes this data, moving it across borders, for cloud hosting, analytics vendors, or an overseas parent company, triggers transfer rules. Under the GDPR model there are two broad routes: where the destination benefits from an adequacy decision, data may flow without additional steps; where it does not, the exporter must put appropriate safeguards in place, typically standard contractual clauses or binding corporate rules, or rely on one of a narrow set of derogations.[22] A sleep app whose EU users' data is hosted on US servers needs one of these mechanisms before the data leaves the EEA.

The defining feature of this layer is that it is conditional: data may move provided the correct mechanism is used. The adequacy-versus-safeguards framing is, however, the EU and UK lens, and other regimes impose their own and sometimes more demanding transfer controls. Saudi Arabia requires a documented risk assessment before sensitive data is transferred abroad on a continuous or widespread basis and withholds several of its transfer exemptions where the data is sensitive; China requires a formal security assessment for certain categories of data.[23] The United States, which has no general transfer regime, has since 2025 prohibited or restricted transfers of bulk personal health data to designated countries of concern, and its definition of personal health data reaches basic physical measurements, vital signs and logs of exercise habits, so that wearable-derived data falls squarely within it.[24] The

[21] HIPAA, 45 CFR 160.103 (covered entity), 160.102(b) and 164.502(a)(3) (business associates); Cal. Civ. Code 1798.140(ae)(1)(C), (1)(E) and (2)(B) (sensitive personal information), 1798.140(v)(1)(K) (inferences) and 1798.121 (right to limit); Washington My Health My Data Act, RCW 19.373.010(8)(b)(xiii) (effective 31 March 2024); Nevada SB 370 (2023), NRS 603A.430(1)(d); Connecticut Public Act 23-56, CGS 42-515 and 42-526 (consumer health data provisions effective 1 October 2023); Maryland Online Data Privacy Act 2024 (effective 1 October 2025); FTC Act s. 5, 15 USC 45; Health Breach Notification Rule, 16 CFR Part 318, as amended by 89 FR 47028 (effective 29 July 2024).

[22] GDPR, Arts. 44-49 (Chapter V).

[23] UAE: ICT Health Law (Federal Law No. 2 of 2019), Art. 13. Saudi Arabia: Regulation on Personal Data Transfer outside the Kingdom (Version 2.0, August 2024), Arts. 4(2) and 7(1)(b). China: Cybersecurity Law, Art. 39 (Art. 37 before 1 January 2026); PIPL, Arts. 38 and 40; CAC, Provisions on Promoting and Regulating Cross-Border Data Flows (22 March 2024), Art. 7.

[24] US Department of Justice, Preventing Access to U.S. Sensitive Personal Data and Government-Related Data by Countries of Concern or Covered Persons, 28 CFR Part 202 (90 FR 1636, 8 January 2025; effective 8

same export can therefore meet very different gates depending on jurisdiction, and an organisation operating across markets cannot assume that satisfying the European transfer regime satisfies the others.

Layer 3 - Health-Sector Localisation

On top of the general transfer rules, several jurisdictions localise health data specifically. The UAE ICT Health Law requires health data connected to services provided in the UAE to remain in-country unless a resolution of the relevant health authority, issued in coordination with the Ministry of Health and Prevention, permits otherwise. This is a hard localisation requirement rather than a "use the right safeguard" question, although Ministerial Decision No. 51 of 2021 sets out permitted cases, such as treatment abroad, insurance and research, each with conditions attached.[25] At emirate level this is reinforced by health-authority instruments. Abu Dhabi's Healthcare Information and Cyber Security Standard (ADHICS, version 2, 2024) applies to any entity in the emirate that generates, stores, processes or transmits health information, and the Department of Health permits offshore hosting only by exemption; the Dubai Health Authority's Policy for Data and Health Information Protection and Confidentiality (2022) prohibits the transfer of protected health information outside the UAE except within the ICT Health Law exemptions and with prior DHA approval. Saudi Arabia has no health-specific localisation rule. Residency expectations reach health data indirectly: NHIC telehealth guidance states that data storage "should be hosted within Saudi Arabia's geographical boundaries"; participation in the national health-information platforms, which are Saudi-hosted, is mandatory for providers and insurers; and public-sector health entities fall within the government-data localisation rules of the CST cloud regulations. Private-sector transfers are governed by the PDPL and its Transfer Regulation, which add conditions to the cross-border transfer of sensitive data; China and others impose their own localisation and transfer controls.

The contrast with the previous layer is the essential point: Layer 2 governs where data may go, whereas Layer 3 governs where it must stay. Whether a given consumer app is caught turns on its connection to the UAE health sector: a standalone wellness app may fall under the general data protection regime, but once its data is tied to UAE health services or it serves a regulated health or insurance entity, the localisation rule can require the

April 2025), ss. 202.205 (bulk threshold of more than 10,000 US persons for personal health data), 202.241 (definition of personal health data), 202.301 (prohibited data-brokerage transactions) and 202.601 (countries of concern: China, Cuba, Iran, North Korea, Russia and Venezuela). See also the Protecting Americans' Data from Foreign Adversaries Act of 2024, 15 USC 9901 (data brokers).

[25] UAE ICT Health Law (Federal Law No. 2 of 2019), Art. 13; Ministerial Decision No. 51 of 2021; Department of Health Abu Dhabi, Healthcare Information and Cyber Security Standard, version 2 (DOH/SD/ICSO/ADHICS/V2/2024, May 2024), scope statement and s. 6; Dubai Health Authority, Policy for Data and Health Information Protection and Confidentiality (DHA/HRS/HISHD/PP-11, 10 August 2022), ss. 4.15.1 and 4.24.3-4.24.4. Saudi Arabia: PDPL (Royal Decree M/19 of 2021, as amended by Royal Decree M/148 of 2023), Arts. 1, 6(4), 22, 23, 29 and 30; Implementing Regulation (September 2023), Arts. 11, 16, 25, 26 and 32; Regulation on Personal Data Transfer outside the Kingdom (September 2023, Version 2.0 August 2024), Arts. 3, 4 and 7; CST Cloud Computing Service Provisioning Regulations v4 (Decision 506/1445, in force 10 October 2023); NCA Essential Cybersecurity Controls (ECC-2:2024), Appendix C (localisation controls transferred to NDMO); NHIC, Telehealth Application Guidelines, "Data Storage Considerations". See also Latham & Watkins/Practical Law, Data Localization Laws: Overview (Saudi Arabia) (law stated 1 January 2026).

data to remain within the UAE, irrespective of whatever transfer safeguards the operator might otherwise have in place.

Localisation of health data is not solely a Gulf or Asian phenomenon. The European Health Data Space Regulation, adopted in 2025, does not localise health data generally, but it requires health data access bodies, trusted health data holders and their processors to store and process personal electronic health data in the Union when carrying out secondary use operations, subject to an exception for countries covered by an adequacy decision, and it expressly preserves national rules requiring in-Union storage for primary use.[26] It therefore reflects a wider European inclination toward keeping health data close to home, which makes data residency a structural design consideration for any organisation building health-adjacent services for an international user base.

Layer 4 - IoT-Specific Law

If the application pairs with a device (a sleep-tracking ring, a smart scale, a connected wearable), IoT-specific regimes stack on top of everything already discussed. IoT-specific laws are directed at connected devices themselves rather than at personal data in general, and they typically cover matters such as device and network security, registration or type-approval of devices, and access to and portability of the data those devices generate. They vary considerably in kind. The UK PSTI Act and California's SB-327 are relatively narrow device-security laws; the EU Cyber Resilience Act and Data Act focus on product security and on access to the data that connected products generate; Saudi Arabia regulates connected devices through the CST IoT Regulations (2024), which cover licensing, device conformity and spectrum and defer to the general data rules; and the UAE operates a binding IoT regulatory regime, issued by the telecommunications regulator under the Telecommunications Law, that requires registration of IoT service providers and imposes classification-based data storage and handling requirements.[27]

Comparative Overview

The table below maps the four layers across the jurisdictions in scope. It is intended as a high-level orientation to the analysis that follows rather than as a substitute for jurisdiction-specific advice. Cutting across them is an AI-governance overlay, binding in the EU through the AI Act and in China through its algorithmic recommendation, deep synthesis and generative AI measures, and elsewhere expressed largely in voluntary instruments such as the NIST AI Risk Management Framework and the OECD AI Principles, which increasingly links profiling and automated decision-making to the health inferences this paper is concerned with.[28]

---

[26] See: EDPB-EDPS Joint Opinion 03/2022 on the Proposal for a Regulation on the European Health Data Space (12 July 2022), Executive Summary and paras 104-111; Regulation (EU) 2025/327 (European Health Data Space), Arts. 86, 87 and 90 (Art. 87 applies from 26 March 2029: Art. 105).

[27] UAE TDRA (then TRA) IoT Regulatory Policy (22 March 2018) and IoT Regulatory Procedures, version 1 (6 March 2019), issued under the Telecommunications Law (Federal Law by Decree No. 3 of 2003); the Policy text was not accessible at the time of review and its storage requirements are taken from published summaries; UK PSTI Act 2022; EU Cyber Resilience Act (Reg. (EU) 2024/2847); California SB-327 (Cal. Civ. Code 1798.91.04-1798.91.06, operative 1 January 2020, as amended 2022); US IoT Cybersecurity Improvement Act of 2020, Pub. L. 116-207 (federal agencies).

[28] NIST AI Risk Management Framework (AI RMF 1.0, 2023); OECD AI Principles (2019, updated 2024).

Jurisdiction	General privacy law	Health-data-specific instrument	IoT-specific instrument	Transfer & localisation posture
European Union	GDPR	GDPR Art. 9; European Health Data Space (EHDS)	Cyber Resilience Act; Data Act; ETSI EN 303 645 (industry standard)	Chapter V transfers (adequacy / SCCs / BCRs); member-state health hosting rules (e.g. France HDS certification: EEA storage; Germany SGB V s. 393: EU/EEA/adequacy-country cloud only)
United Kingdom	UK GDPR	UK GDPR + ICO guidance (no single health statute)	PSTI Act 2022	UK GDPR transfer regime as amended by the Data (Use and Access) Act 2025 (adequacy regulations / IDTA and Addendum; "not materially lower" data protection test from 5 February 2026)
UAE (federal)	PDPL (Decree-Law 45/2021)	ICT Health Law (Federal Law 2/2019) + Ministerial Decision 51/2021; DoH ADHICS v2 (2024); DHA data and health information protection policy (2022)	TDRA IoT Regulatory Policy (2018) + Procedures (2019)	Health data localised by default (ICT Health Law Art. 13, subject to Ministerial Decision 51/2021 cases); PDPL transfers on adequacy approved by the UAE Data Office (whose functions passed to the Artificial Intelligence and Data Authoririry in June 2026) or Art. 23 safeguards; IoT registration and classification-based storage
ADGM	ADGM Data Protection Regulations 2021	DP Regulations (special categories, s. 7); federal ICT Health Law also applies in free zones	No dedicated IoT law; the TDRA IoT Regulatory Policy applies to IoT service providers operating in the UAE	Transfers on adequacy / appropriate-safeguards basis (ss. 41-42)
DIFC	DIFC Data Protection Law 2020	DP Law (special categories, Art. 11); federal ICT Health Law also applies in free zones	No dedicated IoT law; the TDRA IoT Regulatory Policy applies to IoT service providers operating in the UAE	Transfers via adequacy / appropriate safeguards (Arts. 26-27)
Saudi Arabia	PDPL	PDPL sensitive-data rules (Arts. 1, 6(4), 23); NHIC/SHC health-sector policies (security and exchange, no binding localisation, though NHIC	CST IoT Regulations (2024)	No general localisation of personal, sensitive or health data; PDPL Art. 29 transfer conditions with adequacy / SCCs / BCRs / certification and mandatory risk assessment for continuous or widespread transfers of

Jurisdiction	General privacy law	Health-data-specific instrument	IoT-specific instrument	Transfer & localisation posture
		telehealth guidance expects in-country hosting)		sensitive data; government data localised under CST cloud regulations
Canada	PIPEDA (+ provincial)	Provincial health privacy laws (e.g. Ontario PHIPA)	No dedicated federal IoT law	PIPEDA: transfers permitted with accountability / comparable protection, no approval; Quebec: privacy impact assessment and written agreement before communication outside Quebec; public-sector localisation in Nova Scotia only
United States	Sectoral / state (no comprehensive federal law)	HIPAA (covered entities and business associates); state consumer-health-data laws (Washington, Nevada, Connecticut, Maryland), Washington and Nevada expressly reaching inferred data	California SB-327; federal IoT Cybersecurity Improvement Act (govt)	No federal localisation; transfers of bulk personal health data to countries of concern prohibited or restricted (DOJ rule, 28 CFR Part 202, 2025); data brokers barred from transfers to foreign adversaries (PADFAA 2024)
China	PIPL	PIPL Art. 28 (medical health = sensitive PI); NHC Population Health Information Measures (2014) and Health and Medical Big Data Measures (2018)	No dedicated IoT law; Cybersecurity Law and MLPS classification apply to connected products (GB/T 22239-2019 IoT extension requirements); MIIT IoT security standards	Localisation for CIIOs, large-volume processors, state organs and medical/health institutions; export via CAC security assessment, standard contract or certification by data type and volume

Why the Patchwork Leaves a Gap

Almost none of these instruments was designed for inferred health data generated by consumer apps and devices. IoT laws target device security and registration; health-sector laws target collected clinical data inside formal health systems; transfer rules assume that an organisation already knows it holds health data and is consciously moving it. Inferred health data produced by a "wellness" application falls into the space between these regimes: too consumer-facing for the health-sector framework, not obviously "health" enough to prompt an Article 9 analysis in practice, yet capable of engaging all four layers at once. The gap is less one of legal coverage than of design and recognition. The patchwork therefore leaves inferred consumer health data covered in principle but

unrecognised in practice, which is precisely where this paper is situated.

This is not only the present authors' reading. The EDPB and EDPS have themselves described the regulation of health data as an already multi-layered set of provisions, spanning the GDPR, Member State health laws, the ePrivacy Directive, the Data Governance Act, the Data Act, the AI Act (both then still proposals), and the Medical Devices Regulation, whose interplay needs clarifying.[29] The overlap and the uncertainty they identify are the same conditions that allow inferred health risk to escape governance in practice, which is the subject of Section 5.

4. What Regulators Expect: Core Compliance Mechanisms

Beneath the jurisdictional variation, a set of common expectations can be identified once health inference is in play. They are drawn from three kinds of source: binding provisions that recur across the regimes surveyed (special category or sensitive-data conditions, consent standards, impact assessments, transparency duties and controls on automated decision-making); regulator guidance that applies those provisions to inference, connected products and AI; and court decisions on what counts as health data.

Recognising Inferred Special Category Data

The first expectation is that organisations apply the reveals-or-concerns analysis to the outputs of their systems, not only to their inputs, and recognise that an inference can be health data even where every source field was ordinary.[30] In practice this requires a deliberate step that most data inventories omit: cataloguing not just what is collected, but what can be derived, and treating derived health conclusions as special category data subject to the full regime.

This expectation rests on binding law and its authoritative interpretation rather than on regulatory preference. The Court of Justice has held that processing which allows health information to be revealed falls within Article 9(1) whether or not that information is correct (Meta Platforms)[31] and even where the link to the individual is only probable (Lindenapotheke).[32] In its guidance, the ICO states that profiling which infers health status is processing of special category data "regardless of how confident you are that the inference is correct". Washington's My Health My Data Act and Nevada's consumer health data law make the same point in statute by reaching data "derived or extrapolated" from non-health information.[33]

Privacy by Design and by Default in Connected Products

Regulators expect inference-awareness to be built into product architecture rather than added after deployment. For connected

[29] EDPB-EDPS Joint Opinion 03/2022 on the Proposal for a Regulation on the European Health Data Space (12 July 2022), Executive Summary and paras 8, 17 and 28-29.

[30] ICO, "What is special category data?" (Guide to the UK GDPR).

[31] Case C-252/21 *Meta Platforms and Others v Bundeskartellamt*, paras 68-73.

[32] Case C-21/23 *Lindenapotheke*, paras 83-84.

[33] Article 29 Working Party, Annex on health data in apps and devices (5 February 2015), pp. 3-5; ICO, "What is special category data?" (inferences: "regardless of how confident you are that the inference is correct"); Washington My Health My Data Act, RCW 19.373.010(8)(b)(xiii); Nevada, NRS 603A.430(1)(d).

products this means protective defaults, genuine data minimisation, and restraint in what the system is permitted to derive, rather than default settings that place the burden on the user to opt out of inference they cannot see.

The expectation has a statutory footing in most of the regimes surveyed. The GDPR and UK GDPR require data protection by design and by default (Article 25); the DIFC Data Protection Law (Article 14(3)) and the ADGM Data Protection Regulations (section 23) adopt the same principle by name; the UAE PDPL requires appropriate measures both when determining the means of processing and in default settings (Article 7(2) and (3)); and China's PIPL requires handlers to adopt classified management and technical security measures (Article 51). The Saudi PDPL frames the duty as one of organisational, administrative and technical measures and of impact assessment before sensitive data is processed, rather than as a named design principle. The ICO's consumer IoT guidance applies the principle to connected products directly: manufacturers "must consider data protection and privacy issues upfront at the design stage and throughout the lifecycle" of the product.[34]

Consent Engineering Done Properly

Where consent is the chosen basis, it must meet the legal standard set out below. This is a higher bar than a click-through on general terms can meet, and meeting it requires coordination between legal, product, and design functions rather than legal drafting alone. Consent that does not extend to inference cannot legitimise inference.

Under the GDPR and UK GDPR consent must be freely given, specific, informed and unambiguous, demonstrable, presented separately from other matters and as easy to withdraw as to give, and for special category data it must be explicit. The EDPB has confirmed that consent bundled across purposes is not freely given and that deceptive design patterns undermine it. Comparable standards apply under the DIFC law (Article 12), the UAE PDPL (Article 6: clear, unambiguous and easily withdrawable), the Saudi Implementing Regulation (explicit consent for sensitive data), China's PIPL (separate consent for sensitive information) and, in the United States, the CCPA (agreement obtained through dark patterns "does not constitute consent") and Washington's My Health My Data Act (opt-in consent that cannot be obtained through general terms or deceptive designs, with separate consent for sharing).[35]

Impact Assessments Calibrated to Inferred Data

Data protection impact assessments must evaluate what a system can derive, not only what it collects. A DPIA that catalogues

[34] GDPR and UK GDPR, Art. 25; DIFC Data Protection Law 2020, Art. 14(3)-(4); ADGM Data Protection Regulations 2021, s. 23; UAE PDPL, Art. 7(2)-(3); PIPL, Art. 51; Saudi PDPL, Art. 19 and Implementing Regulation, Arts. 23 and 25; ICO, Guidance for consumer Internet of Things products and services (11 June 2026), "How do we apply a data protection by design and default approach?".

[35] GDPR and UK GDPR, Arts. 4(11), 7 and 9(2)(a), Recitals 32, 42 and 43; EDPB, Guidelines 05/2020 on consent (version 1.1, 4 May 2020), paras 13 and 42-44; EDPB, Guidelines 03/2022 on deceptive design patterns in social media platform interfaces (version 2.0, 14 February 2023); DIFC Data Protection Law 2020, Art. 12; UAE PDPL, Art. 6; Saudi Implementing Regulation, Art. 11(2); PIPL, Art. 29; Cal. Civ. Code 1798.140(h) and (l); RCW 19.373.010(6) and 19.373.030.

collected fields while ignoring inferred outputs misses the very risk that matters most.[36]

The law already prescribes when an assessment is required. Under the GDPR and UK GDPR a DPIA is mandatory where processing is likely to result in a high risk, including the large-scale processing of special category data and the systematic evaluation of personal aspects by automated means on which decisions producing legal or similarly significant effects are based; the Article 29 Working Party's criteria treat sensitive data, evaluation or scoring, the matching of datasets and innovative technology as indicators. The UAE PDPL requires an assessment before using modern technologies that pose a high risk, and in any event for systematic profiling with serious effects or large volumes of sensitive data (Article 21); the Saudi Implementing Regulation makes the processing of sensitive data itself a trigger (Article 25(1)); and China's PIPL requires an impact assessment before sensitive information is processed or used for automated decision-making (Article 55). What none of these provisions does is tell an organisation that its outputs are sensitive data. That is the step this paper adds: an organisation that can recognise where its systems may infer health status should treat that capability as a trigger and assess it, rather than waiting for the inference to be labelled as health data by someone else.[37]

Transparency About Inferences

People should be told what is being concluded about them, not merely what is being collected from them. Transparency that stops at the input layer leaves the most consequential processing invisible, and notices written for a legal audience rather than a human one do not discharge the obligation. The expectation is meaningful comprehension of the inferences a service draws, expressed in terms a user can actually understand.

The obligation is statutory. The GDPR and UK GDPR require information to be provided in a concise, transparent and intelligible form and, where automated decision-making including profiling is used, "meaningful information about the logic involved" and its envisaged consequences; the UAE PDPL gives data subjects a right to information about purposes and automated decisions (Article 13); the Saudi PDPL requires a privacy policy to be available before collection (Article 12); China's PIPL requires clear notice of purposes and categories (Article 17) and transparency and explanation of automated decisions (Article 24); and the CCPA requires notice at collection of the categories of personal information collected, of which inferences are expressly one. The ICO's guidance on explaining AI-assisted decisions sets out how the standard can be met in practice.[38]

[36] GDPR Art. 35; Article 29 Working Party Guidelines on Data Protection Impact Assessment (WP248 rev.01, 4 October 2017), endorsed by the EDPB.

[37] GDPR and UK GDPR, Arts. 35(1) and (3) and 36; Article 29 Working Party, Guidelines on Data Protection Impact Assessment (WP248 rev.01, 4 October 2017), endorsed by the EDPB; ICO, "When do we need to do a DPIA?"; ICO, Guidance for consumer Internet of Things products and services (11 June 2026), "Do we need to do a data protection impact assessment (DPIA)?"; UAE PDPL, Art. 21; Saudi Implementing Regulation, Art. 25(1); PIPL, Art. 55.

[38] GDPR and UK GDPR, Arts. 12(1), 13(2)(f), 14(2)(g) and 15(1)(h) (in the UK GDPR, Art. 13(2)(f) now refers to the safeguards in Art. 22C); UAE PDPL, Art. 13; Saudi PDPL, Art. 12; PIPL, Arts. 17 and 24; Cal. Civ. Code 1798.100(a) and 1798.140(v)(1)(K); ICO and The Alan Turing Institute, Explaining decisions made with AI (first published 2020).

AI and Profiling Governance

Predictive systems that infer health must be governed for fairness, explainability, and lawful basis, with meaningful human oversight wherever the consequences for individuals are significant.

In the EU this is now law for high-risk systems. The AI Act requires a risk management system, data governance including the examination of bias, transparency towards deployers and effective human oversight (Articles 9, 10, 13 and 14); it imposes duties on deployers, including competent human oversight (Article 26); and it requires a fundamental rights impact assessment before deploying systems used for life and health insurance risk assessment and pricing (Article 27). These obligations apply to Annex III systems from 2 December 2027. Independently of the AI Act, GDPR Article 22(3) and UK GDPR Article 22C guarantee human intervention, representations and the right to contest a solely automated decision, and the ICO's guidance on AI and data protection addresses fairness, lawfulness and transparency in AI systems.[39]

Taken together, these mechanisms already exist in the law of most of the jurisdictions surveyed; what varies is how squarely each has been applied to inference. Consent at the point of collection is not treated as sufficient on its own where the processing goes on to derive new information, and organisations are assessed on whether their systems are designed and governed to manage what they derive about people, not merely what those people knowingly hand over.

5. The Governance Gap: Why Inferred Health Risk Escapes Existing Controls

The convergence described in the previous section does not, in practice, close the gap between what the law expects and what organisations actually do. Many organisations recognise special category data in principle, maintain policies that reference Article 9, and conduct impact assessments as a matter of routine, and yet still fail to catch inferred health risk in operation. The reason is that inference happens at points in the system where privacy review rarely reaches. This section sets out the recurring sources of that gap, each of which the framework in Section 6 is designed to address.

The first source is the distance between recognition and detection. An organisation may acknowledge that inferred health data is special category data, but acknowledgment at the level of policy does not translate into detection at the level of systems. Inference typically occurs downstream of where privacy review sits, so the resulting data is never flagged as health data in the first place and never enters the processes that would govern it. The control exists on paper; it simply never engages.

The second source is design-stage blind spots. Health inference is frequently built into a product as a feature (a wellness score, a risk indicator, a personalisation signal) without any data protection impact assessment or

[39] EU AI Act (Reg. (EU) 2024/1689), Arts. 9, 10, 13, 14, 26 and 27, applicable to Annex III systems from 2 December 2027 (Regulation (EU) 2026/1744); GDPR, Art. 22(3); UK GDPR, Art. 22C; ICO, Guidance on AI and data protection (updated 15 March 2023); NIST, AI Risk Management Framework 1.0 (January 2023); OECD, Recommendation of the Council on Artificial Intelligence (2019, revised May 2024).

special category analysis being triggered, because the team building it does not perceive what it is building as health-data processing. By the time the implications are recognised, the inference is embedded in the product and expensive to unwind.

The third source, discussed in Section 2, is profiling and model opacity. Health conclusions often emerge from models that no single team fully owns or understands. The fourth source, also discussed in Section 2, is vendor and ecosystem fragmentation. SDKs, analytics partners, and cloud layers generate inferences outside the controller's direct line of sight.

The fifth source is functional silos. Privacy, AI, product, security, and risk functions each see only part of the picture, and inferred health risk is a problem that lives in the spaces between them. The privacy team may not see the model; the data-science team may not see the legal characterisation; the product team may not see either. No single function is positioned to see the whole, and so the whole is not seen. The cumulative result of these five conditions is that inferred health risk is not so much consciously accepted as quietly missed, which is precisely why it requires a structured, cross-functional response rather than a stronger policy.

6. Practical Governance Framework for Organisations

The following framework sets out seven implementation-ordered steps for embedding inference-aware governance into the design, development, and operation of digital services.

Step 1: Map Where Ordinary Data Can Yield Health Inferences

Identify not only what data is collected but what health-related conclusions the service can or does derive from it, and where in the technical stack those inferences are generated: in the application, in an analytics layer, in a model operated by a vendor, or across a combination of these. This inference map is the factual foundation for every step that follows.

Step 2: Assess Special Category Thresholds

Determine, market by market, which mapped inferences cross into a protected category, applying the reveals-or-concerns test to outputs as well as inputs. Check whether sectoral rules attach on top, such as the UAE ICT Health Law, the Chinese NHC measures or HIPAA.[40] Document the assessment: it establishes the obligations that attach and demonstrates the diligence regulators expect.

Step 3: Conduct Inference-Calibrated Impact Assessments

Assess what the system derives, not only what it collects; a generic DPIA is insufficient. Treat inferred health data pooled or combined at scale as a likely high-risk trigger, and check whether any AI component falls within the high-risk categories of the EU AI Act.[40] Feed the findings back into product and design decisions rather than leaving them in a compliance document.

[40] GDPR, Art. 35; EU AI Act, Art. 6 and Annex III (applicable from 2 December 2027 under Regulation (EU) 2026/1744).

Step 4: Embed Privacy by Design and Clean Consent Engineering

Build in protective defaults, minimisation of both collected and derived data, restraint in what the system is permitted to infer, and consent flows that are honest about inference rather than engineered to extract agreement. Retrofitting these properties after launch is more expensive and less effective.

Step 5: Govern AI and Profiling Systems

Subject any system that infers or predicts health to fairness, explainability, lawful-basis and oversight controls, with a clear owner accountable for the inferences it produces. Where models are operated by vendors, secure these controls contractually and verify them in practice.[41]

Step 6: Coordinate Privacy, AI, Cybersecurity, Product, and Risk Functions

Because inferred health risk lives in the spaces between functions, establish shared ownership across the privacy, data-science, product, security and risk teams, with a forum that makes the whole picture visible to someone accountable for it. The coordination model already used for other cross-cutting risks can usually be extended to this one.

Step 7: Provide Transparency and Data-Subject Controls Over Inferred Data

Tell people what is being inferred about them, not only what is being collected, and provide meaningful access, objection, correction and deletion controls over derived data. The test is whether an ordinary user can understand and exercise these controls in practice.

Step	Key action	Output
1. Map inferences	Identify the health conclusions the service can derive and where in the stack they arise	Inference map
2. Assess thresholds	Apply the reveals-or-concerns test to outputs, market by market; check sectoral rules	Documented classification of inferences by market
3. Impact assessment	Run DPIAs on derived data; check for EU AI Act high-risk status	Inference-calibrated DPIA feeding into design
4. Design and consent	Protective defaults, minimisation, limits on inference, honest consent flows	Privacy-by-design controls and compliant consent
5. AI governance	Fairness, explainability and oversight controls; named owner; vendor contracts	Documented, accountable AI controls

[41] EU AI Act, Arts. 9, 10, 13, 14, 26 and 27; GDPR, Arts. 22 and 28(3) (processor contracts); UK GDPR, Art. 22C; Article 29 Working Party, Guidelines on Automated individual decision-making and Profiling (WP251rev.01, 6 February 2018), endorsed by the EDPB; ICO, Guidance on AI and data protection (2023); NIST AI Risk Management Framework 1.0; OECD AI Principles.

6. Coordination	Shared ownership across privacy, data science, product, security and risk	Cross-functional forum with clear accountability
7. Transparency	Disclose inferences; extend user rights to derived data	Inference-aware notices and user controls

7. Conclusion: Governing Health Data Beyond Healthcare

Health information is now produced outside healthcare as a by-product of ordinary digital services. The threshold question for a calorie tracker, a sleep app, a wellness platform or an insurer's pricing model is therefore not whether it holds medical records, but whether what it derives reveals health status. Under the GDPR and the regimes that follow its approach, the answer will often be yes.

Every regime surveyed treats health-related information as sensitive, but they do so differently. Only some of these regimes address inferred data expressly: the CJEU's case law, the ICO's guidance, Washington's and Nevada's statutes and California's treatment of inferences as personal information. In the others the position follows from the breadth of the definitions rather than from an explicit rule, and it should be analysed as such.

Identifying inferred health data and determining which obligations follow are two different exercises. The first turns on what the data reveals and is broadly common across markets. The second is jurisdiction-specific. Localisation in the strict sense is narrower than often assumed. The UAE ICT Health Law localises health information connected to UAE health services; China localises health data held by medical and health institutions and data held by critical infrastructure operators and large processors; the EHDS will localise secondary-use processing within the Union from 2029; France's and Germany's hosting rules confine health data to the EEA or adequate countries; Saudi Arabia localises government data and reaches health data through guidance and national platforms rather than a statute; and the UK, Canada (outside Nova Scotia's public sector) and the United States impose no general health-data residency requirement. Satisfying a transfer mechanism does not satisfy a separate sectoral rule.

For organisations building or operating such services, the proportionate priorities are these. First, map inferences as well as collected data and classify the outputs against each applicable regime; Second, settle the lawful basis and any special category condition for the act of inference itself, not only for collection. Third, run the impact assessment on what the system derives, and treat any component that prices insurance, screens employees or recognises emotion as a candidate for the EU AI Act's high-risk regime. Fourth, design hosting and transfer architecture around the strictest localisation rule actually engaged, rather than assuming that one transfer mechanism suffices everywhere. Fifth, give the privacy, product, data-science and security functions a shared view of what the system infers, with a named owner.

White Label Consultancy works with organisations on these questions, helping them identify inferred health data and put in place the governance, product and transfer controls that the applicable regimes require.

Authors



Nicholai Pfeiffer

Managing Partner
White Label Consultancy

np@whitelabelconsultancy.com



Dr. Magdalena Góralczyk

Data Protection Partner
White Label Consultancy

mg@whitelabelconsultancy.com



Darya Ghasemzadeh

Associate
White Label Consultancy

dgh@whitelabelconsultancy.com



Nauani Benevides

Consultant
White Label Consultancy

nbe@whitelabelconsultancy.com

Contact

hello@whitelabelconsultancy.com
+45 71 74 74 54

NORWAY

Main Office Oslo

+47 4141 2168

Fjordalléen 16
0250 Oslo

DENMARK

Copenhagen Office

+45 71 74 74 65

Dampfærgevej 27
2100 København Ø

POLAND

Warsaw Office

+48 515 07 99 77

Ul. Marszałkowska 58/15
00-545 Warszawa

UNITED ARAB EMIRATES

Dubai Office

+971 50 4536616

Dubai World Trade Center
Dubai

WHITELABELCONSULTANCY.COM