Securys®
Global Data Privacy Experts

Helping Tech Companies Navigate Data & AI Privacy

Contents

Executive Summary

Health tech has been a growing sector for many years now, and the advent of AI is only adding fuel to the fire – just see the recent launch of [ChatGPT Health](#) by OpenAI, in response to the fact that health queries are one of the most popular uses of its platform.

But this opportunity comes with lots of new risks. Any health technology necessarily involves access to highly sensitive personal data, often across multiple jurisdictions, and organisations building or deploying it face complex patchworks of laws, liabilities and obligations that can differ substantially from market to market and region to region.

In just the last two years we've seen regulations such as the EU AI Act and India's Digital Personal Data Protection (DPDP) Act, along with substantial UK data reform and many new US state-level health privacy laws added to more established compliance regimes like the GDPR and HIPAA. This white paper outlines the key principles and pitfalls involved in the safe handling of health data, and serves as a practical, risk-based guide to successfully navigate them.

At Securys our goal is to help health tech scale-ups mitigate risks, meet varying data/AI privacy laws, and turn robust data compliance into a competitive advantage. Because proactive governance is not just about avoiding fines; it's about transforming your data from a risk into an asset and using it to drive sustainable and trusted global growth for the benefit of your company and its customers.

Introduction: A Growing Opportunity, A Rising Risk

The global health tech and digital health sector continues to experience explosive growth. Market projections show an increase from **US$199 billion in 2025 to over US$573 billion by 2030**[1], fueled by innovations in AI-driven apps, wearables, and remote care. This growth reflects how patients and providers are embracing technology for better health outcomes and operational efficiency. However, alongside this opportunity is a heightened focus on data privacy and security. In a 2024 survey of healthcare tech leaders, **"data privacy and security concerns" were cited among the top challenges** facing the sector[2]. High-profile data breaches and the misuse of sensitive health information – such as the [Change healthcare ransomware attack](#) that affected nearly 200m users and created costs in the billions - have put regulators and consumers on alert.

As a result, regulatory scrutiny over data handling, especially over the handling of sensitive data such as health data – broken out into a special category by most data protection legislation – is intensifying worldwide. Authorities are issuing larger fines and demanding demonstrable compliance. Under the EU's GDPR, fines for data protection violations can reach up to 4% of global annual revenue or €20 million (whichever is greater). Similar penalty regimes exist or are emerging in other jurisdictions. In 2025 the UK strengthened its enforcement under the new Data Use and Access Act by aligning e-privacy (cookie) violation fines with GDPR levels (up to £17.5 million or 4% of global turnover)[3]. In 2025, for example, NHS software provider [Advanced was fined £3m](#) by the ICO over security failings that led to a ransomware attack on the NHS.

The cost of getting Privacy wrong is, therefore, steep. The **average cost of a data breach in 2025 reached $4.44 million globally and a record $10.22 million in the US**, the highest of any country[4]. In the UK, the average breach costs several million pounds, and the healthcare industry remains the costliest sector for breaches (over $7 million on average)[5]. For a scaling health tech company, an incident of that magnitude could be devastating diverting millions away from R&D and - perhaps even more damagingly – tarnishing hard-won user trust. In severe cases, regulators can force suspension of a company's operations altogether.

On the positive side, investing in strong data privacy practices yields business benefits. Health tech businesses that are proactive about privacy can **increase brand value, build consumer trust, and sharpen their competitive edge**. When users know their most private data is handled responsibly, they are more likely to adopt and remain loyal to a platform. Likewise, enterprise clients and partners – especially those who themselves are heavily regulated - prefer vendors who meet high privacy and security standards. In short, good privacy is good business, especially in healthcare, where trust is paramount.

> **Myth 1 – "Popular AI tools are automatically safe for health data."**
>
> *Many assume they can freely use cutting-edge AI like ChatGPT with patient or consumer health information.* **Reality:** Consumer-grade ChatGPT (even the new ChatGPT Health beta) is **not** HIPAA-compliant and **should not be used with protected health information (PHI)**[6]. OpenAI offers to sign Business Associate Agreements (BAAs) only for certain enterprise/API customers – the standard ChatGPT service for individuals explicitly **lacks** a HIPAA BAA. In fact, OpenAI has *not launched* ChatGPT Health in the EU/UK as of early 2026, likely due to unresolved GDPR concerns such as required DPIAs. Always vet any AI tool's privacy posture before integrating it with sensitive health data.

Context: What's changed recently?

Data and AI privacy laws are evolving rapidly across key markets, raising the bar for compliance.
Notable developments since 2024 include:

- **EU:** The **EU AI Act** was passed in 2025, with phased obligations – outright bans on
  certain AI practices by Feb 2025, new rules for general-purpose AI by Aug 2025, and
  requirements for high-risk AI systems coming into force in 2026–2027. These add to
  the EU GDPR's existing strict rules on health data (a special category requiring explicit
  consent or equivalent safeguards) and its international data transfer rules.

- **UK:** The **UK Data (Use and Access) Act 2025** amended UK GDPR and related laws. It
  increased fines for electronic privacy (PECR) violations to GDPR levels[7], introduced
  "recognised legitimate interests" that simplify certain processing, and reformed rules on
  international transfers and automated decision-making transparency[8]. These changes
  aim to make UK data law more innovation-friendly while keeping strong protections.

- **US:** With no single federal privacy law covering all health data, the US. landscape is led
  by sectoral laws and state legislation. **HIPAA** remains critical for healthcare providers
  (covering Protected Health Information), but new state laws expand privacy to
  consumer health apps. For example, **Washington's My Health My Data Act (WMHMDA)**
  came into force, broadly defining "consumer health data" and requiring consent for its
  collection/sale (and even prohibiting certain geo-location targeting near medical
  facilities). **California** has also been active: in 2025 the California Privacy Protection
  Agency finalised regulations requiring **privacy risk assessments and annual
  cybersecurity audits** for larger businesses, and set rules for Automated Decision-Making
  Technology (ADMT) with compliance dates in 2026–2027[9].

- **India:** India's new **Digital Personal Data Protection (DPDP) Act 2023** and draft Rules
  (2025) established a modern privacy framework. Companies face *risk-based obligations*
  depending on scale – "significant data fiduciaries" (large data processors) must do
  additional compliance like Data Protection Impact Assessments. Data exports will
  follow a government-defined "negative list" of banned destinations, meaning cross-
  border transfers are allowed unless they are to a blacklisted country. Consent and
  breach notification requirements are being operationalised under these rules; while the

GDPR has six lawful bases for processing personal data, the DPDPA relies heavily on consent as the primary basis for processing personal data. There is a further set of certain legitimate uses (which include legal obligation, vital interest and employment purposes) but these do not map readily to GDPR's lawful bases. Consent management is likely to be a crucial and challenging task for most companies, therefore, considering the large number of data principals that even smaller businesses cater to.

- **Middle East:** Several jurisdictions have introduced or updated data laws. The **United Arab Emirates** implemented a federal Personal Data Protection Law (PDPL) in 2022 largely aligned with global norms and maintains additional sectoral rules such as the UAE's **ICT in Health Fields law**, which imposes strict controls (including local storage requirements) on health data. The UAE PDPL limits cross-border transfers to countries deemed "adequate" or otherwise requires contractual and organisational safeguards[10]. Meanwhile, Dubai's and Abu Dhabi's financial free zones (DIFC, ADGM) have their own data protection regimes (with DIFC requiring DPIAs for high-risk processing, similar to GDPR). Across the Middle East, data exports and data localization (especially for health or sensitive data) are key compliance themes.

- **Caribbean:** A wave of new privacy legislation is sweeping through the Caribbean, aiming to match international standards. **Barbados** enacted a comprehensive Data Protection Act in 2021 explicitly modeled on the EU GDPR[11]. **Jamaica's** Data Protection Act took effect in December 2023 and includes GDPR-like features such as extraterritorial scope (applying to foreign companies handling Jamaican data) and a 72-hour breach notification rule[12]. These laws designate health information as sensitive data, require appropriate consent and safeguards, and restrict cross-border transfers to countries with adequate protection or via binding agreements[13]. Other nations (Trinidad & Tobago, the Bahamas, etc.) are updating older frameworks to align with modern expectations. Health tech firms cannot afford to ignore the Caribbean, especially if their apps attract users there or process data across those borders.

- **Other Global Changes:** Canada's privacy regime is tightening at the provincial level – **Quebec's Law 25** came into force, mandating Privacy Impact Assessments for many projects and upping fines, even as federal reform (Bill C-27) stalls. **Brazil** is enforcing its LGPD with new rules: international transfers will require Standard Contractual Clauses by August 2025 and regulators are prioritizing issues like biometric and children's data. **Australia** passed initial Privacy Act amendments in late 2024, introducing tougher penalties (and even a new criminal offense for malicious disclosure of data) and bolstering transparency around automated decisions. **China** rolled out a certification scheme for cross-border personal data transfers (effective January 2026) to complement its stringent security assessments and standard contract requirements under the PIPL. In short, around the world, the direction is clear: more requirements, more enforcement, and more expectations that companies **demonstrate** accountability in how they handle personal and health data.

Against this backdrop, what must health tech scale-ups focus on? The following sections outline **six key Privacy and AI Governance challenges** that these companies face, along with practical steps to address them. We then discuss the benefits of getting it right – and how Securys can partner with you to build and execute a robust privacy and AI governance programme.

Key Privacy and AI Governance Challenges for Health Tech Scale-ups

Health tech companies often begin with a laser focus on product and growth, but as they mature, **privacy and compliance challenges** quickly multiply. Based on industry research and our experience, six key challenge areas demand attention:

1. Navigating a Complex Regulatory Environment
2. Understanding Your Obligations Under Each Law
3. Maintaining Transparency & Accountability at Scale
4. Consent Collection and Management
5. Ensuring Supply Chain & Third-Party Robustness
6. Governing the Use of AI Ethically and Legally

These challenges are interconnected. Successfully navigating them will help you protect data, preserve user trust, and expand into new markets with confidence. Let's examine each in detail.

1. Navigating a Complex Regulatory Environment

Perhaps the single biggest misconception among emerging health tech companies is underestimating the scope of privacy laws that apply to them. Many assume that compliance is only about the health-specific regulations they know (like HIPAA in the US. or the Medical Device Regulation in the EU) or the law of the country where they are based. In reality **data protection and privacy laws have extraterritorial reach** – they can apply if you **offer apps or wearables in a jurisdiction or collect data from its residents**, even if you have no physical presence there.

For example, imagine a health app developer headquartered in Country X who makes a wellness app available globally via app stores. They might be fully aware of their obligations under Country X's laws, but **unaware that they simultaneously have obligations under the privacy laws of every other jurisdiction where their app is downloaded**. A small team in one country could inadvertently be violating, Brazil's LGPD or the EU's GDPR simply because users in those regions signed up.

Another example: a mood-tracking app that collects information on users' daily emotional states may not seem "medical," but mood data **reveals mental health status, which is considered sensitive personal data under laws like the GDPR** – triggering stricter data protection requirements. We have seen startups surprised to learn that their benign-sounding data (e.g. step counts, heart rate, mood logs) are classified as protected health data in many jurisdictions.

The first step is awareness. As a health tech provider you should ask yourself:

- **Can you list every country where your device or app is available for sale or download?**
- **Do you know where all your users, clients, or patients are located globally?** (Your user base may span more regions than you think.)
- **For each of those countries, can you name the relevant data protection or privacy law that applies?** (e.g. GDPR in Europe, CCPA in California, DPDP Act in India, etc.)
- **Under each law, what types of data you handle are considered "personal data" or even more restricted "sensitive" personal data?** (This is crucial – many laws single out health/biometric/genetic data for heightened protection.)

If you can't easily answer the above, it's a sign to deepen your regulatory mapping. A simple **regulatory checklist** like this can reveal gaps in your compliance coverage.

Being global means coping with *fragmentation* in definitions and rules. What counts as "health data" or sensitive personal data differs: the EU GDPR defines any data about health (including inferred data like wellness app readings) as "special category" personal data, while in the US., new laws like Washington's WMHMDA define "consumer health data" very broadly to include any information that could link to past, present, or future physical or mental health status (even if not collected in a clinical context). Some jurisdictions treat precise location or genetic data as sensitive; others may not. These definitional differences affect **what legal basis you need to process the data, whether you need explicit consent,** and so forth.

In addition, multiple regimes have a scope that stretches well beyond their own geographical borders. The GDPR and many GDPR-inspired laws (like those in Barbados or Jamaica) have extraterritorial clauses, if you're offering goods/services to people in their country or monitoring their behavior, or offering services to citizens of their country while they are abroad, their law likely applies[14]. Brazil's LGPD and India's DPDP Act also cast a wide net over foreign companies processing local data even inside their borders. Understanding these jurisdictional gotchas is essential to avoid "surprise" liabilities.

Finally, consider data residency and localization requirements. Certain countries mandate that health data be stored locally or impose heavy conditions for exporting it. For instance, **UAE's health data law** prohibits the storing or processing of health records outside the UAE without permission[15], and **China's regulations** require a government security assessment before transferring large volumes of personal or sensitive data abroad. Ignoring these precepts could not only draw fines but also get your app blocked from the country by the relevant authorities.

Practical Tip: Map your data flows against a world map of laws. Identify where your users are and which laws are engaged. This might result in a matrix of obligations – but it will highlight common themes (as we address in the next section) that you can tackle in a unified way. Engaging privacy counsel or specialists in your key markets (EU, US, UK, India, etc.) can help clarify your exposure. As you plan market expansion, *bake in regulatory considerations from the start*. It's much easier to incorporate compliance (e.g. add proper consent flows or storage choices) while designing the product than to retrofit them under regulatory pressure later on.

2. Understanding Your Regulatory Obligations

Every privacy law imposes a slightly different set of duties on organisations, but they share core **principles**. As a health tech data controller or processor, you'll need to comply with requirements around: **lawful bases for processing, user notice/transparency, data minimization, purpose limitation, data quality, security safeguards,** and handling of **data subject rights**. Understanding these obligations in each jurisdiction is crucial to determining your overall risk exposure and operational priorities.

Some critical obligations to evaluate in your key markets:

- Legal Basis & Consent: On what grounds are you processing health data? GDPR (EU/UK) typically requires **explicit consent** from users to process sensitive health data, unless another narrow exception applies. Other regimes (like India's DPDP Act) also put consent at the center for personal data use. By contrast, HIPAA in the US. allows data use by providers for treatment/payment/operations without patient consent, but it tightly restricts disclosures to third parties and uses for marketing. Ensure you identify an appropriate legal basis under each law (consent, contract necessity, legitimate interests, etc., as available) and that your user terms and notices reflect that. For sensitive data, consent or explicit authorisation is often the safest route globally.

- Transparency (Privacy Notices): Virtually all laws require telling individuals what data you collect, why and how it will be used, and with whom it will be shared. GDPR sets a

high bar for detailed privacy notices; India's law likewise demands clear notice at collection; in China you have to break out every process that uses personal data separately, with the data used and justifications listed for each one. If you use AI algorithms on user data, some laws (e.g. the EU, and soon the UK) also mandate explaining the logic of any automated decisions made in plain language[16]. Wherever you're operating, you should make sure your online and in-app privacy notices are comprehensive and accessible. Keeping privacy notices up-to-date in a fast-evolving product is an ongoing challenge, but non-compliance can lead to penalties.

- **Data Protection Impact Assessments (DPIAs):** Many jurisdictions now require a form of privacy risk assessment for high-risk processing (which health data processing often is). The **EU GDPR mandates DPIAs** for processing likely to result in high risk to individual rights – processing health data at scale, or using AI on sensitive data, both typically qualify. The UK mirrors this requirement. **Quebec's Law 25** in Canada makes PIAs mandatory for a wide range of new projects involving personal information. California's upcoming rules will force larger tech companies to perform **privacy risk assessments** for certain activities by 2026. Even where not explicitly required by law (e.g. in some APAC countries), doing DPIAs is considered best practice and is looked upon favorably by regulators. These assessments help you identify and mitigate risks proactively – e.g. uncover if your new AI diagnostic feature could adversely affect users or if your data sharing with a research partner has unchecked risks. **Failure to conduct a required DPIA can lead to fines** (under GDPR regulators have fined companies for ignoring this obligation). More importantly, skipping risk assessments means you may overlook serious issues until it's too late.

- **Cross-Border Data Transfer Compliance:** Health tech companies often rely on cloud providers or need to send data back to a central server from users around the world. Most privacy laws restrict **international data transfers** – i.e. moving personal data out of the country/region – unless certain conditions are met. The EU/UK require that the destination country has an "adequate" privacy regime or that you implement approved safeguards like **Standard Contractual Clauses (SCCs)** (or the UK IDTA). The US. to EU data flows now have an option via the new EU-US Data Privacy Framework (for certified companies), but otherwise SCCs are needed; you will also need to carry out documented **Transfer Impact Assessments** (TIAs). Other countries have their own rules: India plans a "negative list" of banned destinations (implying other transfers are allowed by default); China requires government security assessments or certifications for large transfers, and explicit consent must be obtained from all individuals to transfer their data internationally at all; Brazil mandates SCCs or similar agreements by August 2025. Laws in the **Caribbean** (e.g. Barbados, Jamaica) generally require that the receiving country has adequate protection or that you obtain consent or contracts for transfers[17]. In practical terms, if you're using a centralised database or cloud hosted in another country, you likely need to implement these legal mechanisms (SCCs, etc.) and document them. Neglecting transfer rules can result in major fines (the largest GDPR fine to date – $1.3 billion against Meta – was for unlawful EU->US transfers[18]).

- **Data Security & Breach Response:** All jurisdictions insist on appropriate **security measures** to protect personal data (though specifics vary). Health data, being highly sensitive, typically calls for encryption, strict access controls, and robust cybersecurity practices (potentially aligning with standards like ISO 27001 or SOC 2 for best practice). Beyond preventing breaches, you must also prepare for the worst: **incident response plans** and **breach notification procedures**. GDPR requires notifying the regulator within

72 hours of a significant personal data breach. US state laws have varying notification timelines (usually within 30–60 days to individuals, and HIPAA says no later than 60 days to affected individuals and the OCR, with additional rules if over 500 people are affected). Australia's Notifiable Data Breaches scheme and many others also require prompt notification to authorities and victims in cases of serious breaches. Failure to report in time can compound your legal troubles. Companies therefore need to know the rules in each market and have a **breach response procedure** in place so that if (or when) an incident occurs, they can respond swiftly and in compliance with each region's requirements.

- **Data Subject Rights Management:** Privacy laws grant individuals rights over their data – the right to access it, correct it, delete it, object to certain uses, etc. The GDPR pioneered broad rights (access, rectification, erasure, portability, objection, restriction, and not to be subject to purely automated decisions). Many of these rights are mirrored in other laws: for instance, **California** and other U.S. states give rights to access, delete, and opt-out of sale/sharing of personal info; India's law will grant rights to access, correction, and grievance redressal; Jamaican law includes rights similar to GDPR like access and correction. Health tech firms must have processes (often via support or automated tools) to handle such **Data Subject Requests** efficiently and within legal timeframes (the GDPR gives you one month to respond). If an EU user requests a copy of all their wellness data or deletion of their account info, can you easily fulfill that kind of request? If a UK user objects to you using their data for research, do you have a mechanism to handle that? Setting up a **DSAR (Data Subject Access Request) workflow** early on is wise – it only gets harder as you accumulate more data and users.

Staying on top of these obligations is admittedly challenging, especially with differences across jurisdictions. But mapping out the commonalities can help. For example, *nearly everywhere*: have a lawful basis, get consent for sensitive data, protect data well, conduct risk assessments, be ready for breaches, and respect user rights. The differences tend to be in the details (e.g. 72 hours vs. 7 days for breach notice; explicit vs. implicit consent in some contexts; which types of processing need a DPIA). We recommend creating an **obligations matrix** for your company that lists key requirements in each major jurisdiction of operation – that way you can design controls to meet the strictest requirement and know where to tweak policy for local nuances.

Finally, don't forget that **non-compliance has real consequences** beyond fines: regulators can issue enforcement notices that halt certain processing activities (imagine being ordered to stop processing EU data while you fix issues – effectively cutting off that market). They can also mandate compensation to individuals or cause class-action lawsuits. And losing user trust may manifest as users leaving your platform if word gets out that you mishandled data. Thus, understanding and meeting your obligations is not just a legal checkbox, but fundamental to sustaining your business.

3. Maintaining Transparency & Accountability at Scale

A hallmark of successful scale-ups is rapid growth – more users, more data, more features, and expansion into new territories. However, as a health tech organisation grows, **keeping track of data flows and maintaining transparency** becomes exponentially harder. Startups often operate in "build fast" mode, which can lead to ad-hoc data practices that don't scale well under regulatory scrutiny.

When you have ten employees and a few hundred users, it's feasible to know where all personal data is stored and who has access to what. But once you've grown to tens of thousands of users and integrate with numerous third-party services... do you **still have an accurate map of**

what data you collect, where it's stored, and how it moves through your systems? Many organisations discover at this point that they do not. This lack of visibility is risky. If you don't know your own data flows, you can't be transparent to users or regulators about them.

A common issue is **insufficient internal expertise and training**. Health tech firms often lack dedicated privacy officers or IT security staff in the early stages of their growth. As a result, employees may unwittingly introduce vulnerabilities – e.g. using unsecured devices for work, mishandling sensitive emails, setting weak passwords, or falling for phishing scams. (In the UK, an astonishing *90% of businesses* identified phishing attacks as the most common cyber threat they faced. Human error remains one of the top causes of data breaches globally[19] - far more prevalent and damaging that more technically-led cyber attacks. Scaling up without raising privacy and security awareness internally is a recipe for incidents.

Moreover, once you operate in multiple jurisdictions, **incident response obligations** become a complex web. If a breach occurs, you may need to notify different regulators (and possibly impacted individuals) on different timelines with different content requirements. For instance, under GDPR you have 72 hours to notify the Data Protection Authority; under Australia's regime, you must notify individuals "as soon as practicable" if a breach likely causes serious harm; under some US. state laws, you notify the state attorney general for large incidents affecting residents. This puts pressure on organisations to have **robust breach detection and response systems**. You can't afford to discover a breach a month after it has happened – by then, you might already be in violation of a 72-hour rule.

When these challenges (poor data mapping, low staff awareness, multi-jurisdiction incident rules) are not addressed, the consequences include **operational delays and compliance failures**. If you're unclear on what data is collected and where, it becomes extremely hard to respond to a data breach or a user's request. We've seen companies scramble during an incident, asking basic questions: "Which database was that data in? Was it encrypted? Who do we need to inform?" Such delays can mean missing notification deadlines, which can lead regulators to levy additional fines or take stricter action. Not to mention, a disorganised breach response **undermines user trust** – customers perceive (rightly) that the company doesn't have its act together in protecting their data.

To maintain accountability at scale, **processes and documentation are key**. Implement a regular data mapping exercise: update internal business processes (usually held, for anything connected with personal data, in a dedicated record of processing activities or RoPA) whenever you launch a new feature or onboard a new vendor. Build a *culture of transparency* – some companies, for example, have internal "data champions" in each of their teams who ensure new projects a through a privacy review or DPIA and that privacy notices are updated. Conduct periodic **privacy training** for all employees (and more advanced training for engineers, product managers, and customer support who regularly deal with sensitive data). Importantly, have an **Incident Response Plan** that is tailored to privacy/security incidents. This plan should include: how to escalate a potential breach internally (so it reaches the privacy/compliance team quickly), a predefined breach response team (including legal, technical, PR, management), communication templates for notifying users and authorities, and a playbook for containing and investigating the incident. Test this plan with drills. If you operate across borders, your plan needs a handy table of breach notification contacts and timelines across each jurisdiction. When every minute counts, being prepared can be the difference between a contained incident and a regulatory nightmare.

In summary, scaling responsibly means **institutionalising privacy and security**. It's moving from ad-hoc efforts to repeatable processes – much like how a growing company formalises its HR or finance practices. By doing so, you minimise the chance of serious compliance slip-ups, and you **strengthen your position to earn trust** from users, partners, and regulators alike.

4. Consent Collection and Management

Consent is a cornerstone of data privacy in healthcare contexts. Health tech organisations often handle data that by law or ethics requires consent – especially when it involves health information or when personal data is used for secondary purposes like marketing. However, obtaining and managing user consent is trickier than it sounds, and many companies get it wrong.

Under the GDPR (and many similar laws), the default lawful basis for processing **sensitive personal data** (which includes health data) is **explicit consent of the data subject** unless another exception applies. This means users must knowingly agree to the specific use of their health data. Likewise, if a health app wants to show personalised ads or share data with third-party partners for research or product improvement, consent is often required (unless anonymization is in play, which has its own strict criteria).

A challenge arises because health tech apps often serve dual purposes: one primary (e.g. wellness tracking) and others secondary (e.g. monetising via ads or partnering with insurance for discounts). We see some startups attempt to cover all these uses in one go – burying consent in a general *Terms and Conditions* or assuming that a user's acceptance of terms implies consent to data processing. **This is a mistake.** In jurisdictions like the EU, **bundled consent** (forcing users to agree to data processing as a condition of using the service, when not strictly necessary) is not valid. Consent must be **freely given, specific, informed, and unambiguous**. Hiding a consent clause in legal fine print or pre-ticking a box for the user violates these principles.

Our research found some health tech apps that **failed to expressly ask for user consent** before processing health information, or only mentioned their data use in the privacy policy that users rarely read. Others take an "all-in-one" consent approach during sign-up – essentially an ultimatum: *"By creating an account, you agree we can use your data for anything in our privacy policy."* This leaves users in the dark about what they're really agreeing to, leaving them feel a lack of transparency and control (leading to mistrust or backlash), and it doesn't meet regulatory standards in many places, leaving the organisation exposed to legal challenges or enforcement actions.

Health data adds another layer of complexity: sometimes **consent is hard to obtain in a user-friendly way**. If an app continuously collects sensor data (heart rate, sleep patterns), asking a user to tap "I consent" every time would be obnoxious. Yet one-time blanket consent isn't good either. The solution often lies in good UX: obtaining broad consent at onboarding for necessary processing, then using *just-in-time notices* and granular controls for additional uses. For example, a fitness app might get initial consent to process health metrics for core functionality, but later, if it wants to share a user's step count with a wellness brand for a reward programme, it should pop up a clear **opt-in choice** for that specific sharing. Granular consent options (separate toggles for say, "Allow my data to be used to personalise ads" in addition to "Allow my anonymised data to be used for medical research") both help users feel in control and consulted, and help keep you compliant.

Different jurisdictions have varying rules around consent. Europe and many countries require explicit opt-in for sensitive data and direct marketing. The **United States (outside of HIPAA)** historically was more "opt-out" based for general personal data use, but things are changing – e.g. several states now mandate opt-in consent to sell sensitive personal data, and the WMHMDA in Washington requires opt-in consent to collect consumer health data in the first place for certain entities. **India's DPDP** Act emphasises consent (with a requirement for clear, plain language consent requests and an easy way to withdraw consent). **Brazil's LGPD** similarly demands specific consent for the collection and use of sensitive data.

Another area that needs attention is **children's data** – if your health tech product might be used by minors (e.g. a smart fitness band for kids or a mental health app used by teenagers), note that laws like the GDPR require parental consent for under-13 (under-16 in some EU countries) in most cases. The UK's Age-Appropriate Design Code and similar "children's codes" in other countries also impose stricter requirements on transparency and consent for younger users.

To handle consent properly:

- Design a **consent interface** that is clear and not misleading. Avoid pre-ticked boxes or vague language. Use simple descriptions for what the user is agreeing to.

- **Document the consents** you obtain – which user gave what consent and when – because you may need to prove this, if challenged. Many privacy laws put the onus on the company to demonstrate that valid consent was obtained.

- Provide a **mechanism to withdraw consent** easily. GDPR and others explicitly say users should be able to change their mind as easily as they gave consent. If a user turns off data sharing, your systems should honour that and stop processing the data for that purpose.

- Review your user journey through a regulatory lens: Are you asking for too much upfront? Often, it's better to only ask for what's needed for core service at sign-up, and defer other optional consents to later when you can better explain the value to the user (and they have a context for the request).

Getting consent right pays dividends in user trust. When users feel in control – e.g. they know *"I can use this app even if I decline marketing cookies or data sharing, I just won't get personalised ads"* – they are more comfortable engaging with your product. On the flip side, if they discover you were doing something with their health data that they weren't clearly told about, it can lead to upset and disappointment and the speedy deletion of the app- and maybe even a complaint to the regulator or a lawsuit.

5. Ensuring Supply Chain & Third-Party Robustness

No tech company is an island; health tech firms rely on a **supply chain of third-party services** from cloud hosting, data storage, analytics tools and payment processors to wearable device manufacturers, AI model providers and more. But every third party that touches your data or systems is a potential source of risk. In fact, some of the largest healthcare data breaches have stemmed from vulnerabilities in the supply chain (for example, a vendor's misconfigured server or a compromised API connection – just look at the catastrophic WannaCry ransomware attack had on the UK NHS in 2017, when the cryptoworm exploited API connections between computers that had not been kept up to date).

When you entrust sensitive personal data to a vendor, **you are still ultimately responsible for its protection** in the eyes of regulators and users. GDPR explicitly holds controllers liable for the compliance of their processors (and requires strict due diligence and contracts). Even under HIPAA, if a business associate (vendor) breaches PHI, the covered entity faces enforcement as well if they didn't have proper safeguards/agreements. In the EU, lack of proper **data processing agreements** or failure to vet processors is one of the most common compliance failures – it's noted that insufficient processor safeguards rank among the top causes of regulatory fines.

Key risks in the supply chain include:

- **Unauthorised access or data leaks via vendors:** If your cloud storage provider doesn't encrypt data or your CRM SaaS has a bug, your user data could be exposed. Any weakness in their security is effectively a weakness in yours.

- **Sub-processor sprawl:** Perhaps you have a contract with a reputable vendor, but that vendor uses further sub-processors (e.g. subcontractors, or it outsources support to another firm). Do you know about those, and are they held to the same standards? It's common for startups to sign up to a service without realizing data might flow to the service's overseas affiliates or subcontractors.

- **Software Development Kits (SDKs) and tracking tools:** Many mobile health apps incorporate third-party SDKs (for analytics, crash reporting, social media integration, etc.). Some of these SDKs have been caught collecting more data than they should, or sending data to unexpected parties. For example, an innocuous library might be transmitting user device info to an ad network unbeknownst to you. This creates hidden privacy issues. - **International transfers via vendors:** Your company might be based in one country, but your vendors might store data on servers around the world. If a vendor moves EU personal data to the US and you didn't put SCCs in place, that's a violation. Similarly, if an Indian law prohibits sending data to X country and your vendor does it, you both are in breach.

- **Incident response and accountability:** If a vendor suffers a breach, will they even tell you in time for you to meet your notification obligations? Your contract should require it, but smaller vendors might not have mature incident detection, so you could find out long after the fact.

All these risks mean that **vetting and managing suppliers is as important as securing your own systems**. Health tech organisations should implement a **vendor risk management programme** that at minimum:

- **Inventories** all third-party providers and tools in use.

- **Assesses** the data they handle and the criticality of the services they provide.

- **Checks contracts,** Data Processing Agreements (DPAs) or business associate agreements that impose security, confidentiality, and cooperation duties have been put in place and that they include provisions such as: the vendor must implement specific security controls, must notify you within X hours of any breach, must not sub-contract without approval, must assist in DPIAs or audits, and will delete/return data upon termination.

- **Conducts privacy/security due diligence** before onboarding a vendor, and periodically (e.g. annually or when renewal comes up). This might involve sending them a questionnaire or reviewing their compliance attestations (ISO certs, SOC2 report, etc.). If a vendor will be handling extremely sensitive data, you might even consider a third-party audit or penetration test of their platform (or choose a vendor that has publicly undergone such audits).

- **Monitors and limits data sharing by** following the principle of least privilege when integrating vendors. Only send the minimum necessary data. For instance, if using a cloud email service to send appointment reminders, do you need to include sensitive medical details in that data transfer? Maybe not – minimise what passes through third parties.

- **Stays aware of fourth-party risk by** asking vendors to disclose their sub-processors. Many big cloud companies list their sub-processors online. You may need to approve these or at least know who else is in the chain.

When supply chain weaknesses aren't addressed, the fallout can be severe. Consider a scenario: your analytics provider suffers a cyberattack and hackers siphon user health data – regulators will come knocking at *your* door, since you chose that provider. Users will blame your app, not the white-label vendor behind it. Additionally, we've seen cases of **joint liability** where both the primary company and the vendor were fined by EU authorities because both failed in their duties to protect data. And beyond legal impact, a breach via a third party still erodes trust in your brand. That's why due diligence and robust vendor management are truly part of your core privacy programme, not an afterthought.

Finally, remember that supply chain risk includes not just digital service providers but any partners you share data with. For example, if you share pseudonymised health data with a university for research, you need an agreement and assurance they'll safeguard it (and not re-identify individuals, etc.). Or if you integrate with a wearable device maker and exchange data, you need to align privacy practices.

In summary, **your privacy posture is only as strong as your weakest link**. Make strengthening those links a priority by building privacy and security into your procurement and partner selection processes. Health tech users are entrusting their data not just to you but indirectly to everyone you work with – choose wisely and verify consistently.

6. Governing the Use of AI Ethically and Legally

From AI chatbots offering medical advice, to algorithms that analyse patient data for early disease detection, Artificial Intelligence is revolutionising health tech. But eagerly integrating AI/ML models to add value for users., especially when that involves the use of sensitive health data or decision-making that might affect a person's well-being, raises a host of **ethical and compliance concerns** that are being increasingly scrutinised by both regulators and the public.

Key issues include: **privacy** (how is sensitive data used to train or run the AI?), **bias and fairness** (does the AI produce discriminatory or inaccurate results for certain groups?), **transparency** (can you explain how the AI makes decisions or recommendations?), and **accountability** (who is responsible if the AI's output causes harm?).

From a **privacy** perspective, one risk is that companies might feed large amounts of personal health data into AI systems without proper safeguards. For instance, using real user data to train a machine learning model could inadvertently expose that data or lead to unintended secondary uses. If that training data is not anonymised properly, it could be breached or misused. Additionally, if you're using an external AI API (say a cloud AI service) and sending user data to it, that needs to be assessed like any other data transfer to a vendor – does the API provider maintain confidentiality? Do you have user consent to process their data with AI in this manner?

The **accuracy and reliability** of AI in health contexts is another concern. If an AI-driven feature gives a user health recommendations or risk scores, any error could have real consequences (unwarranted panic, or conversely false reassurance). *If* these errors stem from biased data or algorithms, it could also lead to regulatory issues. For example, an AI might under-diagnose a condition in women because it was trained mostly on male data – thus potentially violating anti-discrimination laws or upcoming AI regulations about bias. Organisations that deploy AI without fully understanding its workings risk **"black box" outcomes** that might harm users.

Regulators are responding. The **EU's AI Act** will specifically regulate AI systems deemed "high-risk," a category that is expected to include many healthcare and life-supporting AI applications. High-risk AI (like a diagnostic tool or a system that influences treatments) will require rigorous **risk assessments, documentation, transparency, human oversight, and in some cases even notification to authorities** before deployment. Non-compliance could result in fines up to 6% of global turnover under the AI Act – even higher than the fines levied by the GDPR. The AI Act also bans certain AI practices outright (e.g. social scoring, manipulative techniques) and imposes requirements on general-purpose AI providers. While the EU AI Act is EU-focused, its extraterritorial scope means if you deploy or even just make an AI system available in Europe, you have to comply.

Elsewhere, **the US is inching toward algorithmic accountability**: the California CPPA's regulations on Automated Decision-Making (ADMT) will force companies to disclose when decisions are algorithmically made and possibly allow consumer opt-outs[20]. The FTC has also warned it will use its powers to punish "unfair or deceptive" AI practices (for example, if an AI is biased or if a company lies about what its AI does). **China** already has regulations requiring transparency and user choice for recommendation algorithms and requires security reviews for "algorithms impacting public opinion", which is not directly health-related, but it shows the global trend.

Beyond strictly legal requirements, **ethical AI** is crucial for patient safety and trust. If your AI wellness coach gives someone dangerous advice, you could face not just lawsuits but public backlash and irreparable damage to your credibility. We've seen instances where generative AI "hallucinates" medical info – obviously problematic if presented to users as factual. Liability for AI outcomes is a gray area: if an AI suggests a course of action that leads to harm, could your company be held liable? Possibly, especially if due diligence in testing and oversight was lacking.

To govern AI properly:

Conduct AI Impact Assessments: Similar to DPIAs, these evaluate the potential risks of your AI systems and help you identify biases, privacy issues, error rates, and impacts on rights. The EU and others may mandate these for high-risk AI use cases, so get ahead by doing them now.

Data management for AI: Ensure any personal data used in developing or training AI is permitted and protected. Anonymise or pseudonymise wherever possible. Maintain a **record of what data went into training** and under what lawful basis.

Bias mitigation: Actively test your algorithms for bias or disparate outcomes. Use diverse training data and consult domain experts. If issues are found, retrain or put constraints in place. Document these efforts – under future AI laws you might need to prove you took steps to avoid bias.

Transparency and explainability: Strive to make your AI's functionality explainable to users. Even if the algorithm is complex, provide users with understandable reasons for outputs, e.g. "Our algorithm suggests a higher risk of X because it noted [factor1] and [factor2] in your data." Also, label AI interactions clearly (don't make a chatbot pretend to be human).

Human-in-the-loop: For any significant decisions or recommendations (especially those affecting health or eligibility for services), consider having a human review or an override mechanism. Many regulations encourage or require *human oversight* for high-impact automated decisions. Even if not required, it's a good safety practice.

Monitoring and iteration: Once an AI feature is live, monitor its outputs and user feedback. Set up a process to handle complaints or corrections (like if a user says the AI's suggestion was wrong or harmful). Continuously improve the model or rules as needed.

Define roles and responsibilities: Establish who "owns" AI governance in your organisation. Some companies set up an AI Ethics Committee; others assign the Chief Data Officer or similar to oversee AI compliance. Involve multidisciplinary perspectives – technical, legal, clinical (if it's health advice), etc.

In short, you should treat your AI with the same rigour as you would a core product offering – because to regulators and users, it is just that. The goal is to reap AI's benefits (personalization, scalability, efficiency) **without undermining privacy, fairness, or safety**. The companies that succeed in doing this will stand out as trustworthy innovators in the health tech field. Those that rush AI features to market without these safeguards may quickly find themselves in regulatory hot water or at the wrong end of a viral news story about an AI mishap.

Privacy maturity as a gateway to regulated buyers and regulated pathways

For health tech scale-ups, a compliant data privacy stance is increasingly a **commercial prerequisite,** not just a legal one. Regulated buyers (governments, national health services, large providers and payers) and regulators for medical devices expect evidence that you can manage sensitive health data safely, lawfully, and consistently across your product lifecycle. Companies that can demonstrate privacy-by-design, robust governance, and audit-ready evidence typically face **less procurement friction**, **shorter due diligence cycles**, and **fewer delays** in regulated approvals.

Why regulated buyers care

Public sector and highly regulated healthcare buyers operate under stringent confidentiality, cybersecurity, and data protection duties. Their vendor onboarding processes often treat privacy and security as "go / no-go" criteria. Weaknesses in consent, transparency, vendor governance, cross-border transfers, or incident response can lead to **exclusion from tenders**, prolonged security assurance loops, or onerous contractual terms.

Typical evidence buyers ask for

- A documented privacy management system (roles, policies, Records of Processing/ROPA, review cadence)

- Risk assessments (DPIAs/PIAs) for high-risk processing, including AI/automated decision-making where relevant

- A clear data flow map (what data are used, where they're stored, who has access, which vendors/processors touch them)

- International transfer governance (SCCs/IDTA, TIAs, and a central register of transfers)

- Security controls and assurance (e.g., encryption at rest/in transit, access controls, logging, vulnerability management)

- Incident response playbooks and tested breach notification procedures

- Evidence of staff training and secure SDLC practices

Privacy and security in medical device pathways (UK/EU)

If your product's functionality crosses into **Software as a Medical Device (SaMD)** or includes AI that influences clinical decisions, privacy and security move from "good practice" to essential elements of demonstrating safety, effectiveness, and lifecycle risk management. In the UK, MHRA guidance highlights that many software (including AI) products are regulated as medical devices. [GOV.UK]

A strong privacy posture helps you:

1. document intended use and data processing clearly;

2. evidence risk controls for sensitive data;

3. demonstrate secure development practices; and

4. support post-market surveillance and incident management.

NHS call-out: what "NHS-ready" typically means in practice

If you want your technology to be adopted by NHS organisations, buyers commonly expect you to align to NHS baseline standards and be able to evidence them during procurement. Two recurring requirements are:

- **Data Security and Protection Toolkit (DSPT):** NHS states that organisations with access to NHS patient data and systems must use the DSPT to provide assurance against the National Data Guardian's 10 data security standards. [Data Security and Protection Toolkit+1]

- **Digital Technology Assessment Criteria (DTAC):** NHS England's DTAC is used to assess digital health products against minimum baseline standards, including **data protection** and **cyber security**, to support faster and more consistent procurement decisions. [NHS Transformation Directorate+1]

- In practice, this means having an evidence pack that maps your policies, controls, DPIAs, vendor governance, and incident response to these frameworks—and being able to maintain that evidence annually as your product evolves.

US federal healthcare and FDA expectations

In the United States, privacy maturity materially affects your ability to sell into providers, payers, and public sector programmes. Even where HIPAA does not directly apply (e.g., some consumer health apps), enterprise healthcare buyers frequently require **HIPAA-aligned controls** and vendor due diligence as part of procurement.

Where your product is a regulated medical device, the FDA has elevated expectations for cybersecurity design and documentation in premarket submissions, reflecting the patient safety and data protection implications of connected devices. [U.S. Food and Drug Administration+1]

For many scale-ups, the fastest route to "federal-ready" is to build an audit-ready foundation: risk assessments, vendor governance, security evidence, and transparent user controls—so you can respond quickly to due diligence questions, reduce rework, and move through approvals and contracting with fewer delays.

Commercial advantages of doing this early

- **Shorter sales cycles:** fewer back-and-forth assurance questionnaires and faster security sign-off

- **Higher win-rates:** meeting baseline procurement standards avoids early disqualification

- **Lower cost of compliance:** privacy-by-design reduces expensive retrofits during approvals

- **Reduced operational risk:** clearer data flows and incident playbooks improve response and resilience

- **Stronger trust signals:** demonstrates maturity to regulators, partners, and investors

Benefits of Proactive Data & AI Privacy Practices

Having discussed the major challenges, it's clear that prioritizing data privacy and AI governance is not just about avoiding negatives, but also about **enabling positive outcomes**. Let's briefly highlight the benefits a health tech company can gain by getting privacy right, before we look at ways that Securys can support you in this journey.

Embracing strong data privacy and AI governance measures can seem resource-intensive, but it offers significant payoffs for health tech companies. Here are some key benefits:

- **Boosting Customer Trust:** Privacy is fundamentally about trust. When users know that their sensitive health information is handled with care – kept secure, not misused, only shared with consent – they are more likely to engage deeply and continuously with your product. Trust can be a differentiator in health tech; it turns users into advocates. Conversely, a privacy lapse can drive users away overnight. Prioritizing privacy thus helps you **build a loyal user base** that feels safe sharing data necessary for delivering value. This trust extends to enterprise clients too (hospitals, insurers) – they will favor partners who have demonstrable privacy credentials.

- **Scaling with Confidence:** By adopting **privacy-by-design** principles early and embedding compliance into your operations, you set up a foundation that can support rapid growth. As you expand to new countries or onboard thousands of new users, you can do so confidently knowing your practices meet the regulatory requirements in those markets. This reduces the risk of expansion surprises (like having to re-engineer a feature post-launch because it violated a local law). Essentially, strong privacy practices make your business more agile and resilient, enabling you to seize opportunities (new services, geographies) that others might shy away from due to regulatory complexity.

- **Competitive Advantage:** In a crowded digital health marketplace, companies often compete on features and price – but privacy can be a key value proposition too. An app that is **"privacy-first"** or has earned certifications (like ISO 27701 for privacy) can attract more users and B2B customers, especially as awareness of data issues grows. Also, having robust privacy and security can shorten sales cycles with enterprise customers who conduct detailed due diligence. You avoid costly rework later (like having to retrofit compliance after a big client demands it). In short, good privacy is an **investment that pays off** by setting you apart and preventing future costs.

- **Regulatory Peace of Mind:** While not as flashy as user growth, the peace of mind that comes from knowing you're on top of compliance is valuable. It frees up the team to focus on innovation rather than firefighting legal issues. It also positions you well for regulatory changes – e.g. if you've already got DPIAs and AI oversight processes internally, you'll adapt faster to new laws requiring the same. Companies with a strong privacy culture are less likely to suffer major breaches or fines, which protects their financial health and reputation. Think of privacy investment as **risk insurance** – it minimises the chance of catastrophic hits that could derail your scale-up journey.

- **Strengthened Partnerships and Ecosystem Position:** Health tech solutions often integrate into a broader healthcare ecosystem (connecting with clinics, insurance providers, pharma research, etc.). Demonstrating robust data governance can open doors to partnerships – hospitals or pharma companies will be more willing to work with you if they trust your handling of data. In some cases, being able to prove compliance (with standards like HIPAA, GDPR, etc.) is a prerequisite to even enter into a business development discussion in the healthcare space. Thus, privacy maturity directly correlates with your ability to form lucrative partnerships.

- **Future-proofing for AI Innovations ("AI by Design"):** As AI becomes integral to health tech, companies that bake in **"AI ethics and privacy by design"** will be ahead of the curve. By establishing governance now, you ensure that as your AI capabilities grow (e.g. more predictive analytics, automated coaching, clinical decision support), they do so on a solid ethical foundation. This can prevent nightmare scenarios and build stakeholder confidence that your AI is beneficial, not harmful. Plus, regulators are more likely to give leeway or work constructively with organisations that show a proactive stance on responsible AI.

In summary, focusing on privacy and data protection is not a cost center – it's a value creator. It builds **trust, brand equity, and operational excellence**, all of which ultimately drive growth. One might say privacy is the bedrock upon which a sustainable, reputable health tech business is built.

Now, with challenges identified and benefits clear, **how can health tech companies practically implement these principles?** The next section outlines concrete steps and how Securys can assist in each area to make privacy and AI governance achievable for organisations of all sizes.

How Securys Can Help: Making Privacy & AI Governance Practical

At Securys, we specialise in helping SMEs – including health tech scale-ups – navigate the complexities of data privacy and AI ethics. Our mantra is **Privacy made practical**® meaning we deliver solutions that are pragmatic, business-aligned, and proportionate to your risks. Whether you're just starting your privacy journey or need to refine mature processes, we take a **risk-based approach** to prioritise the areas of greatest vulnerability and value for your organisation. Our services are designed to build trust and compliance *without* stifling innovation. Here's how we can partner with you:

- **Comprehensive Privacy Audits & Strategy:** We begin by **building privacy into your products and operations from the ground up**. Our team can conduct a thorough **privacy audit** of your current practices – reviewing everything from how you collect data in your app, to how it's stored, secured, and shared. We benchmark this against regulatory requirements and industry best practices. The outcome is a *gap analysis* and a tailored **roadmap** for compliance that aligns with your business goals. This gives you a clear action plan on what to fix or improve, prioritised by risk. Essentially, we help you chart a course from your current state to an optimised privacy posture, so you know where to invest effort for maximum impact.

- **Data Mapping & Discovery:** Often the first step to solving privacy challenges is knowing exactly what data you have and where. Securys consultants can facilitate **data discovery and mapping workshops** with your team. We'll catalogue your data flows, systems, and processing activities, creating a **data inventory** that identifies personal data elements, sensitive health data, data on vulnerable groups (such as children) across your business. This map is invaluable for meeting documentation obligations (GDPR Article 30 records, etc.) and for making informed decisions about risk. It also supports easier responses to DSARs and breaches because you have a clear picture of your data landscape.

- **DPIAs, PIAs & Other Assessments:** Navigating regulatory obligations like DPIAs can be daunting. Our experts have deep experience running **Data Protection Impact Assessments (DPIAs)** under GDPR/UK law and similar **Privacy Impact Assessments (PIAs)** required by laws like Quebec's Law 25, as well as Legitimate Interest Assessments (LIAs), required if you use "legitimate interest" as a lawful basis for processing data, and AI Conformity Assessments (AICAs), required by the EU AI Act if you're using AI in a high risk use case (which includes use in health, HR and educational contexts). We'll work with you to **assess any high-risk processing.** If you're launching a new AI diagnostic feature, for example, or planning to share data with a research partner, we'll conduct a DPIA to identify privacy risks and recommend mitigations, as well as an AICA to cover off the risks of using AI. We ensure the DPIA process is proportionate (not a check-box exercise) and yields actionable insights to reduce risk. If needed, we can also assist with any required **regulatory consultations** (e.g. if a DPIA shows high residual risk, we help liaise with authorities). In the US. context, we can help you prepare for California's forthcoming risk assessments and even annual cybersecurity audits – shaping your internal processes so that when these rules kick in, you'll be ready. Ultimately, by engaging **Securys for you** assessment work, you'll gain both compliance documentation and peace of mind that you've got pre-emptive governance in place on any new projects.

- **Consent & Transparency Design:** Securys can help you design and implement **user-friendly consent mechanisms and privacy notices** that meet legal standards and

enhance user experience. Our team can review your current consent flows (e.g. sign-up forms, cookie banners, in-app prompts) and rewrite or reconfigure them to be clear, granular, and compliant. We also assist in creating **layered privacy notices** – succinct summaries with links to full details – so users are informed without being overwhelmed. Additionally, we advise on setting up **preference centers** where users can easily manage their consents and communication preferences over time (for example, giving users a dashboard to opt in/out of research data use, marketing emails, etc.). This not only keeps you compliant across jurisdictions (no more bundled or perpetual consents), but it also builds trust by putting users in control. If you have issues with minors' consent or parental consent verification, we can recommend solutions for that as well. In short, we bring best practices to ensure your **consent model is robust yet user centric**.

- **Third-Party & International Data Transfer Governance:** To tackle supply chain and cross-border challenges, Securys offers support in **vendor management and data transfer compliance**. We can help you implement a process to vet and contract with suppliers: including developing standard **contractual clauses** and checklists to ensure each vendor agreement covers security, sub-processor approval, audit rights, breach notification, and relevant jurisdictional requirements (like GDPR Article 28 terms or HIPAA BAAs as needed). We also assist with drafting and implementing **Standard Contractual Clauses (SCCs)** or the UK **International Data Transfer Agreement (IDTA)** for your data flows from EU/UK to third countries. If you plan to rely on mechanisms like the EU-US Data Privacy Framework or others, we'll guide you through that. Our team can perform **Transfer Impact Assessments (TIAs)** to evaluate cross-border data risks – which is increasingly expected by EU regulators post-Schrems II. For regions like India or China with specific rules, we'll help you adapt (e.g. ensuring compliance with India's negative list approach or preparing for China's certification/assessment protocols). We can if you wish effectively act as your **privacy engineering team for data flows** – making sure that moving data across borders or to partners doesn't become your Achilles' heel.

- **Security & Incident Response Readiness:** While you likely have IT teams focused on cybersecurity, Securys brings a privacy perspective to ensure security measures align with protecting personal data. We can review your **information security controls** against standards like ISO 27001/27701 and healthcare-specific practices. If gaps are found (say, missing encryption on certain data stores, or no policy for device management), we'll recommend pragmatic fixes. Crucially, we help you develop and test **Incident Response (IR) plans** that incorporate privacy breach handling. This includes creating breach escalation protocols, drafting **breach notification templates** for different jurisdictions (so you have them ready), and even running tabletop exercises with your team to simulate a data breach scenario. We want to ensure that if the worst happens, you can react swiftly and in compliance (72-hour regulator notices, etc.). Additionally, we can assist in setting up processes for handling **data subject requests** efficiently – e.g. templates and workflows for responding to access or deletion requests within deadlines. Being prepared not only helps avoid fines but also minimises business disruption in a crisis.

- **AI Governance Framework:** If AI is part of your product strategy, Securys can help establish an **AI governance framework** tailored to your organisation. This may involve defining internal roles like an *AI Product Owner* and an *AI Risk Manager* or leveraging your Data Protection Officer to also oversee AI ethics. We assist in creating **AI use policies** that set out guidelines for acceptable AI use cases, data training standards, and bias testing protocols. We can help you set up an **AI model registry** – a documented inventory of models, their purposes, training data, and risk level – which is very useful for accountability (and may be required under future AI laws). Our team stays abreast

of AI regulatory trends (EU AI Act, FTC guidelines, etc.), so we inject that compliance foresight into your AI projects. We also offer **AI Impact Assessments or AI Conformity Assessments (**similar to DPIAs, as mentioned above), to evaluate new AI deployments for ethical and privacy risks. With Securys's help in AI governance, you can innovate with AI confidently, knowing there's a safety net of oversight to catch issues like bias, drift, or non-compliance early.

- **Privacy Training & Culture Building:** Tools and processes alone aren't enough – people make privacy happen. Securys provides **training programmes** crafted for different stakeholder groups. For engineering and product teams, we deliver workshops on privacy by design, secure coding, and data minimization principles so they can bake privacy into development. For marketing and analytics teams, we focus on topics like lawful basis for campaigns and pseudonymization techniques. For clinical/medical staff interfacing with data (if you have any), we reinforce patient confidentiality and proper data handling. We also conduct engaging **executive briefings** for leadership, translating privacy and AI compliance into business risk terms that the C-suite cares about. Our training isn't generic; we use examples and scenarios from actual client engagements and, specifically, the **health tech sector** to make it relevant (e.g. what to do if a celebrity's fitness data is in your system and someone internally tries to peek at it – that's a privacy violation scenario we'd cover). By equipping your team with this knowledge, we help cultivate a **culture of privacy** where everyone feels responsible for protecting data, rather than it being just the compliance officer's job.

Working with Securys means you gain a partner who not only understands the legal requirements of your organisation but also understands the **health tech context and the resource constraints of a scale-up**. We aim to provide right-sized, practical measures – balancing strong protection with the agility you need to compete. Our services can be provided as one-off projects (like conducting a DPIA for a new launch) or as an ongoing retained privacy officer support.

Sample 12-week plan

To illustrate a path forward, here's a snapshot of a **12-week action plan** we often recommend for health tech SMEs looking to uplift their data privacy and AI governance:

- **Weeks 1–2: Discovery & Scoping** – We'll help you confirm the scope of your data environment: identify all jurisdictions you operate in, catalogue products/features and the types of personal data they handle, and flag any obvious high-risk areas (e.g. an AI module or a planned new data sharing initiative). By the end of week 2, you'll have a clear map of "what data, where, and which laws apply."

- **Weeks 3–4: Risk Assessments & Quick Wins** – We conduct DPIAs/PIAs on the most critical processing activities identified. Simultaneously, we begin supplier audits for key third parties (reviewing their contracts and practices). We'll deliver mitigation measures for any urgent risks (e.g. "enable encryption on X database now" or "pause launching in Country Y until consent flow is fixed").

- **Weeks 5–6: Policy and Documentation Refresh** – During this phase, we revamp your consent forms, privacy notices, and internal data handling policies. We also help you update your Records of Processing Activities (ROPA) and prepare draft SCCs/transfer assessments for your international data flows. Essentially, we generate the documentation and user-facing content that brings you in line with any regulatory requirements.

- **Weeks 7–8: Implement Technical Controls & IR Plans** – Now it's time to put improvements into action. We guide your IT team in implementing baseline security measures (maybe hardening access controls, enabling audit logs, etc.). We also finalise and test your Incident Response Plan, including breach simulation drills. By week 8, you'll have stronger defenses and a rehearsed plan for incidents.

- **Weeks 9–10: Establish AI and Data Governance** – We work with your product/data science team to formalise AI governance: setting up that model registry, documentation for algorithms, and human review processes for any automated decisions of significance. If needed, we also help designate a Data Protection Officer or similar role and set up a governance committee to review privacy-impacting plans going forward.

- **Weeks 11–12: Training & KPI Handover** – In the final stretch, we deliver targeted training sessions to your teams to ensure they understand the new policies and the importance of compliance. We also help you publish a privacy roadmap and define KPIs (key performance indicators) – for example, "% of new features going through DPIA" or "Avg. time to fulfill DSAR" – so you can measure and maintain progressfile://file_000000008fcc71f4a7406e636d9b0a09/. By the end of week 12, privacy is not a one-time project but an ongoing part of your operations, with people accountable for keeping it on track.

This intensive programme can be adjusted to your needs, but it illustrates how quickly improvements can be made with expert guidance. In three months, we've seen clients transform from having ad-hoc, reactive privacy measures to having a structured privacy programme that impresses investors, customers, and regulators alike.

Conclusion and Call to Action

Handling sensitive health data and deploying AI solutions across multiple jurisdictions is undoubtedly challenging, but it is manageable given the right approach. By understanding the regulatory landscape, addressing key privacy challenges proactively, and embedding strong governance practices, health tech companies can not only avoid pitfalls but actively **build trust and competitive advantage**. Privacy and AI compliance should be seen not as obstacles to innovation, but as enablers of sustainable innovation. A company that respects user data and is transparent about its practices will foster loyalty in a domain where trust is paramount to user engagement.

We've outlined the critical areas – from consent to breach response to AI ethics – and provided practical steps that you can follow to improve your data risk stance. This journey might seem complex, but you do not have to navigate it alone. **Securys is here to support you every step of the way.** We bring deep expertise in global data protection and a pragmatic mindset tailored for agile, high-growth companies like yours. Our team can help you implement the recommendations in this paper efficiently, drawing on proven methods and real-world experience with clients in the health tech and broader tech sectors.

By partnering with Securys, you gain not just advice, but hands-on assistance to *operationalise* privacy and AI governance. The outcome is a robust compliance posture that satisfies regulators in the UK, EU, US, India, Middle East, Caribbean and beyond – *and* sends a powerful message to your customers, investors, and partners that you take data stewardship seriously. In an era where users are becoming more privacy-conscious and laws are getting stricter, this is an investment in the longevity and integrity of your business.

Take the next step towards making privacy practical for your organisation. We encourage you to reach out to Securys for a consultation or to discuss your specific needs. Whether you require a baseline privacy health check, help with a DPIA for a product or a assessment for a new AI feature, a plan for training your staff in better data protection awareness, or a full privacy programme build-out, we are ready to assist.

Contact us to learn how we can tailor our services to your context and help you turn data privacy and AI governance into an asset and business strength, rather than a risk to be managed. Together, we can ensure that your health tech innovation not only improves lives, but does so in a way that is legal, ethical, and worthy of your users' trust.

Let's make privacy a competitive advantage for your health tech journey – connect with Securys and get started today.[21]

Appendix 1: Global Privacy & AI Compliance Matrix

To underscore the varying obligations across key jurisdictions – and why a one-size-fits-all approach won't work – below is a **quick-reference matrix** highlighting some major requirements in different regions. This kind of overview can help inform your compliance strategy in each market:

Jurisdiction	Sensitive Data Definition	Risk Assessment / DPIA	Cross-Border Transfers	Breach Notification	AI-Specific Obligations
EU/EEA (GDPR)	"Special category" data includes health, genetic, biometric data	DPIA required for high-risk processing; must consult DPA if residual risk is high	Restricted unless to adequate country or via safeguards (SCCs); EU–US DPF available for certified US entities	Notify DPA within 72h of serious breach; notify individuals if breach poses high risk	EU AI Act incoming: bans on certain AI (2025); rules for general AI (2025); strict obligations for high-risk AI systems by 2026–27
UK (UK GDPR + DUA 2025)	Similar "special category" definition (health data highly protected)	DPIA required for high-risk processing (mirrors EU GDPR)	Transfer mechanism required (UK SCCs/IDTA); new UK–US "data bridge" expected to ease UK-US transfers	Notify ICO within 72h of notifiable breach; notify affected individuals if high risk	New law mandates transparency for Automated Decision-Making (ADM); PECR (cookie) fines aligned to GDPR levels; further reforms in

					progress[22]
US (HIPAA + state laws)	**HIPAA:** PHI (identifiable health info in healthcare contexts); **State laws:** e.g. "Consumer health data" under WA law (broadly includes wellness info)	No general federal DPIA yet. **California CPPA regs** will require risk assessments for certain data uses starting 2026[23]. Some sectoral guidance (e.g. FDA for certain health tech)	No federal cross-border rules, but contracts and due diligence expected. Data exports largely governed by company policy and any sectoral regs. (Some states may introduce transfer restrictions in future.)	Varies by state: all states have breach notification laws (often ~30-60 days to notify individuals); HIPAA breaches affecting >500 must be reported to HHS and media within 60 days[24]	No comprehensive AI law yet. California's upcoming **ADMT rules (2027)** will require transparency and opt-out for automated decisions [25]. FTC can enforce against biased or unfair AI. WMHMDA in Washington bans certain uses (like geofencing around clinics to target ads).
India (DPDP Act 2023)	Defines "personal data" broadly; no explicit "sensitive"	"Significant Data Fiduciaries" (large data handlers) will have to conduct Data Protection Impact Assessments and audits – risk-based approach. Others have lighter obligations	Government to whitelist/ blacklist countries for transfers ("negativ	Breach reporting criteria to Data Protection Board	No dedicated AI law yet. However, draft policy discussio

	category, but children's data and some sensitive processing have extra safeguards		e list" model). If destination is not banned, transfers allowed but with expectations of similar protection	to be defined by rules (likely mandatory reporting of significant breaches). Timelines TBD in Rules, but organisations expected to report expeditiously	ns emphasise responsible AI and potential future guidelines. For now, general IT and data laws apply; strong public/consumer protection laws can be used if AI causes harm or bias.
Middle East (e.g. UAE)	UAE PDPL: health data is considered sensitive; *plus* separate **health data laws** (e.g. Dubai Healthcare City regulations, etc.) impose strict rules on	Federal PDPL adopts risk-based approach (DPIAs not explicitly mandated except in certain contexts). However, **DIFC** (Dubai Int'l Financial Centre) data law explicitly requires DPIAs for high-risk processingfile://file_000000008fcc71f4a7406e636d9b0a09/.	UAE PDPL allows transfers to countries with adequate protection; otherwise requires contracts or regulator approval[26]. Many GCC countries similarly restrict exports unless certain	Under UAE PDPL, breaches likely need to be reported to the regulator (awaiting further guidance). DIFC law requires breach notification to DIFC Commis	AI governance in Middle East is evolving. The UAE has an AI ethics charter and sectoral guidelines, but no binding AI law yet. Expect future requirements as Gulf states

			safeguards or permits are in place.	sioner of Data Protection. Timelines vary (e.g. "as soon as practicable").	adopt more AI in healthcare (e.g. possible licensing or transparency rules).
Caribbean (e.g. Barbados, Jamaica)	New laws align with GDPR definitions: health data is sensitive personal data requiring higher protection[27].	Generally no mandatory DPIA in most (Barbados doesn't mandate, but encourages risk assessments; Jamaica's law doesn't explicitly mandate DPIAs but Office of Information Commissioner may issue guidance). Still, best practice is to do PIAs especially for new tech deployments.	Many Caribbean laws **model GDPR's approach**: cross-border transfer allowed only to adequate jurisdictions or with safeguards (standard clauses or individual consent)[28]. Jamaica's law, for instance, requires similar standards for international transfers.	**Breach Notification:** Jamaica requires breaches to be reported to authorities and affected persons within 72 hours[29]. Barbados requires notification "as soon as reasonably practicable." These laws closely mirror EU timelines.	No AI-specific statutes yet in Caribbean privacy laws. However, general provisions on automated processing and fair treatment exist (Barbados gives the right to object to solely automated decisions [30]). As AI usage grows, regulators may issue guidelines. Companies should self-

					regulate AI impacts in anticipation.

Note: The above matrix is a simplified snapshot. Within each jurisdiction, there are many nuances and possibly other sectoral laws (e.g. specialised health privacy laws in certain US. states or in countries like Singapore, Australia, etc.). Health tech firms should seek local legal advice for specifics. Nonetheless, the matrix illustrates why a global health tech company must keep track of differing legal requirements – from breach notice windows to AI rules – and incorporate them into their compliance programme.

Appendix 2: Regulated Buyer Readiness Checklist

A practical self-assessment for health tech scale-ups selling into regulated healthcare markets

This checklist is designed to help health tech companies assess whether their **data privacy, security, and AI governance posture** is ready for engagement with regulated buyers such as national health services, government agencies, large healthcare providers, or for progression through medical device and federal healthcare compliance pathways.

It reflects the types of evidence most commonly requested during procurement, assurance, and regulatory review.

---

1. Organisational governance & accountability

☐ Clear ownership of data protection and privacy (e.g. DPO, privacy lead, or equivalent)
☐ Defined roles and responsibilities for data protection, security, and AI governance
☐ Up-to-date privacy policies and internal data handling standards approved by leadership
☐ Regular review cycle for privacy and security governance (not one-off compliance)

---

2. Data mapping & transparency

☐ Comprehensive inventory of personal and health data processed
☐ Documented data flows showing collection, storage, processing, sharing, and deletion
☐ Clear identification of sensitive data (health, biometric, genetic, children's data)
☐ Accurate and accessible privacy notices aligned to actual data practices

---

3. Lawful basis & consent management

☐ Lawful basis identified for each category of processing in each key jurisdiction
☐ Explicit consent mechanisms in place for sensitive health data where required
☐ No bundled or "forced" consent for optional processing (e.g. marketing, analytics)
☐ Ability for users to withdraw consent easily and for systems to honour that choice
☐ Evidence of consent capture and auditability

---

4. Risk assessments & assurance

☐ Data Protection Impact Assessments (DPIAs / PIAs) completed for high-risk processing
☐ AI or automated decision-making risk assessments where relevant
☐ Clear mitigation actions documented and implemented
☐ Escalation process defined where residual risk remains
☐ Assessments reviewed when products or data uses change

---

5. Security & incident readiness

☐ Baseline security controls implemented (access control, encryption, logging)
☐ Secure development lifecycle practices in place

☐ Tested incident response plan covering data breaches and security incidents
☐ Jurisdiction-specific breach notification timelines understood and documented
☐ Evidence of regular security and privacy training for staff

---

6. Supplier & supply-chain assurance

☐ Inventory of all third parties and vendors with access to personal or health data
☐ Data processing agreements in place with appropriate privacy and security clauses
☐ Sub-processor visibility and approval mechanisms
☐ Periodic vendor risk reviews or audits
☐ Controls in place for SDKs, analytics tools, and embedded third-party components

---

7. International data transfers

☐ Clear record of where data is stored and processed globally
☐ Transfer mechanisms implemented where required (e.g. SCCs, IDTA, TIAs)
☐ Jurisdiction-specific transfer rules understood (EU/UK, India, Middle East, Caribbean)
☐ Central register of international transfers maintained and reviewed

---

8. AI governance & responsible use (if applicable)

☐ Inventory of AI models and automated decision systems in use
☐ Defined purpose and scope for each AI use case
☐ Controls for data used in training and inference
☐ Bias, accuracy, and performance evaluation processes
☐ Human oversight for high-impact or health-related decisions
☐ Clear user transparency about AI involvement

---

9. Regulated buyer & NHS readiness

☐ Evidence aligned to public sector procurement expectations
☐ NHS DSP Toolkit position understood and achievable
☐ Alignment with NHS DTAC data protection and cyber security criteria
☐ Ability to respond quickly to security and privacy due-diligence questionnaires
☐ Audit-ready documentation available on request

---

10. Medical device & federal healthcare pathways (where relevant)

☐ Clear understanding of whether the product qualifies as SaMD
☐ Privacy and security risks addressed as part of product lifecycle management
☐ Documentation supporting safety, effectiveness, and post-market monitoring
☐ HIPAA-aligned controls in place where selling into US healthcare environments
☐ Evidence suitable for FDA, MHRA, or equivalent regulatory review

How to use this checklist

- **Mostly ticked?** You are likely well-positioned to engage regulated buyers and progress through formal assurance processes.

- **Several gaps?** These are common for growing health tech companies and can usually be addressed quickly with a structured, risk-based approach.

- **Unsure how to evidence items?** That uncertainty itself is often what delays procurement or regulatory approval.

References

[1] Digital Health Market worth US$573.5 billion by 2030 with 23.6% CAGR | MarketsandMarkets™

https://www.prnewswire.com/news-releases/digital-health-market-worth-us573-5-billion-by-2030-with-23-6-cagr--marketsandmarkets-302547830.html

[2] KPMG global tech report: Healthcare insights

https://kpmg.com/pt/en/insights/2025/03/kpmg-global-tech-report-healthcare.html

[4] [5] [19] Research shows data breach costs have reached an all-time high | CyberScoop

https://cyberscoop.com/ibm-cost-data-breach-2025/

[3] [7] [8] [16] [21] [22] UK Enacts Data Use and Access Act 2025 | ReedSmith

https://www.reedsmith.com/our-insights/blogs/technology-law-dispatch/102lu0t/uk-enacts-data-use-and-access-act-2025/

[9] [20] [23] [25] California Finalizes Regulations to Strengthen Consumers' Privacy

https://cppa.ca.gov/announcements/2025/20250923.html

[10] [15] [26] United Arab Emirates Allows Cross Border Data Flows of Personal Data

https://www.trade.gov/market-intelligence/united-arab-emirates-allows-cross-border-data-flows-personal-data

[11] [12] [13] [14] [17] [27] [28] [29] [30] A Regional Perspective on Privacy: Caribbean Data Protection Laws and the Case for Reform in The Bahamas - Lexology

https://www.lexology.com/library/detail.aspx?g=f19f3d93-5ee7-4b96-b6bf-58c489c0d463

[6] Is ChatGPT HIPAA Compliant? Updated for 2025

https://www.hipaajournal.com/is-chatgpt-hipaa-compliant/

[18] GDPR Compliance Failures Lead to Surge in Fines - Sentra

https://www.sentra.io/blog/gdpr-compliance-failures-lead-to-surge-in-fines

# About Securys

Our experience

Securys believe privacy matters because people matter. We have delivered global information security, privacy, data governance and ethical AI services to clients around the world and use our wide and deep experience of cyber, data protection, regulation and governance to bring a strongly practical approach to helping organisations of all sizes protect themselves and their stakeholders.

From the Arctic Circle to Australia Securys has supported enterprise and SME clients alike across sectors as diverse as financial services, commodities extraction, healthcare and luxury retail, as well as supporting the non-profit sector in the UK and beyond.

Our team

Based on experience with enterprise clients in over 60 jurisdictions worldwide, Securys draws on the expertise of an international team of multilingual consultants to engage with key stakeholders at all levels of a business in order to develop a profound understanding of the way clients work.

Collectively our team has a wealth of relevant data protection and information security certifications, including CIPP/E, CIPP/A, CIPP/US, CIPM, FIP, CISSP, ISSMP, CISA and AIGP. Our privacy and information security management framework is certified by BSI to comply with ISO27001 and ISO27701. We are corporate members of the International Association of Privacy Professionals. More importantly, we have decades of collective experience in the management and governance of organisations, so we know how to put the theory into practice.
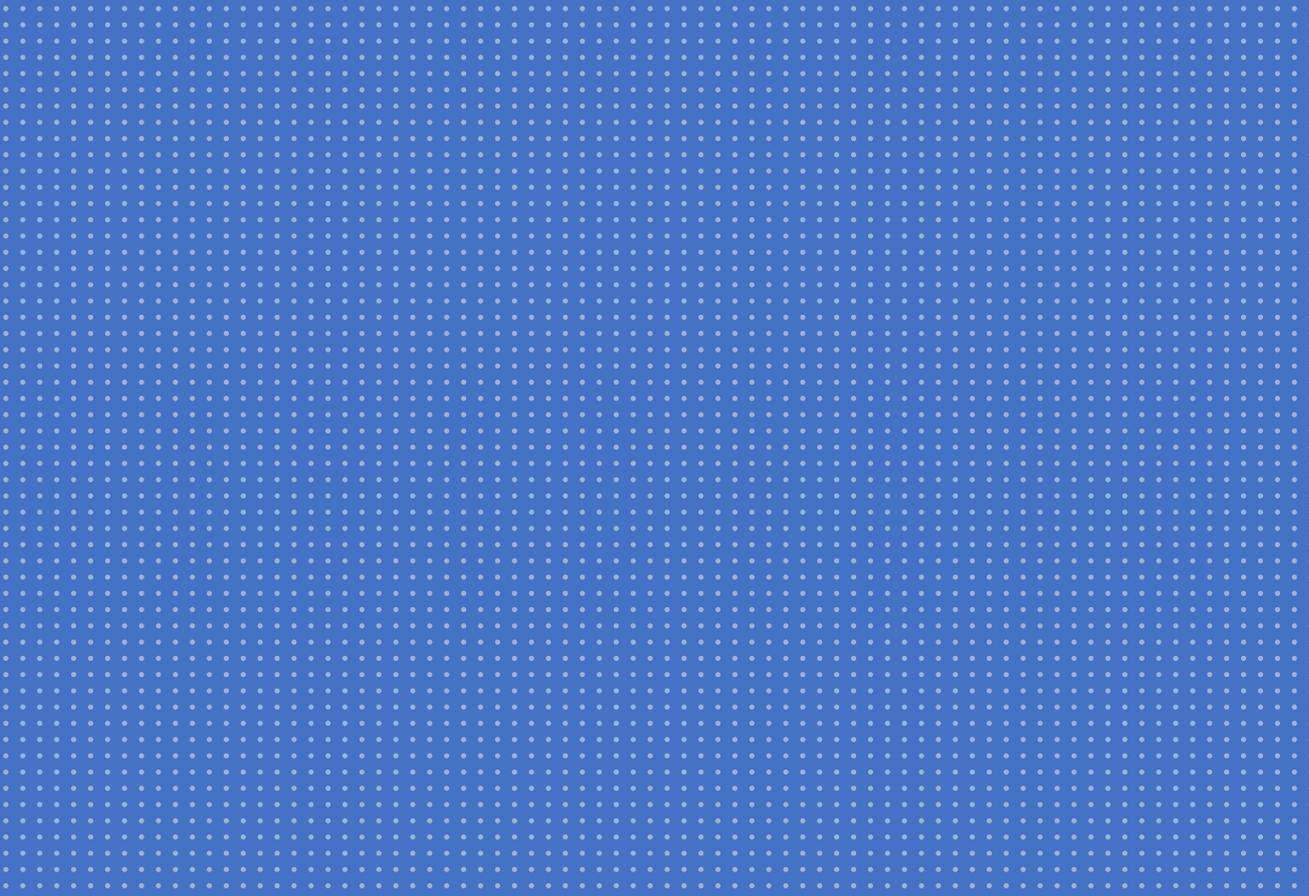
A global presence, locally delivered

Headquartered in London – with offices in Europe, the Caribbean, India and the US –
we have practical on-the-ground experience in most of the world's privacy regimes with
our team of international experts operating across dozens of jurisdictions.

Europe

United Kingdom
4th Floor
91 Goswell Road
London
EC1V 7EX

Ireland
28 Upper Fitzwilliam
Street
Dublin 2
Ireland

Caribbean

Jamaica
9th Floor
PanJam Building
60 Knutsford Boulevard
Kingston
Jamaica W.I.

Saint Lucia
20 Micoud Street
Castries
Saint Lucia, W.I.

Asia

India
308A Commerce House
Nagindas Master Road
Fort
Mumbai
400001, India

USA

Trust Center
1209 Orange Street
Wilmington
Delaware 19801
United States of
America

Securys®
Global Data Privacy Experts