Balancing Security and Safety: A Look at Emergency Access Solutions for Implantable Medical Devices (IMDs)

Austin James[1], Veronica Schmitt[2], Janine Medina[3], Danny "Danhammer" Hetzel[3,4], Lucas Potter[5] and Xavier Palmer[5]

[1]University of East London, UK
[2]Noroff University, Norway
[3]Biohacking Village, USA
[4]Husaria Security, USA
[5]BiosView Labs, USA

hello@theaustinjames.com

Abstract: Wireless implantable medical devices (IMDs) have transformed clinical care but have also created a cyber-physical attack surface in which security failures can cause direct physiological harm. This tension is most acute in emergency access: clinicians must rapidly interrogate or reprogram devices under conditions that invalidate normal authentication. Security measures can delay critical care; mechanisms guaranteeing rapid access introduce vulnerabilities. This paper offers a critical review treating IMD emergency access as a safety-critical socio-technical exception regime, not merely a technical backdoor. We synthesise major access mechanisms—proximity triggers, break-glass credentials, external mediators, and manufacturer-mediated approaches—and analyse them through threat models, safety failure modes, human factors, and governance structures. We emphasise how cognitive load, workflow disruption, and long device lifecycles amplify risk during urgent interventions. The analysis shows that these mechanisms redistribute risk among patients, clinicians, manufacturers, and regulators, and that both security and safety failures can be catastrophic. We argue that no single mechanism can simultaneously maximise safety, security, and accountability across diverse contexts. We conclude by specifying the governance interventions required to make emergency access defensible in practice.

Keywords: Implantable medical devices, IMD, Emergency access, Cyber-physical security, Safety-critical systems, Break-glass mechanisms, Governance and risk redistribution

1. Introduction

Software-driven IMDs, such as pacemakers, ICDs, and neurostimulators, now rely heavily on wireless communication. This shift has improved therapy and follow-up care but has also expanded the cyber-physical attack surface in ways that directly affect patient safety (Halperin et al., 2008; James et al., 2025; Straw, Ashworth, and Radford, 2022). A central unresolved challenge is emergency access. In acute situations—malignant arrhythmias, device malfunctions, or therapy mis-delivery—clinicians may need to immediately stop, interrogate, or reprogram an ICD. Yet security controls designed to prevent unauthorised manipulation can delay these lifesaving interventions, while mechanisms that guarantee rapid access introduce new vulnerabilities (Denning et al., 2009; Tandon et al., 2024).

Early research demonstrated that assumptions about "security through proximity" or proprietary interfaces were insufficient, showing that unauthorised access was feasible under realistic conditions (Halperin et al., 2008). Finkle (2016) reports that cybersecurity researcher Jay Radcliffe has identified vulnerabilities in the Animas OneTouch Ping insulin pump that could allow an attacker to spoof wireless communications and trigger unauthorised insulin delivery. Innes (2019) also reports that Arizona clinicians have begun preparing for the possibility of cyberattacks on medical devices, highlighting the clinical implications of compromised pacemakers and other implantables. Subsequent work highlighted how IMD resource constraints, long lifecycles, and dependence on external programmers complicate the application of conventional cybersecurity controls (Denning et al., 2009; Marinescu et al., 2025). Schmitt and Alli (2020) argue that healthcare cybersecurity failures often stem from the application of generic security models in highly specialised clinical environments. More recent analyses broaden the threat landscape to include insider misuse, governance gaps, and lifecycle degradation across decades of deployment (James et al., 2025). Despite these findings, emergency access is still typically treated as a secondary feature embedded within broader security proposals, which has the practical effect of framing it as a security–usability optimisation problem rather than as a distinct safety-critical challenge.

This framing is inadequate. Evidence from digital health and human-factors research shows that emergency care occurs under extreme cognitive load, time pressure, and incomplete information, in which even well-designed controls can fail unpredictably (Tandon et al., 2024; Sendelbach and Funk, 2013; Bonney, 2011,

p. 240; Sutton et al., 2020). In these conditions, security failures (e.g., blocked access) and safety failures (e.g., misconfiguration) are clinically indistinguishable. Moreover, decisions about who may override an IMD in an emergency embed implicit assumptions about authority, responsibility, and acceptable harm. Governance research demonstrates that emergency access mechanisms redistribute risk among patients, clinicians, manufacturers, and regulators rather than eliminating it (Beckers et al., 2016; Ison et al., 2025; Díaz-Martínez et al., 2025).

This paper, therefore, treats emergency access not as a usability exception or a technical backdoor but as a socio-technical exception regime within a safety-critical cyber-physical system. By synthesising technical mechanisms, failure modes, realistic threat assumptions, and governance constraints, the review clarifies why no single mechanism can simultaneously maximise safety, security, and accountability across diverse clinical and regulatory contexts. Rather than proposing an optimal solution, the analysis identifies the structural tensions that shape emergency access and argues that governance and standards must drive future progress as much as device-level innovation.

2. Methodology

This review employs a critical narrative synthesis to evaluate the emergency access landscape for IMDs, integrating principles from technical security engineering, human factors research, and governance analysis. The methodology recognises emergency access as a sociotechnical exception within a regime in which technical performance is inseparable from human and environmental variables.

2.1 Search Strategy and Corpus Construction

A systematic search was conducted across Scopus and Google Scholar. The core search string targeted the intersection of IMDs, access control, emergency use, and security–safety concerns: (“implantable medical device” AND (“access control” OR “authentication”)) AND (“emergency” OR “break-glass”) AND (“security” OR “safety”). Supplementary searches incorporated device-specific terms (e.g., 'pacemaker', 'ICD', 'insulin pump'). The timeframe spanned 2008–2025 to reflect developments following Halperin et al.’s seminal work. The aim is to establish a baseline of IMD-specific security literature; the review intentionally draws on insights from adjacent domains. This inclusion compensates for the relative scarcity of empirical studies specifically addressing IMD-human-factors interactions in emergency scenarios by extrapolating established principles from digital health and safety-critical cyber-physical systems.

2.2 Analytical Framework: Thematic Coding

A qualitative thematic lens was applied to the corpus, focusing on how technical mechanisms interact with operational constraints. Rather than relying on singular experimental data, claims regarding clinical workflow and cognitive load are derived from cross-disciplinary thematic patterns:

Technical Constraints: Evaluation of mechanism feasibility based on IMD power and computational limits.

Human Performance: Extrapolation of error probabilities under stress from broader digital health studies.

Structural Tensions: Identification of irreducible trade-offs where security controls measurably increase time-to-intervention.

2.3 Evaluative Scope and Theoretical Deductions

As a critical perspective piece, this methodology does not claim exhaustive systematic coverage of all IMD security literature. Instead, it uses a socio-technical evaluation to identify structural tensions—such as the "security-safety inversion"—treating them as inherent properties of complex systems. Claims regarding clinical workflow disruption and the impact of cognitive load are presented as theoretically grounded deductions synthesised from technical constraints and human-factors principles. This approach moves the analysis beyond descriptive summary to specify the governance interventions required for defensible emergency access.

3. System Boundaries and Definitions

Clear system boundaries prevent conflation of technical, clinical, and governance challenges. This section defines the scope and key terms used throughout the analysis.

3.1 Implantable Medical Devices (IMDs)

Jeremy and Nicole (2010) define IMDS as devices that are permanently or semi-permanently implanted in a patient to treat an underlying medical condition. This category includes cardiac implants, neurostimulators, and implantable drug delivery systems. These devices share several defining characteristics: constrained computational and energy resources, reliance on external programmers or controllers, long operational lifetimes, and deployment within complex physiological environments (Maisel and Kohno, 2010; Marinescu et al., 2025). Wearable, external, or intermittently attached medical devices are explicitly excluded, even when they interface with implanted components. Although such systems raise important cybersecurity and usability concerns, their failure modes, consent assumptions, and access pathways differ fundamentally from those of permanently implanted devices. These exclusions ensure that the analysis remains grounded in the structural constraints of IMDs rather than drifting into speculative or non-implant contexts, where emergency access is based on fundamentally different assumptions.

3.2 Emergency Access

Emergency access in implantable medical devices can be understood as a set of capability-limited interaction pathways among the clinician, the device, and the supporting infrastructure operating under conditions of uncertainty and degraded assumptions. These pathways are instantiated through various concrete mechanisms—such as magnet responses, programmer-mediated sessions, or telemetry-driven alerts—that expose only a bounded subset of device functionality in ways that depend on device type, manufacturer-specific design, and clinical environment.

Let emergency access in implantable medical devices be represented as a family of time-critical, capability-limited interaction pathways:

$\prod = \{\pi_1, \pi_2, \ldots, \pi_n\}$

Each pathway $\pi \in \prod$ is characterised by:

$\varsigma \in \varsigma_{set}$: Triggering conditions, such as magnet placement, programmer invocation, or contextual inference based on physiological or environmental signals.

$\tau \in \mathbb{T}$: Access latency, defined as the time required to achieve a clinically meaningful action.

$s \in \mathbb{S}$: Scope of accessible functionality, capturing whether the pathway permits, for example, read-only interrogation, therapy suspension, or constrained reprogramming.

$\mathcal{R} \in \mathbb{R}$: Recoverability, understood as the ability to re-establish a safe and controlled system state following intervention under conditions of uncertainty and degraded assumptions.

Let B_N denote the behaviour set available during normal operation, and B_{AUX} denote auxiliary mechanism-specific behaviours (e.g., fixed-rate pacing under magnet application). Each emergency pathway π exposes a behaviour subset:

$B_\pi \subseteq B_N \cup B_{AUX}$

Emergency access analysis then focuses on the structure of $\prod$ and the distribution of (ς, $\mathbb{T}$, $\mathbb{S}$, $\mathbb{R}$) across devices, mechanisms, and clinical contexts.

Framing emergency access in terms of these pathways foregrounds access latency, functional scope, and recoverability as central design and governance dimensions and captures the heterogeneous, partial, and human-mediated nature of emergency interactions with implantable medical devices.

3.3 Security and Safety

'Security' refers to the protection of IMDs against unauthorised access, manipulation, or disruption, whether malicious or accidental. "Safety" refers to avoiding an unacceptable risk of physical harm to the patient that arises from device behaviour, malfunctions, or delayed interventions. In IMDs, these concepts are inseparable: a security control that delays access can directly cause physiological harm, while a safety-orientated override

can introduce exploitable vulnerabilities or pathways to misuse (Halperin et al., 2008). This review treats security and safety as coequal, interacting risk domains, rather than as independently optimisable objectives. To move beyond purely descriptive treatments, we formalised the security-to-safety relationship as a minimal, three-variable system. Let any emergency access mechanism be represented as a tuple:

$M = (S_1, S_2, S_3)$,

where:

S_1 (Security strength) denotes resistance to unauthorised access or manipulation.

S_2 (Safety responsiveness) denotes the time required to achieve successful clinical intervention, and

S_3 (Accountability clarity) denotes the degree to which actions are attributable, auditable, and reviewable.

Under this formulation, emergency access mechanisms can be compared not as discrete implementations but as positions within a constrained trade-off. Empirical and theoretical work across IMD systems suggests that these variables are not functionally independent. In particular, for all known mechanism classes, increasing S_1 beyond a certain threshold results in S_2 degradation under emergency conditions.

This formulation does not claim a strict mathematical optimality range; instead, it defines a testable system property: mechanisms that strengthen authentication, validation, or control layers beyond a given threshold introduce latency, complexity, or dependencies that measurably increase time-to-intervention. This model provides a basis for analysing emergency access as a constrained optimisation problem with clinically bounded tolerances rather than as a purely technical design challenge.

3.4 System Boundary Exclusions

To maintain analytical rigour and focus, this review explicitly excludes several domains: namely, artificial intelligence applications in diagnostic or control systems; population-scale, remote-only cyberattack scenarios lacking immediate physical or clinical consequences; the design or benchmarking of novel cryptographic protocols; and any empirical performance evaluation of specific commercial devices beyond established case studies. These deliberate exclusions focus the paper's findings on the fundamental structural and governance challenges inherent in securing emergency access to medical technology, rather than diverting attention to speculative technological solutions or complex engineering details.

3.5 Stakeholders and Responsibility

Emergency access mechanisms inherently redistribute risk and responsibility among patients, clinicians, manufacturers, and regulators, not because clinical roles or training differ, but because each mechanism embeds implicit assumptions about who is trusted and what actions are permissible. For example, a proximity-triggered override implicitly assigns risk to the patient by treating anyone physically close to them as an authorised actor. Governance and socio-technical systems research shows that such redistributions are a structural property of safety-critical systems rather than a consequence of any particular design choice (Beckers et al., 2016; Ison et al., 2025). Making these allocations explicit provides the foundation for the taxonomies, threat models, and failure mode analyses that follow.

4. Issues in the Technical Foundations for IMD Security

The design and evaluation of any emergency access mechanism for IMDS must be grounded in an understanding of its inherent technical constraints and the cryptographic primitives available to engineers. These factors fundamentally bound the solution space and defined the core security-safety tension.

4.1 The Primacy of Resource Constraints

IMDs are examples of systems that have very few resources. Their defining characteristics—miniaturisation, a non-replaceable battery lifespan of 5–15 years, and limited computational power and memory—create a design paradigm where security cannot be an afterthought but must be meticulously balanced with device longevity and core therapeutic functions. Traditional, resource-intensive cryptographic operations, such as standard asymmetric encryption methods like RSA, are often prohibitively power- and computation-intensive for direct implementation on the implant itself (Samaraweera et al., 2025). Consequently, a robust security scheme incorporating authentication, encryption, and authorisation may not be feasible if it significantly degrades battery life or delays life-critical operations (Camara, Peris-Lopez, & Tapiador, 2015). This limitation is

not merely an engineering challenge but also a direct security vulnerability; adversaries can exploit power-hungry protocols to launch Denial-of-Service (DoS) attacks that rapidly deplete the IMD's battery.

4.2 Evolution of Cryptographic Strategies for IMDs

In response to these constraints, the field has evolved from non-existent or weak security to a suite of lightweight, context-aware cryptographic strategies. Research has moved beyond simple, static keys toward more dynamic and resilient foundations:

Lightweight Cryptography: The imperative is for "lightweight security mechanisms" tailored for devices (Sun, Lo and Lo, 2017). This includes optimised implementations of standards such as the Low-Power Advanced Encryption Standard (AES) for Implantable Cardiac Devices and the use of elliptic-curve cryptography (ECC), which offers stronger security per bit than RSA, making it more suitable for constrained environments (Liu, Richardson, and Van der Spiegel, 2021).

Physiological Key Generation: A significant innovation is the use of the patient's own body as a source of security. By deriving cryptographic keys from unique, time-variant physiological signals such as the electrocardiogram (ECG) or photo-plethysmogram (PPG), these schemes enable secure access without pre-shared secrets (Karimian et al., 2017; Rahimi M.S., 2021). Techniques such as fuzzy vault and fuzzy commitment are employed to handle natural signal variability, allowing a legitimate programmer in proximity to generate the same key as the IMD for authentication and session establishment.

Proximity and Context as Security Parameters: Cryptographic protocols are increasingly integrated with physical and contextual safeguards. Distance-bounding protocols measure the signal's time of flight to verify that a programmer is physically near the patient, thereby mitigating relay or man-in-the-middle attacks (Camara, Peris-Lopez and Tapiador, 2015; Tabassum et. al., 2018). Furthermore, the concept of "context-aware security modes" proposes dynamically adjusting security postures – for example, relaxing authentication in a verified emergency room setting, which requires cryptographic systems that can support such policy shifts (Riegler, Sametinger, and Rozenblit, 2022).

These technical constraints explain why 'ideal' security solutions are impractical and why the trade-offs that follow are structural, not merely design choices

5. Emergency Access Mechanism Taxonomy

This section classifies emergency access mechanisms by their trust anchors and operational assumptions, following recent research on cyber-physical systems and governance (Beckers et al., 2016).

5.1 Physical Proximity-Based Mechanisms

Physical proximity mechanisms enable emergency access by allowing interaction with the implanted device from a short distance, typically via magnets, inductive coupling, or short-range electromagnetic triggers. Activation places the device into a predefined emergency mode, such as therapy suspension or simplified programming access. These mechanisms rely on physical presence as the primary trust anchor, implicitly assuming that those in close proximity during emergencies are authorised clinicians. Their principal advantages are minimal latency and independence from external infrastructure, which align with clinical urgency. However, recent analyses emphasise that proximity alone is a weak authorisation signal in crowded or uncontrolled environments, particularly given the severe consequences of accidental or unauthorised activation (Halperin et al., 2008; James et al., 2025).

5.2 Credential-based “Break-Glass” Access

Credential-based break-glass mechanisms permit emergency access when specific credentials, override codes, or emergency procedures are invoked. These mechanisms intentionally relax standard access controls under predefined emergency conditions, often accompanied by enhanced logging or post-event review. While this model preserves an explicit notion of authorisation, it also introduces vulnerabilities associated with credential availability, memorability under stress, and long-term exposure over the multi-decade lifetimes of IMDs. Human-factors research indicates that reliance on rarely used credentials during emergencies increases error probability and delays intervention, undermining safety objectives even when security controls function as designed (Denning et al., 2009; Tandon et al., 2024).

5.3 External Mediator Devices

External mediator approaches introduce an intermediary device, such as a programmer, shield, or policy-enforcing controller, that brokers emergency access between clinicians and the IMD. The mediator may authenticate users, enforce access policies, or temporarily relax constraints while maintaining an audit trail. This class delegates trust to the mediator, leveraging its greater computational capacity and flexibility relative to the implant itself. While such separation can reduce on-device complexity and attack surface, it introduces a dependence on the availability, integrity, and proper configuration of external devices, which may be absent or compromised in emergency settings. Cyber-physical systems research highlights that mediator failures often manifest as access denial rather than graceful degradation, amplifying safety risk (Marinescu et al., 2025).

5.4 Manufacturer-Mediated Emergency Access

Manufacturer-mediated mechanisms rely on device vendors to authorise or enable emergency access—possibly through proprietary tools, remote authentication, or emergency support channels. The manufacturer, acting as an external authority, centralises trust. While these methods can facilitate precise control and traceability, they also introduce latency, connectivity dependence, and jurisdictional complexity, which are ill-suited to acute clinical care. Recent governance-focused reviews indicate that manufacturer-mediated access introduces unresolved issues regarding liability, clinician autonomy, and equitable availability across healthcare systems (James et al., 2025).

5.5 Context-Conditioned Access Mechanisms

Context-conditioned mechanisms infer authority based on environmental or situational cues, such as location, time, or detected physiological distress. Access may be granted automatically or semi-automatically when predefined conditions are met. These mechanisms shift trust from explicit actors to contextual inference, reducing clinician burden but introducing risk from false positives and false negatives. Evidence from digital health and human-centred systems research suggests that contextual signals are often ambiguous or noisy during emergencies, making an incorrect inference a significant safety hazard rather than a marginal edge (Tandon et al., 2024).

5.6 Cross-cutting Taxonomic Dimensions

Emergency access mechanisms across different classes of connected devices, including implantable medical devices, are fundamentally characterised by four critical, interdependent dimensions: the trigger modality (how access is initiated, e.g., physical action, credential, or inferred context), the primary trust anchor (what establishes confidence in the access attempt, e.g., proximity, identity, or vendor authority), the degree of infrastructure dependence (whether the system operates locally or requires external network support), and the inherent reversibility and auditability of emergency actions. While technological configurations across these dimensions are numerous, a key finding from recent systems and governance research is that no single design eliminates risk; rather, each unique combination reallocates the inevitable risks associated with potential misuse or failure modes across various stakeholders, including patients, clinicians, manufacturers, and regulators.

5.7 Implications of the Taxonomy

This revised taxonomy underscores that emergency access mechanisms are not interchangeable technical options but embodiments of distinct assumptions about trust, urgency, and responsibility. By stabilising these categories before evaluation, the taxonomy provides a structured basis for the threat modelling and safety failure analyses that follow without implying that any single class offers a universally defensible solution.

6. Threat Modelling and Adversarial Assumptions

A realistic threat model for IMD emergency scenarios must extend beyond remote hacking. Consequential harm often stems from non-malicious failures and operational stress, not just adversarial sophistication (James et al., 2025). This section evaluates how these mechanisms fail under realistic adversarial and clinical conditions. Specifically, the critical risks to emergency access mechanisms are multi-faceted: Local proximity attackers can exploit short-range channels, demonstrating that physical presence is not a sufficient security boundary, especially when emergency protocols mandate permissive access (Halperin et al., 2008; Rasmussen et al., 2009). The more dominant threats relate to human factors and long-term deployments, including insider misuse and overreach, in which authorised personnel, under acute stress and cognitive load, inadvertently misconfigure devices, resulting in outcomes equivalent to a malicious act (Tandon et al., 2024; Sendelbach &

Funk, 2013; Bonney, 2011, p. 240; Sutton et al., 2020; Gündoğan & Erdağı Oral, 2023; Jämsä et al., 2021; Sarioğlu & Amarat, 2025). This vulnerability is compounded by accidental and stress-induced misuse, in which degradation in human performance under high-pressure conditions increases the likelihood of procedural errors (Tandon et al., 2024). Furthermore, the decades-long service life of IMDs introduces significant lifecycle and supply-chain degradation, making vulnerabilities due to outdated cryptography or poor key management a more dominant long-term concern than novel attack vectors (James et al., 2025). Consequently, effective policy and design must prioritise resilience against predictable human error, misaligned trust assumptions, and operational drift, rather than focusing predominantly on technical measures against less probable, high-sophistication remote cyberattacks.

6.1 Implications for Emergency Access Design

Under this revised threat model, the dominant risks to emergency access mechanisms arise from misplaced trust assumptions, stress-induced misuse, and lifecycle drift, not from adversaries with exceptional technical capabilities. Consequently, mechanisms that prioritise strong authentication or cryptographic complexity without addressing emergency usability and governance are unlikely to reduce net damages. These insights motivate the safety failure-mode analysis that follows.

7. Safety Failure Modes in Emergency Access

The IMD safety policy must integrate evidence from digital health, cyberphysical systems, and human-centred design, explicitly treating safety failures as co-equal determinants of patient harm along with security risks. Emergency IMD access is a safety-critical operation that occurs under significant stress, time pressure, and incomplete information. These are all conditions known to increase the likelihood of mistakes in people and systems, especially when they are not following their usual workflows (Tabassum et. al. 2018; Tandon et al., 2024). Security must not compromise robustness and usability in emergencies; an unusable, technically secure system is a direct safety hazard. Two dominant safety failure modes are access denial/latency and misconfiguration. Denials or delays in interventions can result from credential-, mediator-, or manufacturer-mediated access failing due to unavailable or slow credentials, intermediary devices, or external authorisations. Cyber-physical systems research shows that control-layer latency often causes outright failure, which is critical for IMDs, where seconds matter. These failures can occur even if security controls function, exposing a structural incompatibility between security assurances and emergency safety. Rapid parameter alteration in emergencies risks misconfiguration under stress. Research on human-centred systems suggests that unfamiliar interfaces or coarse-grained controls raise the likelihood of errors, a situation worsened by the variety of proprietary IMD interfaces (Beckers et al. 2016).

Usability failures and workflow disruption degrade safety. Complex authentication, precise timing, or unfamiliar tools are functionally equivalent to access denial. Digital health studies confirm that workflow disruption—not technical capability—dominates adverse outcomes in emergency digital interventions (Tandon et al., 2024). This evidence challenges the assumption of infrequent, flawless emergency use. Ethical risks regarding consent and autonomy arise when access overrides patient preferences. Poorly specified policies erode trust, a safety risk that systems research confirms affects acceptance and adherence (Beckers et al., 2016).

Summarily, these failure modes show that tighter security alone cannot ensure safe IMD emergency access. Any measure that blocks unauthorised use but adds delay or ambiguity increases patient harm. A shift from security-only fixes to systemic governance, accountability, lifecycle responsibility, and risk redistribution is essential to preserve institutional trust in these safety-critical systems.

Figure 1 and Table 1 illustrate the irreducible trade-off between access latency and security exposure. No mechanism can optimise both simultaneously; the curve represents a structural security-safety inversion. The central governance question is where to position mechanisms along this frontier in high-stakes contexts.

Table 1: Mapping mechanisms to their failure modes, dependencies and trust assumptions

MECHANISM	PATIENT SAFETY (Y-AXIS)	SYSTEM SECURITY (X-AXIS)	TRADE-OFF PROFILE	NOTES
Magnet-triggered override	Very High	Very Low	Top-left corner	Instant access; vulnerable to physical tampering or accidental activation
Proximity-based triggers	High	Low	Left of curve	Fast and intuitive; lacks authentication or audit trail
Break-glass protocols	Moderate	Moderate	Mid-curve	Requires clinician action; may log override, but still bypasses full authentication
Manufacturer-mediated access	Low	High	Bottom-right curve	Secure and auditable; slow and dependent on external coordination
Cryptographic authentication	Very Low	Very High	Far-right curve	Strongest security; impractical in emergencies without preloaded credentials
Pre-authorised clinician override	High	Moderate	Upper-mid curve	Balances speed and accountability; depends on trust in local personnel
Biometric triggers (experimental)	Moderate	High	Right of curve	Promising but unproven; latency and reliability concerns in crisis

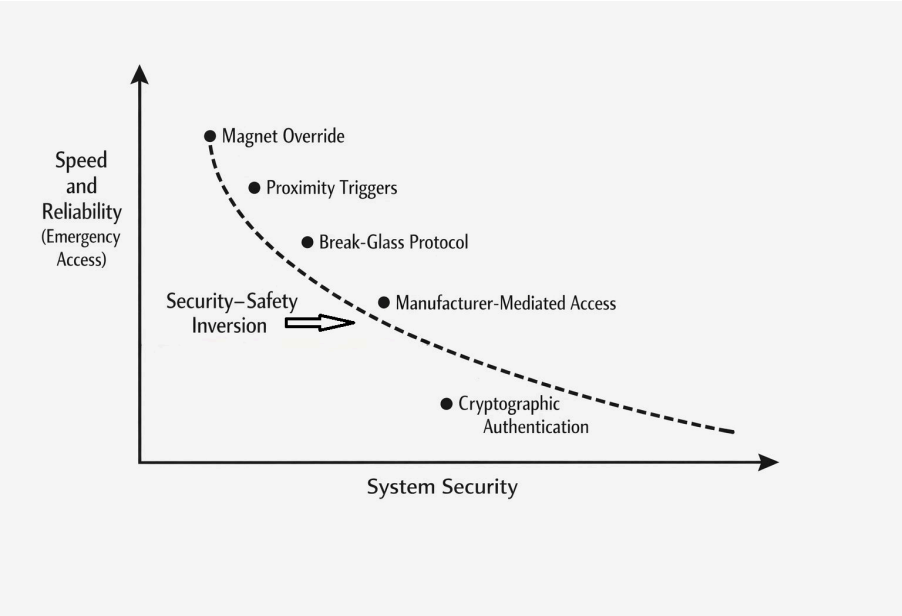

Figure 1: Pareto frontier diagram showing the irreducible structural tensions forcing trade-offs across essential dimensions of IMD operations

Together, Figure 1.0 and Table 1 show that these trade-offs are structural rather than incidental. They arise from the interaction of human performance under stress, IMD resource constraints, and governance expectations. As a result, emergency access design is not a search for optimal points but rather a process of selecting which risks are tolerable, for whom, and under what clinical and organisational conditions.

8. Governance, Accountability, and Risk Redistribution

Emergency access failures are governance gaps, not cryptographic deficiencies. Access denials arise when frameworks fail to mandate emergency access standards; misconfigurations reflect a lack of interface harmonisation and human-factor validation; lifecycle drift stems from weak manufacturer obligations; ambiguous authority reveals a failure to specify who acts under what conditions (Beckers et al., 2016). These are institutional, not technical, failures.

Emergency access is an exception regime that requires normative judgements about who acts, when, and with what consequences (Ison et al., 2025). Design choices encode governance decisions that policy must manage.

Authority allocation during emergencies directly affects safety. Clinician-centric models enable rapid action but impose liability; manufacturer-centric models offer traceability but introduce delays; hybrid models create ambiguity. Clear precedence rules are essential to prevent hesitation (James et al., 2025).

Each mechanism redistributes residual risk: proximity shifts risk to patients; break-glass transfers liability to clinicians; manufacturer mediation increases vendor complexity; context conditioning embeds risk in opaque algorithms. Explicit risk redistribution—not technical optimisation—is the core policy decision (Beckers et al., 2016).

Long device lifecycles and post-hoc logging complicate accountability. Regulatory silence on emergency cybersecurity exceptions incentivises conservative designs that minimise manufacturers' liability at the expense of patient safety (Maisel & Kohno, 2010; James et al., 2025).

Governance frameworks must recognise emergency access as a managed exception regime, aligning legal liability, clinical authority, and regulatory expectations in advance.

9. Comparative Synthesis and Irreducible Tensions

The design and governance of emergency access mechanisms for implantable medical devices (IMDs) prevent convergence on a single optimal technical solution. Instead, comparative synthesis reveals that these systems encode irreducible structural tensions, forcing trade-offs across essential dimensions such as access speed, resistance to malicious use, auditability, and clear governance. This dynamic is a recognised characteristic of tightly coupled safety-critical cyberphysical architectures: optimising one attribute inherently degrades performance in another. Therefore, analysing various emergency access strategies yields not a universally superior technical protocol, but rather a set of “bounded capability envelopes”, in which effectiveness is constrained by fundamental assumptions regarding trust models, urgency, and the assignment of operational responsibility. Any technical proposal that claims universal superiority likely fails to account for these fundamental, interlinked constraints.

A crucial policy insight is the security–safety inversion that occurs under emergency conditions. While controls prioritising security over safety are reasonable during a steady-state operation, their continuous enforcement can become counterproductive during a crisis. Emergency access is infrequent but invoked under extreme circumstances, and the denial or delay of access may pose a greater immediate risk to a patient than the threat of unauthorised use. Recent work in digital health emphasises the dominance of interventional and recovery capacities over pure prevention in crisis management (Tandon et al., 2024; Siddiqi et al., 2020). This inversion undermines conventional cybersecurity evaluation frameworks that emphasise access restriction without adequately accounting for time-critical safety constraints and the real-world operational demands of clinical emergencies.

Furthermore, efforts to manage IMD access inevitably lead to the redistribution of risk rather than its net reduction, with disagreements often rooted in contested views on acceptable risk allocation (Beckers et al., 2016). The core irreducible tensions that persist across all IMD emergency protocols—and those that resist purely technical or artificial intelligence solutions—include the conflict between immediate action and verifiable authorisation; the clash between simplified usability during a crisis and the detailed logging required for auditability; and the friction between centralised security control and necessary decentralised clinical autonomy. The design challenge is not the elimination of risk but the management of inherent tensions within a managed-exception space, where design choices explicitly trade off safety, security, and governance priorities. Future policy directions must focus on establishing explicit frameworks for allocating authority and responsibility, rather than seeking device-level technical optimisation.

10. Conclusion: Limits, Non-Solutions, and Future Directions

Synthesis of Findings

Emergency access to IMDs presents a structural tension: no single technical solution can simultaneously maximise security, safety, and accountability. Designs prioritising security risk fatal delays; those prioritising immediacy increase the potential for misuse. Enhanced technical controls—'perfect' authentication or universal backdoors—are non-solutions. The core constraints are institutional. Without explicit regulatory guidance on emergency exceptions, audit standards, and a clear allocation of responsibilities, proprietary mechanisms will continue to create ambiguity.

Prescriptive Governance Interventions

Addressing the structural tensions identified in this review requires governance interventions that reshape the constraint space rather than attempting to engineer around it. Regulators should establish minimum emergency-access performance standards, including maximum permissible access latency and requirements for predictable behaviour under degraded conditions. Manufacturers must be subject to lifecycle responsibility obligations, ensuring that emergency access remains viable, secure, and auditable throughout the device's multi-decade lifespan. Governance frameworks should mandate post-event auditability, requiring that all emergency overrides—whether successful or failed—are logged, reviewable, and linked to clear accountability structures. Cross-jurisdictional harmonisation is essential for avoiding fragmented expectations that undermine clinician autonomy and patient safety. Finally, procurement and certification processes should incorporate human-factors validation to ensure that emergency access mechanisms can be executed reliably under acute cognitive load and environmental stress.

Closing Statement

These interventions reposition governance as the primary mechanism for managing the unavoidable trade-offs between safety, security, and accountability. Technical innovation remains necessary, but without governance structures that explicitly define acceptable risk, emergency access will continue to fail in predictable and harmful ways.

Acknowledgements

We would like to thank Jay Radcliffe and Dr Jorge Acevedo Canabal for their insights, and the entire Bio Hacking Village for their support throughout this work.

Ethics Declaration: This review did not involve human participants, animals, or personal data and, therefore, did not require ethical approval.

AI Declaration: Grammarly, Microsoft Copilot, and DeepSeek were used to support the preparation of this manuscript. These tools assisted with summarisation and keyword extraction, spelling and grammar refinement, and reference organisation. The authors reviewed and verified all AI-assisted outputs. No AI system was credited as an author or cited as a source. AI tools did not generate or fabricate any data, citations, or analyses.

References

Beckers, K., Landthaler, J., Matthes, F., Pretschner, A. and Waltl, B., 2016, June. Data accountability in socio-technical systems. In International Conference on Business Process Modeling, Development and Support (pp. 335-348). Cham: Springer International Publishing.

Bonney, W. (2011) Impacts and Risks of Adopting Clinical Decision Support Systems. Efficient Decision Support Systems – Practice and Challenges in Biomedical Related Domain. InTeh. http://dx.doi.org/10.5772/16265

Camara, C., Peris-Lopez, P., & Tapiador, J. E. (2015). Security and privacy issues in implantable medical devices: A comprehensive survey. *Journal of biomedical informatics*, *55*, 272–289. https://doi.org/10.1016/j.jbi.2015.04.007

Denning, T., Matsuoka, Y. and Kohno, T., 2009. Neurosecurity: security and privacy for neural devices. Neurosurgical Focus, 27(1), p.E7.

Finkle, J. (2016) Insulin pump vulnerable to hacking. Scientific American. Available at: https://www.scientificamerican.com/article/insulin-pump-vulnerable-to-hacking/ (Accessed: 6 April 2026).

Gündoğan, G. & Erdağı Oral, S. (2023) The effects of alarm fatigue on the tendency to make medical errors in nurses working in intensive care units. Nursing in Critical Care, 28(6), pp. 996–1003. doi: 10.1111/nicc.12969

Halperin, D., Heydt-Benjamin, T.S., Ransford, B., Clark, S.S., Defend, B., Morgan, W., Fu, K., Kohno, T. and Maisel, W.H. (2008) 'Pacemakers and implantable cardiac defibrillators: software radio attacks and zero-power defenses', Proceedings of the IEEE Symposium on Security and Privacy, pp. 129–142. https://doi.org/10.1109/SP.2008.31

Innes, S. (2019) Can medical devices be hacked? Arizona doctors prepare for possibility of cyberattack. SimGHOSTS News & Press. Available at: https://simghosts.org/news/508619/Can-Medical-Devices-be-Hacked-Arizona-Doctors-Prepare-for-Possibility-of-Cyberattack.htm (Accessed: 6 April 2026).

Ison, R., Buckle, P., Nguyen, N., Preiser, R., Vandenbroeck, P. and Klein, L., 2025. Reimagining Systems Thinking as Cybersystemic Researching: An Invitation to a Cyber-Systemic Co-Inquiry. Systems Research and Behavioral Science, 42(5), pp. Pp.1391-1403.

Jämsä J.O., Uutela K.H., Tapper A-M, Lehtonen L. Clinical alarms and alarm fatigue in a University Hospital Emergency Department—A retrospective data analysis. Acta Anaesthesiol Scand. 2021;65:979–985. https://doi.org/10.1111/aas.13824

James, A., Potter, L. and Palmer, X.-L. (2025) 'Understanding cybersecurity threats to implantable medical devices: a review', Proceedings of the 24th European Conference on Cyber Warfare and Security (ECCWS), pp. 230–240. https://doi.org/10.34190/eccws.24.1.3605

Jeremy A.H. and Nicole M.H. (2010). A taxonomy of vulnerabilities in implantable medical devices. In Proceedings of the second annual workshop on Security and privacy in medical and home-care systems (SPIMACS '10). Association for Computing Machinery, New York, NY, USA, 13–20. https://doi.org/10.1145/1866914.1866917

Karimian, N., Guo, Z., Tehranipoor, M. and Forte, D. (2017) 'Highly reliable key generation from electrocardiogram (ECG)', IEEE Transactions on Biomedical Engineering, 64(6), pp. 1400–1411. doi: 10.1109/TBME.2016.2607020.

Liu, Xilin & Richardson, Andrew & Van der Spiegel, Jan. (2020). An Energy-efficient Wireless Neural Recording System with Compressed Sensing and Encryption. 10.48550/arXiv.2009.06532.

Maisel, W.H. and Kohno, T., 2010. Improving the security and privacy of implantable medical devices. New England Journal of Medicine, 362(13), pp. 1164-1166.

Marinescu, M.R., Ionescu, O.N., Pachiu, C.I., Dinescu, M.A., Muller, R. and Şuchea, M.P., 2025. Next-Gen Healthcare Devices: Evolution of MEMS and BioMEMS in the Era of the Internet of Bodies for Personalized Medicine. Micromachines, 16(10), p.1182.

Rahimi Moosavi, S. (2021) 'PPG-KeyGen: Using photoplethysmogram for key generation in wearable devices', Procedia Computer Science, 184, pp. 291–298. doi: 10.1016/j.procs.2021.03.038.

Rasmussen, K.B., Castelluccia, C., Heydt-Benjamin, T.S. and Capkun, S., 2009, November. Proximity-based access control for implantable medical devices. In Proceedings of the 16th ACM conference on computer and communications security (pp. 410-419).

Riegler, Michael, Sametinger, Johannes & Rozenblit, Jerzy. (2022). Context-Aware Security Modes For Medical Devices. 372-382. 10.23919/ANNSIM55834.2022.9859283.

Samaraweera C., Bhattarai A., Peng D., Sharif H., Liu Y. and Wang H. (2025) "Differentiated Communication Strategies for Remote Electrocardiogram Monitoring With a Multilevel Blockchain System," in IEEE Internet of Things Journal, vol. 12, no. 15, pp. 30507-30517, 1 Aug.1, 2025, doi: 10.1109/JIOT.2025.3571794.

Sarioğlu E, Amarat M. The Relationship Between Alarm Fatigue and Medical Error Tendency in Intensive Care Unit Nurses: The Mediating Affect of Role Overload. Nurs Crit Care. 2025 Jul;30(4):e70121. doi: 10.1111/nicc.70121. PMID: 40704558; PMCID: PMC12288113.

Siddiqi, M. A., Doerr, C., & Strydis, C. (2020). IMDfence: Architecting a Secure Protocol for Implantable Medical Devices. *IEEE Access*, *8*, 147948–147964. https://doi.org/10.1109/ACCESS.2020.301568

Schmitt, V. & Alli, N. (2020) Why are we getting it so wrong in healthcare and medical device security? Available at: https://veronica-schmitt.com/2020/07/23/why-are-we-getting-it-so-wrong-in-healthcare-and-medical-device-security/ (Accessed: 6 April 2026)

Sendelbach, S. & Funk, M. (2013) Alarm fatigue: a patient safety concern. AACN Advanced Critical Care, 24(4), pp. 378–386.

Straw I, Ashworth C, Radford N. When brain devices go wrong: a patient with a malfunctioning deep brain stimulator (DBS) presents to the emergency department. BMJ Case Rep. 2022 Dec 26;15(12):e252305. doi: 10.1136/bcr-2022-252305. PMID: 36572446; PMCID: PMC9806045.

Sun, Y., Lo, F.P.-W. and Lo, B. (2017) 'Security and privacy for the Internet of Medical Things enabled healthcare systems: A survey', IEEE Access, 5, pp. 16777–16797.

Sutton, R.T., et al. (2020) An overview of clinical decision support systems: benefits, risks, and strategies for success. NPJ Digital Medicine, 3:17.

Tabassum, A., Safi, Z., AlKhater, W. and Shikfa, A. (2018) 'Cybersecurity Issues in Implanted Medical Devices', 2018 International Conference on Computer and Applications (ICCA). Doha, Qatar: Qatar University.

Tandon, A., Cobb, B., Centra, J., Izmailova, E., Manyakov, N.V., McClenahan, S., Patel, S., Sezgin, E., Vairavan, S., Vrijens, B. and Bakker, J.P., 2024. Human factors, human-centered design, and usability of sensor-based digital health technologies: a scoping review. Journal of Medical Internet Research, 26, p.e57628.