

PRIVACY-ENHANCED THREAT MODELLING FOR WEARABLE HEALTH DEVICES: A CRITICAL REVIEW OF STRIDE, LINDDUN, AND PRIVACY IMPACT ASSESSMENT INTEGRATION

¹NWOKPURU SAMUEL ABAFU, ²CHINONSO JOB, ³ONWE, FESTUS CHIJOKE

¹Department of Computing, Software Engineering, University of Greater Manchester

²University of greater Manchester, United Kingdom.

³Information Technology Department, University of Port Harcourt, Rivers State, Nigeria.

DOI: <https://doi.org/10.5281/zenodo.21060970>

Published Date: 30-June-2026

Abstract: Wearable health devices (WHDs) continuously collect, transmit, and store sensitive biometric and physiological data, placing them at the intersection of three threat modelling traditions that have historically developed in isolation: security-oriented modelling (STRIDE), privacy-oriented modelling (LINDDUN), and regulatory compliance assessment (Privacy Impact Assessment, PIA). This paper presents a critical review of how these three frameworks address—and fail to fully address—the privacy and security risks specific to WHDs. We synthesise findings from systematic reviews, framework extension studies, and empirical case studies published between 2019 and 2025, organised around four recurring gaps identified across the literature: the structural incompleteness of any single framework when applied to continuous biometric monitoring; the persistent difficulty of mapping privacy threats onto specific GDPR obligations; the usability and tooling barriers that limit practitioner adoption of privacy-specific modelling; and the absence of a validated method for combining STRIDE, LINDDUN, and PIA into a single coherent assessment workflow. We conclude that, while each framework individually addresses a necessary dimension of the WHD threat landscape, no published, empirically validated integration of all three currently exists, and we set out the specific design requirements such an integration would need to satisfy.

Keywords: Wearable health devices, threat modelling, LINDDUN, STRIDE, Privacy Impact Assessment, GDPR, health IoT security, privacy engineering.

I. INTRODUCTION

WEARABLE health devices (WHDs)—ranging from consumer fitness trackers to implantable and clinically prescribed biosensors—have moved from novelty to clinical and regulatory significance over the past decade, enabling continuous remote monitoring, early anomaly detection, and personalised treatment adjustment [1]. This same continuity of data collection is the source of their distinctive risk profile: a WHD that streams heart-rate, location, and activity data around the clock generates a persistent attack surface and a correspondingly persistent privacy exposure that episodic, session-based systems do not present [2], [3].

Three threat-modelling traditions have developed largely independently to address different facets of this exposure. STRIDE, Microsoft's security-threat taxonomy (Spoofing,

Tampering, Repudiation, Information disclosure, Denial of N. S. Abafu is with the Department of Computing, Software Engineering. Manuscript prepared for submission; corresponding author email to be inserted by the author prior to submission.

service, Elevation of privilege), provides a structured method for identifying technical attack vectors against a system's data-flow diagram [4]. LINDDUN extends this idea to privacy specifically (Linkability, Identifiability, Non-repudiation, Detectability, Disclosure of information, Unawareness, Noncompliance), targeting privacy harms that are not necessarily security failures in the STRIDE sense [5], [6]. Privacy Impact Assessment (PIA) operationalises the legal and regulatory obligations imposed by frameworks such as the EU General Data Protection Regulation, providing a structured compliance documentation process rather than a technical threat taxonomy [7], [8].

Each framework, applied alone, is well validated within its own scope. The central finding of this review, consistent across the literature surveyed in Section III, is that applying these frameworks separately—rather than as a deliberately integrated workflow—produces systematic coverage gaps precisely at the points where WHDs are most exposed: continuous biometric monitoring, cloud-dependent data pipelines, and automated decision-making subject to specific GDPR obligations (notably Article 22 on automated profiling).

A. Contributions

This review contributes: (i) a synthesis of the documented individual and comparative strengths and limitations of STRIDE, LINDDUN, and PIA specifically in WHD and health-IoT contexts; (ii) identification of four recurring, literature-wide gaps that motivate integration of the three frameworks; and (iii) a set of design requirements, derived directly from the gaps identified, that any future integrated framework would need to satisfy to be considered a genuine advance rather than a re-labelling of existing methods.

B. Review Methodology

This is a structured narrative review. Sources were identified through targeted search of IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, arXiv, and Google Scholar, using combinations of the terms “LINDDUN,” “STRIDE,” “Privacy Impact Assessment,” “wearable health device,” “GDPR,” and “threat modelling,” restricted to 2019–2025 to capture the postGDPR literature, with priority given to peer-reviewed journal articles and the foundational framework-defining papers. The fourteen primary sources synthesised in Table I were each assessed for research methodology, principal findings, and stated limitations, following standard narrative-review synthesis practice.

The remainder of the paper proceeds as follows. Section II positions WHDs as a distinctive threat-modelling context. Section III synthesises the comparative literature on STRIDE, LINDDUN, and PIA. Section IV identifies four recurring gaps. Section V derives design requirements for a future integrated framework, and Section VI concludes.

II. BACKGROUND: WHY WEARABLE HEALTH DEVICES ARE A DISTINCTIVE THREAT-MODELLING CONTEXT

WHDs differ from the desktop- and server-oriented systems for which STRIDE and, to a lesser extent, LINDDUN were originally developed, along three dimensions consistently emphasised in the literature. First, continuous biometric streaming means that the relevant data flow is not a discrete transaction but an unbroken time series, such that even correctly authenticated, authorised access can constitute a privacy harm if the resulting behavioural profile is itself sensitive [3], [9]. Second, WHDs are resource-constrained: limited battery, compute, and bandwidth budgets restrict the cryptographic and monitoring overhead that can be deployed on-device, pushing security-relevant processing into companion mobile applications and cloud back-ends and thereby multiplying the number of trust boundaries a threat model must capture [9]. Third, WHD data is frequently subject to dual regulatory regimes—general data-protection law (GDPR in the EU/UK) and, where the device is positioned as a medical device or processes data for clinical purposes, sector-specific health-data regulation—creating compliance obligations that purely technical threat models do not natively represent [7], [10].

III. COMPARATIVE LITERATURE ON STRIDE, LINDDUN, AND PIA IN HEALTH CONTEXTS

Table I synthesises the primary studies underpinning this review.

A. STRIDE in Health-IoT Contexts

STRIDE's six-category taxonomy was designed for general software systems and maps cleanly onto classical attack vectors—spoofed device identity, tampered firmware, denial-of-service against a monitoring pipeline—but Vakhter et al.'s empirical adaptation for biomedical wearables shows that STRIDE alone better surfaces *adversarial* scenarios than the subtler, often non-adversarial privacy harms that arise from correctly functioning but excessively revealing data flows [9]. Aijaz et al.'s systematic review reaches the same conclusion from a different methodological angle: STRIDE is consistently

reported as stronger for attack classification, with LINDDUN identified as the necessary complement for privacy-specific threats [11].

B. LINDDUN and Its Documented Limitations

LINDDUN's seven-category privacy taxonomy directly targets the harms STRIDE under-serves, but Van Landuyt and Joosen's descriptive study of LINDDUN's own elicitation process finds that analysts applying LINDDUN make implicit, often undocumented assumptions about system boundaries and data flows, introducing inconsistency in how the same system can be modelled by different practitioners [5]. This finding is significant for WHDs specifically, where the boundary between on-device processing, companion-app processing, and cloud processing is itself a frequent source of ambiguity. Robles-Gonzalez et al. extend LINDDUN toward biometric identification and authentication processes, a direct precursor to WHD-specific privacy modelling, but their extension remains validated only for identification/authentication flows rather than the continuous-streaming case that characterises most WHD telemetry [15]. Wuyts et al.'s LINDDUN GO addresses a different limitation entirely: usability. Full LINDDUN's analytical depth comes with a documented tooling and training burden that LINDDUN GO's lightweight, card-based elicitation method is explicitly designed to reduce, at the cost of some analytical granularity [6].

C. PIA and the GDPR Compliance Gap

Azam et al.'s GDPR-perspective survey finds that PIA, while effective at surfacing transparency obligations, is resourceintensive to apply and has not been adequately adapted for the scalability constraints of edge IoT devices such as WHDs [12]. Ekdahl and Nyman's data-flow-based methodology for validating GDPR compliance represents a concrete attempt to operationalise PIA more rigorously, but, by the authors' own account, has not been validated through field deployment [8]. LeClair's comparative study of legal/procedural PIA alignment against user risk perception is notable for a finding with direct relevance to WHDs: formal compliance, as captured by a PIA, does not reliably track what users themselves perceive as risky, implying that PIA compliance and user trust are not interchangeable outcomes [16].

D. Toward Integration: What the Literature Has and Has Not Yet Established

Popoola's doctoral work is the most ambitious attempt in the surveyed literature to combine STRIDE- and LINDDUNstyle threat modelling with machine-learning-based privacyviolation prediction and blockchain-mediated consent management, reporting strong self-evaluated predictive performance for the resulting privacy-violation prediction component [13]. Bugeja et al.'s PRASH framework, while developed for smart-home rather than WHD contexts specifically, demonstrates via simulation that adaptive, context-aware privacy risk modelling can expose genuine privacy/utility trade-offs that static frameworks do not surface [14]. Taken together, these works establish that components of an integrated approach— privacy scoring, adaptive risk modelling, ML-assisted anomaly detection—have each been independently demonstrated, but, consistent with the broader privacy-threat-modelling literature surveyed in a recent IoT-focused framework review [17], no published study has combined STRIDE, LINDDUN, and PIA specifically into a single validated workflow for the WHD domain.

TABLE I: SYNTHESIS OF PRIMARY LITERATURE ON STRIDE, LINDDUN, AND PIA IN HEALTH AND WEARABLE IOT CONTEXTS

Source	Methodology	Theme	Principal Finding	Stated Limitation
Aijaz et al. [11]	Systematic literature review	STRIDE/LINDDUN in healthcare IT	STRIDE stronger for attack classification; LINDDUN excels at privacy-specific threats	Limited empirical validation across studies reviewed
Azam et al. [12]	Survey, GDPR-perspective synthesis	PIA in autonomous/IoT systems	PIA improves transparency but is resource-intensive; scalability to edge IoT under-addressed	Scalability to constrained edge devices not resolved
Popoola [13]	PhD thesis; experimental/empirical	ML- and blockchain-assisted privacy framework	Reports high predictive accuracy for a privacy-violation prediction model (self-reported AUC and consent accuracy metrics)	Scope concentrated on a specific smart-home/health-ecosystem context
Bugeja et al. [14]	Design science + simulation (PRASH)	Contextual privacy risk in smart homes	Demonstrates privacy/utility tradeoffs in adaptive risk models	Limited real-world user testing

Robles-Gonzalez' et al. [15]	Framework extension	LINDDUN for identification/authentication	Extends LINDDUN to biometric identity protection scenarios	Framework-specific; not validated for WHD continuous streaming
Van Landuyt & Joosen [5]	Descriptive/modeldriven analysis	LINDDUN assumption auditing	Identifies inconsistent assumptions embedded in LINDDUN elicitation	Implementation complexity for practitioners
Ekdahl & Nyman [8]	Conceptual/dataflow analysis	PIA–GDPR alignment	Proposes data-flow-based method for validating GDPR compliance	No field deployment reported
Vakhter et al. [9]	Empirical case study	STRIDE adaptation for biomedical wearables	LINDDUN better surfaces privacy harms; STRIDE better surfaces adversarial scenarios	Limited diversity of devices tested
LeClair [16]	Comparative user-risk-perception study	PIA/regulatory alignment	User risk perception diverges from formal threat-model output	Legal/procedural focus; limited technical depth
Wuyts et al. [6]	Framework extension + case study	LINDDUN usability	GO Lightweight card-based LINDDUN variant improves practitioner adoption	Reduced analytical depth relative to full LINDDUN

IV. RECURRING GAPS ACROSS THE LITERATURE

Four gaps recur across the sources synthesised in Section III, independent of which specific framework a given study centres on.

A. Structural Incompleteness Under Continuous Monitoring

Every framework reviewed was designed primarily around discrete, transactional data flows. WHD telemetry is structurally continuous, and Vakhter et al. and the broader wearable-security survey literature both report that continuous streaming generates privacy and security implications— behavioural profiling, long-horizon re-identification risk—that none of STRIDE, LINDDUN, or PIA individually treats as a first-class modelling concern [2], [9].

B. The GDPR Article 22 Mapping Gap

Automated decision-making and profiling, regulated specifically under GDPR Article 22, is identified by Azam et al. as inadequately addressed by LINDDUN's privacy taxonomy and only partially addressed by standard PIA documentation, which tends to focus on data collection and storage rather than downstream automated inference [12]. This is a direct gap for WHDs, whose value proposition increasingly depends on automated health-risk inference from continuous data.

C. Usability and Tooling Barriers

Van Landuyt and Joosen's assumption-auditing findings and Wuyts et al.'s motivation for LINDDUN GO both point to the same underlying problem from different directions: privacy threat modelling, in its analytically complete form, imposes a training and tooling burden that limits practitioner adoption, and the lightweight alternatives developed to address this trade analytical depth for usability rather than resolving the underlying tension [5], [6].

D. Absence of a Validated Tri-Framework Integration

No source identified in this review reports an empirically validated combination of STRIDE, LINDDUN, and PIA applied together to a single WHD use case. The closest precedents—Popoola's ML/blockchain-assisted privacy framework [13] and Bugeja et al.'s adaptive smart-home risk model [14]—each combine some but not all of the three traditions, and neither targets the WHD domain specifically with all three integrated.

V. DESIGN REQUIREMENTS FOR A FUTURE INTEGRATED FRAMEWORK

Based on the gaps identified in Section IV, we set out four design requirements that a future STRIDE–LINDDUN–PIA integration for WHDs would need to satisfy to constitute a genuine advance over applying the three frameworks separately.

- 1) Continuous-flow representation. The framework's underlying data-flow representation must treat behavioural and biometric time series as a first-class modelling object, rather than approximating continuous telemetry as a sequence of discrete transactions, directly addressing the structural-incompleteness gap (Section IV.A).

- 2) Explicit Article 22 threat category. Automated profiling and inference should be elicited as an explicit threat category linked directly to GDPR Article 22 obligations, rather than left to be inferred indirectly from LINDDUN's existing seven categories or a generic PIA data-flow diagram (Section IV.B).
- 3) Tiered usability. Following the precedent set by LINDDUN GO [6], the framework should offer a lightweight elicitation pathway for practitioners without specialist privacy-engineering training, while preserving an option to escalate to full analytical depth for safety- or compliance-critical components (Section IV.C).
- 4) Empirically validated integration protocol, not just conceptual combination. The integration of the three frameworks must be tested against a real or realistic WHD use case, with reported coverage (the breadth of threats identified relative to a known threat catalogue) and practitioner usability outcomes, rather than asserted by design alone (Section IV.D).

VI. CONCLUSION

This review has synthesised the literature on STRIDE, LINDDUN, and Privacy Impact Assessment as applied, individually, to health-IoT and wearable contexts. Each framework is well validated within its own scope: STRIDE for adversarial, attack-oriented threats; LINDDUN for privacy-specific harms not captured by security modelling alone; and PIA for regulatory compliance documentation. However, four gaps recur consistently across the literature when these frameworks are considered for wearable health devices specifically: structural incompleteness under continuous biometric monitoring, an inadequately mapped GDPR Article 22 automated-decision making category, persistent usability and tooling barriers limiting practitioner adoption, and—most significantly—the absence of any published, empirically validated integration of all three frameworks for the WHD domain. We have derived four concrete design requirements from these gaps that future integration work should satisfy, and a companion paper proposes one candidate framework, the Integrated Privacy Threat Modelling (IPTM) approach, explicitly designed against these requirements and presented together with a validation protocol for the empirical testing this review has shown to be currently missing from the literature.

LIMITATIONS

As a structured narrative review rather than a fully systematic review, this paper does not report PRISMA-style screening statistics, and the fourteen primary sources synthesised in Table I reflect a targeted rather than exhaustive search. Some claims regarding self-reported performance metrics (notably from [13]) are reported as the original authors' own findings rather than independently re-verified by this review.

REFERENCES

- [1] A. K. Islam Riad, "Security and privacy analysis of wearable health device," *International Journal of Software Engineering Research*, vol. 12, no. 7, pp. 45–53, 2021.
- [2] "A survey on security and privacy issues in wearable health monitoring devices," *Computers & Security*, 2025.
- [3] S. Khan, S. Parkinson, L. Grant, N. Liu, and S. McGuire, "Biometric systems utilising health data from wearable devices: Applications and future challenges in computer security," *ACM Computing Surveys*, vol. 53, no. 4, p. 85, 2020.
- [4] R. J. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, 2nd ed. Hoboken, NJ: John Wiley & Sons, 2010.
- [5] D. Van Landuyt and W. Joosen, "A descriptive study of assumptions made in LINDDUN privacy threat elicitation," in *Proceedings of the 35th Annual ACM Symposium on Applied Computing*, 2020, pp. 1280–1287.
- [6] K. Wuyts, L. Sion, and W. Joosen, "LINDDUN GO: A lightweight approach to privacy threat modeling," in *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 2020, pp. 302–309.
- [7] European Parliament and Council of the European Union, "Regulation (eu) 2016/679 (general data protection regulation)," *Official Journal of the European Union*, L 119, 2016.
- [8] A. Ekdahl and L. Nyman, "A methodology to validate compliance to the GDPR," *arXiv preprint*, 2019, arXiv: 1901.08387.

- [9] V. Vakhter, B. Soysal, P. Schaumont, and U. Guler, "Threat modeling and risk analysis for miniaturized wireless biomedical devices," *IEEE Internet of Things Journal*, vol. 9, no. 15, pp. 13338–13352, 2022.
- [10] "“whispers from the wrist”: Wearable health monitoring devices and privacy regulations in the U.S.: The loopholes, the challenges, and the opportunities," *Cryptography*, vol. 8, no. 2, p. 26, 2024.
- [11] M. Aijaz, M. Nazir, and M. N. A. Mohammad, "Threat modelling and assessment methods in the healthcare-it system: A critical review and systematic evaluation," *SN Computer Science*, vol. 4, no. 6, p. 714, 2023.
- [12] N. Azam, L. Michala, S. Ansari, and N. B. Truong, "Data privacy threat modelling for autonomous systems: A survey from the GDPR's perspective," *IEEE Transactions on Big Data*, vol. 9, no. 2, pp. 388–414, 2023.
- [13] O. J. Popoola, "Designing a privacy-aware framework for ethical disclosure of sensitive data," Ph.D. dissertation, Sheffield Hallam University, 2025.
- [14] J. Bugeja, A. Jacobsson, and P. Davidsson, "PRASH: A framework for privacy risk analysis of smart homes," *Sensors*, vol. 21, no. 19, p. 6399, 2021.
- [15] A. Robles-Gonzalez, J. Parra-Arnau, and J. Forn' e, "A LINDDUN-based' framework for privacy threat analysis on identification and authentication processes," *Computers & Security*, vol. 94, p. 101755, 2020.
- [16] B. LeClair, "Threat modelling in virtual assistant hub devices compared with user risk perceptions," *arXiv preprint*, 2023, arXiv:2301.12772.
- [17] "PTMF: A privacy threat modeling framework for IoT with expert-driven threat propagation analysis," *arXiv preprint*, 2025, arXiv:2510.21601.