A Moonshot for Trustworthy Medical Software Updates Using Automated Insulin Delivery Systems as a Proving Ground

Josiah Dykstra
Designer Security, LLC
Ellicott City, MD
josiah@designersecurity.com

Shannon Lantzy
Shannon Lantzy, LLC
Silver Spring, MD
shannon@shannonlantzy.com

Eugene Y. Vasserman
Kansas State University
Manhattan, KS
eyv@ksu.edu

***Abstract*—Current medical device software update cycles require months to complete, creating fundamental incompatibilities with modern cybersecurity threats that emerge within hours and leave life-critical devices vulnerable to rapidly evolving attacks. We address this temporal mismatch by proposing a cybersecurity moonshot initiative focused on automated insulin delivery (AID) technologies that could compress update cycles from months to days while maintaining regulatory compliance and patient safety standards. Through systematic analysis of AID ecosystems using the National Academies' ten Cyber Hard Problems [1], we identify unique constraints and opportunities where targeted cybersecurity innovation may be especially feasible and impactful.**

We propose an initiative to develop and field a development and testing pipeline to drastically accelerate AID system patching development through fully- and semi-automated verification engines, pre-validated testing frameworks approved by the U.S. Food and Drug Administration as Medical Device Development Tools, and digital twin validation networks. This work suggests a path to rapid updates for life-critical medical devices without compromising safety or regulatory compliance, establishing a replicable, generalizable framework for rapid threat response in cyber-physical systems.

***Index Terms*—cybersecurity, healthcare, type 1 diabetes, automated insulin delivery, medtech, software updates**

1. Introduction

Medical devices represent a critical frontier where cybersecurity innovation can address fundamental challenges in cyber-physical system security. Automated insulin delivery (AID) systems for type 1 diabetes (T1D) management are an excellent candidate for a cybersecurity "moonshot" initiative: the technologies are robust, and existing ecosystems possess the stakeholder alignment and regulatory clarity necessary to develop transformative cybersecurity solutions. Development toolchains, best practices, and lessons learned from this tightly-scoped exercise can later be generalized, with additional work, to other cyber-physical systems, achieving advances in automated threat response and secure update mechanisms without requiring new legislation, policy shifts, or changes to existing organizational structures.

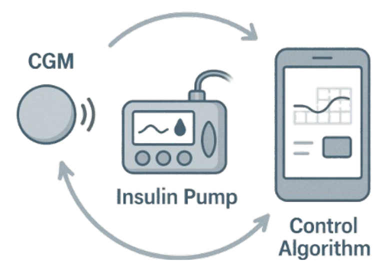

Figure 1. Closed-loop automated insulin delivery (AID) components: continuous glucose monitor (CGM), insulin pump, and control algorithm.

An estimated 1.8 million people in the United States live with type 1 diabetes [2]. Diabetes management and treatment have matured dramatically since insulin was discovered in 1921, and innovation continues to the present day. Digital technologies, particularly for glucose monitoring and insulin delivery, have improved the lives of people with diabetes [3]. Advances in medical technology continues to improve health outcomes while driving economic development [4].

Despite its value, medical technology advances—including in AID systems—often fail to reach the people who need them. This is in part due to the long time needed to develop and test safety-critical systems, and in part due to the difficulty in commercialization [5]. Medical system development and distribution lag behind consumer technology. In one industry survey, 92% of digital health decision makers said that medical device manufacturers "must operate more like consumer technology companies and deliver solutions more quickly" [6].

Among the forefront of innovation in insulin-requiring diabetes care are automated insulin delivery (AID) systems, which combine technologies that measure glucose levels with automatic insulin delivery based on those measurements and the calculations of an algorithm. A simplified system diagram of this control loop is shown in Figure 1. A major reason we call for the use of AID systems as a representative sample are due to their simultaneous simplicity and complexity: an AID system is simple enough that interoperability can be explicitly addressed by because there are few manufacturers controlling the vast majority of market share,[1] and complex enough to admit much more functionality, even to the point where the continuous glucose monitor (CGM) and the insulin pump do not have to be explicitly interoperable, letting an intermediary app (e.g., running on a smartphone) act as a translation layer. This combination is a starting point for testing solutions that can be adapted to both simpler and much more complex systems.

There is a major stumbling block in rapid update development: the critical "software understanding gap" [8], [9], which needs specific challenge problems and benchmarks to resolve. To that end, we use the National Academies of Sciences, Engineering, and Medicine (NASEM)'s 2025 list of Cyber Hard Problems as a framework to systematically analyze the cybersecurity frontier of AID systems and present practical AID solution innovation opportunities [1]. Using these problems as motivation and guidance, we propose an ambitious research program "moonshot" to solve for the cybersecurity fundamental challenges in AID. We suggest that a moonshot focused on AID tech will offer a concrete use case to demonstrate radical improvements. AID systems offer a proving ground for rapid testing of software updates for safety, efficacy, and performance in safety-critical cyber-physical systems, pioneering mechanisms that can be subsequently generalized and expanded to other safety-critical devices and systems. If successful, solving Cyber Hard Problems within the AID system use case will yield incentive-compatible benefits for many stakeholders.

There are six goals for the moonshot: (1) provide increased trust and decreased disease burden for people living with type 1 diabetes; (2) decrease incumbent medical device manufacturers' (MDMs') costs per feature, increasing shareholder value; (3) reduce the barriers for new MDM entrants; (4) increase regulatory review quality while decreasing regulatory burden; (5) provide an exemplar for other medical device technology or critical infrastructure and cyber-physical systems (CPS) to follow; and (6) accelerate development of all life-saving or life-protecting medical devices, so that as soon as technology is available, people can receive it.

The remainder of the paper is organized as follows. Section 2 describes the current state of cybersecurity in diabetes technology and AID. In Section 3, we summarize each of the National Academies' list of Cyber Hard Problems [1] and describe how AID tech currently aligns with them. We outline a moonshot initiative to demonstrate an automated secure software update testing pipeline in AID tech as an exemplar model for other medical devices and systems in Section 4 and conclude in Section 5.

2. Background and Related Work

Current pace of the AID tech development lifecycle. Commercial AID systems have undergone extensive, formal regulatory evaluations and compliance processes as a prerequisite for market entry. These regulatory requirements include secure development, testing, risk assessment, and ongoing maintenance for vulnerabilities [10]. Historically, the development, validation, and distribution of software updates in diabetes technology have been slow and burdensome on manufacturers. One vendor previously had a 9-month test pipeline, a desire for 12 days, and is currently at 90 days [11]. This is in part a result of regulatory evolution based on a historical industry of medical devices as hardware. Until recently, MDMs had a disincentive to make significant updates to software. However, the U.S. Food and Drug Administration's (FDA) updated software and cybersecurity guidance now clearly require software updates. For example, 2016 postmarket cybersecurity guidance asks MDMs to publish the existence of a vulnerability, interim mitigations, and a remediation plan "as soon as possible but no later than 30 days after learning of the vulnerability" and to fix, validate, and distribute the fix "as soon as possible but no later than 60 days after learning of the vulnerability" [12].

Cybersecurity and safety in T1D tech. There is an established history of security researchers evaluating medical technology and finding flaws, including in continuous glucose monitors (CGM) [13] and insulin pumps [14]. For more than a decade, researchers have identified threats and vulnerabilities in the cybersecurity of connected diabetes devices [15]. Demonstrations of vulnerabilities and insecurities have gained the attention of researchers [16]–[18]. Nazzal, Alalfi, and Cordy surveyed techniques for verifying and validating artificial pancreas systems (APS), with a particular focus on model-based and formal methods, to address the critical safety and security challenges posed by these devices [16]. Others have evaluated the reinforcement learning components of APS [19]. In 2025, another team reviewed the security landscape of AID systems, analyzing technical vulnerabilities, legal frameworks, and commercial product considerations, alongside a detailed examination of attack vectors and existing defense mechanisms to identify gaps and propose future research directions for securing AID devices [17]. Complimenting academic scholarship is a robust and growing community of people are building and experimenting with do-it-yourself T1D closed-loop technology, often at a development pace faster than medical device manufacturers can achieve [20].

Commercial technologies such as web browsers have established rapid release cycles which decreases the time software is vulnerable to known bugs. This success is motivational that a similar outcome could be achieved in T1D AID software. Eclipse, a widely used, open-source integrated development environment, moved from annual to quarterly software releases and has found this change to provide more stable bug handling [21]. Google Chrome releases major software updates every six weeks, with minor updates (including security patches) occurring every two to three weeks. Firefox releases major updates every 4 weeks, a steady decline from 16 weeks in 2011 [22]. Minor updates, including security and stability fixes, are released as needed between major releases. Rapid release cycles for life-or-death cyber-physical systems exist as well. Automotive manufacturers have developed the capability of deploying

1. Currently, there are only five major commercial AID systems available in the United States: Tandem Control-IQ+, Medtronic MiniMed 780G, Insulet Omnipod 5, Beta Bionics iLet, and Sequel twiist [7].

software updates over the air, with awareness and clearance from automobile regulators [23].

People with diabetes consider trustworthiness of technology as an important factor for adoption and use [24]. While security is frequently considered eventually, at least one open source project for trustworthy automated insulin insulin delivery has security as a primary design element [25].

To accelerate the development of AID features for people living with type 1 and insulin-requiring type 2 diabetes, the software testing pipeline must be accelerated to the speed of consumer technology without compromising trustworthiness or device performance. Some of the slowness in shipping software features is due to historical regulatory requirements and therefore incumbent MDM processes and systems. The remainder may be attributed to fundamental challenges of securing cyber-physical systems.

3. Alignment with Cyber Hard Problems

In 2023, the U.S. Office of the National Cyber Director (ONCD) sponsored a consensus study at the National Academies of Sciences, Engineering, and Medicine (NASEM) to identify key problems for cyber resiliency [1]. The final report was released in 2025 and included ten "unsolved technical and research problems for which progress toward solution would have a significant impact on the practical security of cyber systems." In this section, we describe aspects of the state of AID tech software updates with regard to each of the NASEM Cyber Hard Problems [1], including how these problems are being considered and addressed.

Each subsection below corresponds to one of the Cyber Hard Problems and offers three parts: (1) a summary of the problem as presented in the NASEM report, (2) a discussion of how this problem manifests and is relevant in the AID technology ecosystem, and (3) the unique opportunities the T1D use case presents for addressing the hard problem.

3.1. Risk Assessment and Trust (CHP1)

CHP1 addresses the fundamental challenge that it is difficult to reliably evaluate the inherent cybersecurity properties of a cyber system, and that residual risks remain even after security engineering is complete. Unlike physical systems where properties like tensile strength can be reliably measured and predicted, cyber systems lack reliable security properties and metrics, making it nearly impossible to assess vulnerabilities directly, predict attack consequences, or establish appropriate trust levels. This problem is compounded by the non-stochastic nature of cyber risk, hidden correlations in supply chains, and the difficulty of measuring security improvements, creating a measurement conundrum where the inability to assess security impairs investment in better security practices.

Relevance to AID Tech: Because AID technology is subject to FDA regulatory oversight, risk assessment is standard practice for commercial devices. While FDA cybersecurity guidance has improved risk assessment practices, AID tech MDMs have varying levels of experience with cyber risk: vulnerabilities in software, exposure via wireless interfaces, supply chain compromises, and more. As the NASEM report emphasizes, risk assessment must encompass vulnerability surfaces, operational consequences of failures (including impacts to the users, e.g., the potential for device errors leading to hypoglycemia, and business impacts such as stock price [26]), and the characteristics of emerging threats like sophisticated adversaries.

Researchers found that threat modeling for medical devices poses unique challenges, particularly when integrating safety considerations [27]. For these critical devices, trustworthiness must be established and maintained over time through transparency, robust performance under stress, and responsive patching when vulnerabilities are discovered. For example, systems like Tandem's Control-IQ and Medtronic's SmartGuard include layered risk mitigations that anticipate both clinical and cyber-physical failure modes [28].

Opportunity in AID Tech: The AID, especially T1D, ecosystem has unique advantages for advancing risk assessment and trustworthiness establishment because it features well-established companies, innovative regulatory environments, highly active user populations, and immediate feedback loops for software testing. Most importantly, cleared medical products represent existing risk tolerance baselines that could provide reference points for developing and validating improved methods of risk assessment and trustworthiness establishment in cyber-physical medical systems.

3.2. Secure Development (CHP2)

CHP2 addresses the problem of insufficient reliable practices for verifying that software and firmware components satisfy their intended security requirements. According to the NASEM report, the core problem is that secure development is perceived as costly, time-consuming, and producing unmeasurable results, creating counter-incentives for development organizations. This is compounded by the difficulty of specifying clear security requirements in the first place, and the need to repeat (potentially lengthy) evaluation processes as systems evolve. Without advances in secure development practices, system developers are forced to accept limits on capability and complexity for systems deployed in consequential circumstances, or conversely, accept security risks when deploying more capable systems.

Relevance to AID Tech: Secure development is a central concern in the design and evolution of AID technologies, which operate as cyber-physical systems (CPS) with life-critical functions. Modern diabetes tech devices are software-driven and rely on wireless protocols, cloud services, and, increasingly, mobile apps on patients' smartphones to provide real-time feedback and automate insulin delivery. This complexity exposes them to cybersecurity threats ranging from unauthorized wireless access to manipulation of firmware. Recent vulnerabilities have illustrated these risks: in 2019, the FDA issued a recall for certain Medtronic MiniMed pumps due to unpatchable flaws in their wireless communication systems [29]. In response to such challenges, manufacturers have begun integrating practices such as cryptographic signing of firmware, secure boot mechanisms, and authenticated over-the-air updates. Tandem's t:slim X2 pump, for example, was the first insulin pump to receive FDA approval for remote software updates—an important milestone in medical device software lifecycle management [30].

Regulatory bodies have increasingly pushed for formal cybersecurity integration, with recent FDA guidance requiring medical device developers to implement a secure product development lifecycle, document threat modeling, produce Software Bills of Materials (SBOMs), and include plans for postmarket patching [10]. These trends reflect an ongoing shift toward treating secure development not as an afterthought, but as a continuous, integral part of the lifecycle of AID technologies—essential for preserving trust and safety in increasingly interconnected systems [31], [32].

Opportunity in AID Tech: AID devices offer several advantages for advancing secure development practices relevant to automated software updates. The continuous, high-frequency interaction between users and devices creates natural opportunities for rapid detection of development flaws that might compromise update mechanisms. Unlike medical devices used sporadically, device issues could be quickly discovered and reported by engaged users. The regulatory framework already requires AID tech manufacturers to demonstrate secure development lifecycles and postmarket update capabilities, providing a foundation for automated approaches without additional compliance burdens. Additional automation could reduce developer burden. The open-source diabetes community has taken an agile approach to secure development demonstrating transparent development practices, with projects like Tidepool Loop demonstrating how public code review and version control can accelerate secure development cycles [33].

3.3. System Composition (CHP3)

CHP3 addresses the challenge of securely integrating larger-scale systems from diverse components and services. The problem is that modern systems are assembled from hundreds or thousands of physical and/or software components—often drawn from different suppliers with varying security levels—by architects and developers with varying knowledge. The goal of secure composition is to enable reliance on separately made security judgments about individual components to support efficient judgments about the composite system. However, there is no comprehensive science of safe composition to guide integration, nor are there generally usable tools to validate compliance with semantic rules that allow safe composition. Success depends on deep technical properties of both the components and the design rules for their combination, making it difficult to achieve aggregate security judgments.

Relevance to AID Tech: A typical AID ecosystem includes CGMs, insulin pumps, smartphones, cloud-based analytics, and mobile applications, each potentially developed, updated, and secured independently. System security and human safety are emergent properties of the entire integrated system, which encompasses hardware, software, machine models and other complex algorithms, and medical applications, rather than properties that can be validated at the level of individual components alone. For example, a vulnerability in a mobile device's Bluetooth interface could undermine an otherwise secure pump [34], [35], while inconsistent data formats or failure modes between a CGM and a control algorithm could introduce additional hazards.

While modular design enables component-level security validation, true interoperability requires dynamic behavioral compatibility that extends beyond simple compositional modularity. The Tandem ControlIQ system's experience with the Dexcom G6 to G7 CGM transition exemplifies this challenge: despite both devices being individually validated and the system's modular architecture, unforeseen interactions between the algorithm and the new CGM's altered noise profile necessitated recalls and fixes, demonstrating that secure composition must account for emergent behaviors that arise from component interactions under real-world operational conditions.

Opportunity in AID Tech: To address these issues, some people in the T1D community have pioneered approaches that embody the principles of secure composition. Open-source projects like OpenAPS [36] and Loop [37] explicitly build composite systems using known, modular components—glucose sensors, insulin pumps, and mobile controllers—with documented interfaces and version-specific integration logic. These developments illustrate a maturing understanding in the T1D domain that secure composition is not just a technical concern but a regulatory and architectural imperative—expanding and reinforcing these collaborations are critical for safe composition in the service of safer and more secure diabetes care.

3.4. Supply Chain (CHP4)

CHP4 addresses the challenge of securely developing and managing large software and hardware systems when components and services come from diverse, often opaque sources. System elements can be sourced from vendors, open-source projects, custom developers, and cloud service providers, each operating under different rules and standards. The fundamental difficulty is that integrating system designers cannot readily and safely leverage this diversity of required components and services. Commercial suppliers generally do not reveal component details to protect trade secrets, making sound security assurance judgments impossible at any stage in the supply chain. Without some degree of transparency, components may rarely be updated to address vulnerabilities, and when updates do occur, extensive testing is required to ensure compatibility with existing system elements.

Relevance to AID Tech: Supply chain security is an increasingly urgent concern in medical technology, where both hardware and software components originate from a wide array of global vendors. Devices are frequently assembled from proprietary components, off-the-shelf hardware like wireless chipsets, third-party off-the-shelf software libraries, and third-party services, including cloud infrastructures. This diversity complicates efforts to ensure end-to-end security, particularly when source code, firmware, or subcomponents are closed or sourced from opaque vendors. Recent research highlights that medical devices often contain outdated or vulnerable libraries [38]. Without clear provenance and lifecycle transparency, vulnerabilities in these supply chains may go unnoticed until exploited.

In response, both regulators and some forward-leaning AID tech developers are beginning to adopt supply chain risk management principles. The FDA's recent premarket guidance mandates inclusion of an SBOM and emphasizes the need for secure update mechanisms, component provenance documentation, and end of life component management [10]. Open-source T1D projects such as OpenAPS [36] improve supply chain visibility by exposing every line of source code and all dependencies to community review. As devices grow more interconnected and reliant on external services, the need to build translucency—not necessarily full transparency—into every layer of the AID supply chain will be essential for maintaining trust and security.

Opportunity in AID Tech: AID devices offer advantages for developing automated patch validation across complex supply chains because they generate continuous operational data that can help detect component-level issues rapidly after updates. The FDA's existing SBOM requirements for AID devices provide structured dependency information that automated patch systems can leverage to assess update compatibility. The active open-source T1D community has demonstrated that transparent dependency management is feasible, while high-frequency user interaction means that supply chain security issues can be discovered and reported quickly, providing essential rapid feedback loops for validating automated patch deployment strategies.

In general, medical device patch management and deployment can be stymied due to operational environments and constraints (e.g., hospital operations, in-clinic management of technology), but patching challenges may be more tractable in AID systems in particular. First, some components of AID systems are disposable, such as Dexcom's 15-day CGM. In such cases, "patches" are not changes to devices "in the field" but rather devices as they are manufactured and provisioned, meaning patch management is fully "owned" by the manufacturer, and occurs within its manufacturing facilities rather than in hospitals. Second, for non-disposable AID components that do require field updates, the personal-device nature of these systems offers deployment advantages over traditional hospital-managed devices. However, recent incidents (e.g., a Google Play store configuration change negatively affecting insulin pump availability and function [39]) highlight that user trust in software updates remains fragile. Increasing the trustworthiness and completeness of the verification and validation pipeline through the moonshot would be particularly valuable for AID systems, where users must rely on frequent updates to personal medical devices without the safety net of clinical oversight.

3.5. Policy Establishing Appropriate Economic Incentives (CHP5)

CHP5 addresses the widespread misalignment of economic incentives in cybersecurity, where suppliers of cyber systems are seldom held liable even for poor-quality products, nor are there sufficient rewards for high quality. The core problem is that cyber systems are opaque and security attributes are hard to measure, so poor security becomes the norm since customers cannot effectively distinguish between good and bad products (a market failure described as the "market for lemons" [40]). Suppliers have avoided effective regulation by arguing that accountability would hinder innovation, while many policy initiatives result in weak process-focused mandates with self-attestation of compliance. This creates a "catch-22" situation where inadequate security practices impair the potential for warranting results, but without this potential, there is less incentive to invest in advancing those practices and tools.

Relevance to AID Tech: Recognizing the direct threat that vulnerabilities pose to patient safety, regulatory bodies are actively working to establish economic incentives that push manufacturers towards a higher standard of cybersecurity. A primary example of this is the FDA's Section 524B mandate [41]. This applies to devices with software that can connect to the internet, and directly impacts manufacturers of AID technology by requiring them to submit a comprehensive plan for postmarket vulnerability management, demonstrate secure design and development processes, and provide an SBOM for premarket approval. The economic incentive here is clear: non-compliance slows or prevents regulatory approval, effectively barring companies from a lucrative segment of the medical device industry.

Legal changes would play a role in incentivizing cybersecurity of AID systems. In January 2025, the U.S. government released proposed changes to HIPAA intended to strengthen cybersecurity, though these changes are still being considered [42]. There are grassroots initiatives to modernize diabetes data rights in the United States [43], including real-time access to CGM data [44].

Beyond direct regulatory and legal hurdles, the market itself is beginning to exert its own economic pressure. Reports indicate a willingness among healthcare buyers to pay a premium for devices with robust security features, creating a competitive advantage for manufacturers who invest proactively in cybersecurity [45]. Conversely, manufacturers of AID devices that experience highly publicized security incidents, such as past vulnerabilities in certain insulin pump systems, face significant reputation damage, potential liability, and a loss of market share [46]. This dynamic incentivizes manufacturers to prioritize cybersecurity as a core component of their product's value proposition and brand integrity.

Opportunity in AID Tech: The AID ecosystem presents unique advantages for demonstrating how aligned economic incentives can enable rapid automated patching in medical devices. Unlike many medical device markets where security failures have delayed consequences, AID device vulnerabilities create immediate, measurable health risks that translate directly into liability exposure and market consequences for manufacturers. The engaged T1D user community amplifies these market forces by quickly identifying and publicizing security issues, while the existing regulatory framework provides a foundation for demonstrating that automated patching systems can satisfy liability and compliance requirements while reducing long-term costs.

3.6. Human–System Interactions (CHP6)

CHP6 addresses the challenge that cyber systems often present confusing and uninformative human interfaces that can encourage unsafe behavior, even for expert users [47]. The core problem is that human interfaces for security-related interactions—ranging from authentication and privacy control to configuration of access policies and responding to security alerts—are poorly suited to their purpose. Inadequate attention to user design for security and privacy, coupled with failed market incentives and regulation, has led to growing problems in establishing and maintaining user security and privacy. This is compounded by the complexity and nuance of many security-related interactions, with errors in human interaction being by far the dominant cause of security breaches. Training and awareness approaches alone are insufficient and the solution requires design adaptations to human interfaces and judicious choices regarding operational workflow structure.

Relevance to AID Tech: Human–system interaction plays a crucial role in the security and adoption of AID technologies, especially as these devices become increasingly connected and reliant on user input. The design of AID software often presumes a baseline of digital literacy that many patients or caregivers may lack. In a study of caregivers of children with diabetes, 50% cited “limited technical skills” as a barrier to adopting open-source AID systems [48]. Similarly, widespread use of closed systems without transparency or rigorous user training leaves patients vulnerable to malicious updates, spoofed apps, or insecure pairing mechanisms. Effective security for AID tech must therefore include human-centric design principles and minimize the burden on users [49].

Insulin management, especially in type 1, is already cognitively burdensome, with patients spending an average of 77 minutes per day actively thinking about their diabetes [50]. This constant engagement creates both opportunities and risks—opportunities for real-time security prompts and layered authentication—but also risks from decision fatigue and interface overload. To achieve this, developers must incorporate insights from social science and behavioral economics, crafting interactions that make secure behaviors the path of least resistance such as simplifying updates, using progressive disclosure for advanced features, and enabling users to verify system authenticity regardless of technical expertise.

Opportunity in AID Tech: While automated software updates should require minimal human–system interaction, AID users interact with their devices multiple times daily and have demonstrated comfort with technology, making them ideal validators for automated update interfaces. The T1D community’s constant engagement provides immediate feedback on whether automated update notifications, progress indicators, and error messages are clear and actionable. This active user base can rapidly identify interface design flaws that might compromise user trust in automated security updates, establishing interface patterns directly applicable to other medical technologies where user confidence in automated security processes is essential.

3.7. Information Provenance, Social Media, and Disinformation (CHP7)

Information provenance refers to maintaining a complete, traceable record of where each piece of information (or software) originated, how it was modified, and through what channels it traveled before being shared. CHP7 addresses both this critical need and the complex challenge that social media platforms can profoundly affect opinions across broad populations while enabling precision targeting at scale to manipulate opinions and perpetrate scams. The fundamental problem is the proliferation of disinformation and the potential degradation of trustworthiness and reliability of online media. Without a focused effort on detecting trustworthy data and information, the integrity of online information ecosystems remains at risk.

Relevance to AID Tech: This problem intersects critically with AID tech through the information ecosystems that patients, caregivers, and clinicians rely on for decision-making and support. T1D management today is highly digital, with individuals turning to social media platforms, YouTube tutorials, patient forums, and Reddit communities for real-time advice on device settings, troubleshooting, and algorithm tuning. While these platforms foster empowerment and shared learning, they also introduce significant risks in the form of misinformation, unverified medical claims, and misleading portrayals of unapproved modifications to insulin delivery systems. For example, recent studies have documented how platforms such as TikTok propagate erroneous diabetes management tips and promote off-label or unsafe practices under the guise of community sharing [51]. Without mechanisms to establish the provenance of health information patients may be left unable to distinguish between trustworthy and dangerous content.

The T1D open-source community offers an instructive case in navigating misinformation risks. Projects like OpenAPS and Loop maintain rigorous documentation, change logs, and peer-reviewed technical write-ups to ensure that shared guidance can be traced to original, identifiable contributors.

Opportunity in AID Tech: For AID technologies, solving CHP7 means not only addressing technical media provenance (e.g., signed updates with verifiable chain-of-custody; see also Section 3.4) but also investing in digital health literacy. Manufacturers and advocacy groups are beginning to engage actively on social platforms to counter disinformation and elevate vetted resources.

3.8. CPS and Operational Technology (CHP8)

CHP8 addresses the unique vulnerabilities of cyber-physical systems (CPS) that include operational technology used in manufacturing, civil infrastructure, and transportation, as well as Internet of Things technology used in telecommunications, homes, and industrial controls. The core problem is that CPS are both vulnerable and consequential. They are vulnerable because they have historically operated offline and inherited engineering traditions with lower security standards than IT systems, often lacking access pathways for updates and offering attack vectors via audio, network, and physical access. They are consequential because they serve as the control fabric for civil infrastructure, industrial controls, manufacturing systems, and embedded national security systems. The challenge is compounded because secure design of CPS requires both hardware and software expertise, meaning fewer trained professionals. Moreover, real-time control requirements create complex interdependencies that complicate re-engineering for security and resilience.

Relevance to AID Tech: AID technology exemplifies the cybersecurity challenges inherent in CPS, where the integration of sensors, actuators, and computing elements creates complex attack surfaces that can directly impact patient safety. Modern AID devices operate as sophisticated CPS networks: continuous glucose monitors use subcutaneous sensors to measure interstitial glucose levels, insulin pumps employ mechanical actuators to deliver precise medication doses, and AID systems combine these components with control algorithms that make real-time therapeutic decisions (see Figure 1). Unlike traditional IT systems where failures typically result in data loss or service disruption, vulnerabilities in AID CPS can lead to immediate, life-threatening consequences such as severe hypoglycemia or diabetic ketoacidosis. The operational technology aspects of these systems—including wireless communication protocols, embedded firmware, and real-time control loops—were often designed with assumptions of software and network isolation that no longer hold true as devices become increasingly connected to cloud services, healthcare networks, and general-purpose smartphones. The challenge is compounded by the real-time operational requirements of diabetes management, where system reboots or temporary service interruptions are not generally acceptable.

Opportunity in AID Tech: The unique integration of real-time sensing, actuation, and control algorithms in AID technology presents a significant opportunity for advancing CPS security. The high-stakes environment where cyber vulnerabilities have direct, life-threatening consequences necessitates robust and innovative security solutions that push boundaries beyond traditional IT security paradigms [52]. The critical requirement for continuous operation compels the creation of novel security measures that ensure both resilience and availability, while regulatory emphasis on secure design principles provides structured impetus for research and development, making AID technology a proving ground for engineering secure and resilient CPS.

3.9. AI and Emerging Capabilities (CHP9)

CHP9 addresses the challenges of contemporary AI systems representing technologies with difficult-to-predict safety and security ramifications. The core problems include dependence on often poorly curated training data, the opacity of neural networks that makes behaviors unpredictable, the stochastic nature of outputs that can be inexact and untrustworthy, and an extraordinarily broad attack surface encompassing training data, network architecture, and operational inputs. The challenge is compounded by automation bias and the violation of core security principles when generative AI systems use the same input pathways for both trusted instructions and untrusted data.

Relevance to AID Tech: People with diabetes must make management decisions in their diet and medication based on the trajectory of their blood sugar levels over time [53]. Modern AID systems rely on complex algorithms, often incorporating machine learning (ML), to predict glucose excursions and adjust insulin delivery. This inherent complexity, coupled with the statistical nature of these models, leads to a degree of opacity, making it difficult to fully audit the decisions made by software. The performance of these systems is heavily dependent on extensive and diverse training data, and any biases or gaps can lead to unpredictable or unsafe patient outcomes (e.g., hypoglycemia or hyperglycemia), and is therefore a critical concern in a high-consequence setting like T1D management.

To address these issues, progress in AID technology necessitates the development of robust safety frameworks specifically tailored to the unique risks of AID systems incorporating machine models, including mechanisms for ensuring data integrity and model robustness against mistakes and attacks. Furthermore, there is a growing imperative for model architectures that integrate structured reasoning alongside purely statistical prediction, allowing for interpretable and explainable AI decisions [54], which can foster greater trust and facilitate clinical oversight. Finally, the implementation of external guardrails and independent oversight mechanisms is crucial to constrain the behavior of ML-powered AID systems, ensuring they operate within safe physiological parameters and adhere to established clinical guidelines, especially as software updates introduce new algorithmic iterations that could potentially alter system performance in unforeseen ways. The FDA is actively exploring methods to identify and tag medical devices incorporating advanced ML/AI, including large language models (LLMs), to promote transparency and facilitate appropriate regulatory review [55].

Opportunity in AID Tech: The challenges of ML-enabled AID systems create dual opportunities. First, the continuous operational data from AID devices enables development of robust validation frameworks for ML algorithm updates, allowing real-time assessment of whether algorithmic changes maintain safety bounds and clinical efficacy. Second, this same high-frequency data stream presents unique opportunities for developing AI/ML-driven automated patch validation systems. The well-defined therapeutic targets of maintaining glucose in range provide clear, quantifiable metrics that machine models can monitor to validate that patches, whether addressing ML components or other system elements, have not introduced functional regressions or safety issues. Real-time validation capability is essential for automated patch deployment, as it enables systems to automatically roll back updates if performance degrades, establishing a framework applicable to other medical technologies where automated verification of update safety is critical for enabling rapid response to cybersecurity threats while managing the complexities of ML-enabled systems.

3.10. Operational Security (CHP10)

CHP10 addresses the challenge of maintaining security in increasingly pervasive large-scale systems where users have little insight beyond vendor self-attestation into the state of security configuration. The problem encompasses prevention, detection, response, and recovery, all of which require significant planning, cost, constant practice, and comprehensive knowledge of networks and applications. As systems grow in scale and interconnection, many organizations struggle with expensive and error-prone detection and response, particularly with recovery from attacks.

Relevance to AID Tech: The operational security challenges in AID technology mirror those of large-scale cyber systems, where the complexity and interconnectedness of devices, cloud services, and data analytics platforms create significant challenges for maintaining security visibility and incident response capabilities. Modern AID ecosystems operate across multiple operational domains: glucose data flows from sensors through mobile applications to cloud-based analytics platforms, insulin delivery commands traverse from algorithmic controllers to pump hardware, and software updates propagate from manufacturer servers to patient devices—often with limited visibility for patients, caregivers, or even healthcare providers into the security posture of these distributed systems [56]. The operational security burden is further complicated by the diverse stakeholder ecosystem, where device manufacturers, cloud service providers, mobile app developers, and healthcare systems each maintain separate security operations with varying levels of maturity and coordination. Unlike enterprise IT environments where security teams have centralized monitoring and control, AID operational security relies heavily on vendor self-attestation and fragmented incident response capabilities that may not adequately address the speed and sophistication of modern cyber threats.

Cybersecurity detection, response, and recovery capabilities in AID operational environments face unique constraints that distinguish them from traditional cybersecurity operations. The life-critical nature of diabetes management means that defensive actions such as device isolation, service interruption, or system rollbacks—common responses in enterprise security—may pose greater risks to patient safety than the original security incident [57]. The 2019 Medtronic pump vulnerability exemplified these operational security challenges: while the manufacturer could detect potentially malicious wireless communications, the inability to remotely patch affected devices left patients and providers with limited options beyond manual device replacement—a recovery process that took months to complete and required extensive coordination between manufacturers, healthcare providers, and patients. Recent initiatives such as the Medical Device and Health IT Joint Security Plan (JSP2) represent efforts to establish industry-wide operational security standards, but implementation remains uneven across the AID technology landscape, particularly for smaller manufacturers and legacy device ecosystems [58].

Opportunity in AID Tech: The intricate, distributed nature of AID technology ecosystems presents an opportunity to pioneer advanced operational security paradigms. The fragmented visibility and reliance on vendor self-attestation within these multi-stakeholder environments underscore the critical need for novel, transparent security solutions that extend beyond traditional enterprise models. The life-critical aspect of T1D management compels the development of innovative detection, response, and recovery strategies that prioritize patient safety and device availability. This pressing constraint fosters research into "warm" recovery mechanisms, in-situ patching, and anomaly detection systems that can provide real-time threat mitigation without compromising therapeutic continuity. Furthermore, the inherent challenges associated with legacy device support and the diverse maturity levels of security operations across the AID technology landscape create a fertile ground for establishing and validating industry-wide operational security standards. Consequently, AID systems for insulin-requiring people with diabetes serves as a *crucial testbed for developing resilient, continuously monitored, and highly coordinated operational security frameworks* that can be adapted to other cyber-physical systems (including other medical devices) where uninterrupted functionality is paramount.

4. A Moonshot for Automated, Secure AID Software Update Pipelines

We propose a moonshot to demonstrate a progressively automated, inexpensive, safe, secure, and FDA-qualified testing pipeline for AID software updates A moonshot aims to achieve an ambitious, concrete goal to dramatically reduce the need for human decision-making in software update pipelines for AID systems on the market.

As a general goal, the testing pipeline will be increasingly automated within widening boundaries, starting with tightly-scoped cases and generalizing as more data is collected regarding the resulting level of assurance. In other words, **automation becomes the focus of developers' work, and their role shifts from *performing* the assurance to *assessing the safety and performance of automated assurance*** (e.g., via stratified sampling and manually checking edge cases).

Based on our analysis in Section 3, we conclude that this moonshot represents a radical yet feasible solution to an important problem [59]. In this section, we characterize the performance criteria for the moonshot, the current bottlenecks, and the technical and regulatory innovations necessary to meet the performance criteria.

4.1. Moonshot Success Criteria

To define tangible success metrics, we draw from Heilmeier Catechism (originally used at the Defense Advanced Research Projects Agency [60] and recently adopted with elaboration by the Advanced Research Projects Agency for Health (ARPA-H) [61]), which suggests outlining midterm and final "exams" to check for success in addition to estimates of cost, duration, and risks. The following subsections address each requirement.

Validation Approach and Success Metrics. Validation of the moonshot approach must occur through a collaborative demonstration of the pipeline's performance, together with willing incumbents, their regulator (i.e., FDA), other stakeholders (e.g., clinicians, patients), third-party technical experts, and moonshot organizers.

In practical terms, the automated testing pipeline must be adoptable by current AID tech incumbents, tolerable to decision-makers, and more secure than today. We suggest three categories of performance: business, regulatory, and cybersecurity. Business success includes adoption by tech manufacturers (i.e., the pipeline actually becomes usable and in use), business benefits from adoption (e.g., return on investment). Regulatory success implies that the pipeline is acceptable by FDA as a replacement for current pipelines, and that regulators would recommend the approach for other medical devices. Cybersecurity success would include increasing the security of AID systems. In Table 1 we offer specific, measurable metrics for success. Below the table, we offer more descriptions of the categories and measurement approaches for each. Following this section, we discuss control criteria that must be controlled for while measuring the success of the moonshot. Precise metrics reflect pragmatic targets that balance feasibility and ambition.

Business Success Criteria ("No-Brainer" Adoptable by Incumbent AID Device Manufacturers):

- Adoptable by at least three of the current on-market AID tech manufacturers and devices (i.e., pipeline applies to Tandem, Abbott, Insulet, Dexcom, Medtronic's tech stacks), as measured by business transactions and use of the pipeline.
- Significant value proposition. Project at least 10x return on investment within three years of adoption, including at least a 90% reduction in postmarket software testing costs and time, while simultaneously improving device security posture. As measured by an ROI model, ROI estimates and projected deltas in benchmarks of software validation and verification speed, quality, and business costs.
- Reduce incumbents' regulatory uncertainty and increase the predictability of FDA's acceptance of software updates, as measured by FDA's qualification of the pipeline.
- Promise of continued innovation. Viable business model supports sustaining the pipeline over lifetime of medical devices. Is not proprietary to, or customized for, specific medical device architectures, as measured by adoption by multiple incumbents with different architectures. Vendors (including open-source tool integration vendors) exist to sell, license, and maintain the pipeline over time. There are viable on-market products with revenue-generating, profitable businesses.

Regulatory Success Criteria (Qualified Medical Device Development Tool):

- FDA acceptance of Medical Device Development Tools (MDDT) qualification [62] for core pipeline components; demonstration of regulatory pathway for day-scale security updates; establishment of precedent for continuous validation approaches in life-critical medical devices; no foreseeable increase in unacceptable failure (i.e., silent irreversible, or severe harm).

Cybersecurity Success Criteria:

- Solves Cyber Hard Problems, as measured by external expert observers' feedback (e.g., through interviews or focus groups with representatives from the original NASEM panel).
- AID tech improves quickly, as measured by the increased speed of software updates, from conception to deployment.
- New entrants (e.g., startups building a new CGM) can rely on the pipeline to reduce cybersecurity uncertainty, as measured by interviews with CEOs and CTOs of devices that are in premarket development.
- The pipeline should be publishable and publicly auditable while maintaining its security [63].
- Automated test generation covering >95% of device functionality; validation time reduction of >90% compared to current manual processes; minimal false negatives in security vulnerability detection during demonstration scenarios.
- The pipeline results in a demonstrated improvement in time-to-deployment for critical security patches; maintained or improved user trust in device security and reliability.

4.1.1. Control Criteria and Risks.

The Moonshot pipeline will do no harm, adhering to the principles of the Hippocratic Oath for Connected Medical Devices [64]. The success of the moonshot will be judged based on a measurable positive effect on control criteria and reduction in unintended consequences.

System-level failure detection. The moonshot pipeline should not decrease the overall system's ability to detect and resolve failure. By overall system, we mean MDMs, regulators, adverse event surveillance, and users. The moonshot will not introduce the possibility of "silent" failures.

Business Success Metrics	Regulatory Metrics	Cybersecurity Metrics
Adopted by >3 incumbents	Qualified as MDDT by FDA	Increase security of AID systems
High expected ROI	Regulatory NPS >90%	Solve NASEM Cyber Hard Problems
Increased speed of testing patches	Fewer recalls	Lower time-to-deployment for critical patches
Lower cost of market entry	Simplified software review	Increase quality of new software
Reduced regulatory uncertainty		Automated testing >95% device functionality Min. false negatives in vulnerability detection

TABLE 1. MOONSHOT SUCCESS PERFORMANCE CRITERIA AND METRICS.

People living with diabetes. Reduction in features or freedom to tailor AID systems to user needs represents a risk. The moonshot pipeline should not reduce the freedom and safety of people with AID, in any way, as measured by interviews with key representatives in T1D communities, leaders calling for patient rights. The moonshot should not lead to degradation in glycemic outcomes or user experience during accelerated update cycles.

Regulatory burden. The Moonshot pipeline should not increase ongoing regulator burden, as measured by regulator feedback, interviews, and feedback during the FDA Medical Device Development Tools (MDDT) review process [62]. The Moonshot should not increase other legal, regulatory, compliance, or brand risk to incumbents, as measured by interviews with incumbent executive leaders, adoption decisions, and other regulator-stakeholder feedback (e.g., leveraging the International Medical Device Regulators Forum (IMDRF)).

New AID technologies. The Moonshot pipeline should not increase barriers to entry for new device manufacturers, as measured by the voluntary use of the pipeline, reasonable cost, its availability, and interviews with startup or new-entry leaders. The moonshot should not reduce the technical freedom of future, potentially better, and disruptive technologies. The moonshot is not intended to prescribe a certain technical stack for future use in AID and/or aim to become a required "standard."

Incumbent business. The Moonshot should not result in unplanned costs for participating medical device manufacturers. Voluntary participation should not result in long-term cost increases or short-term business disruption. A significant risk to undertaking and participating in this moonshot is the risk of being penalized for discovering cybersecurity vulnerabilities in on-market products that result in unplanned action that otherwise would not have been discovered. This risk bears further discussion: see note below.

Note that risk assessment of a new vulnerability may be irrational. A newly discovered vulnerability does not necessarily increase the risk in the system. Rather, it *reduces the uncertainty* about risk. Reducing uncertainty about risk in a latent system should be considered a positive action. Any improvement in the detection of vulnerabilities in a system may cause a *perceived* increase in risk. While the vulnerabilities should be addressed in the normal course of producing software updates, the discovery of the new vulnerability should not be considered a "new" risk. It should not automatically trigger action on the part of the medical device manufacturers. We therefore we believe that *replacing the vulnerability assessment and risk assessment process with more rational, objective, and logically honest approach must be the first step*.

Milestones. While R&D is inherently unpredictable, we suggest a three-year effort with milestones. Inherent in this short timeline is an assumption that all technical solutions exist and are (or will soon be) available as products. **Year 0 Milestones:** Stakeholder feedback report regarding moonshot design, success criteria, and incentive-aligned signals of support. Research and development program designed and funded at a level likely to achieve success. **Year 1 Milestones:** Postmarket cybersecurity decision model published, MDDT proposal submitted. Business ROI model published, with incumbent MDM feedback and current state benchmarks (e.g., costs, timelines, and cybersecurity posture of current testing pipelines). Moonshot success and measurement model submitted for public stakeholder review (e.g., American Diabetes Conference Scientific Sessions). Technical demonstrations, pilots, and convening sessions (e.g., hackathon-style conferences) demonstrate positive initial results; case studies submitted to academic conferences. **Year 2 Milestones:** Prospective side-by-side trials of moonshot versus incumbent MDM current state pipelines. Trial results submitted for publication. **Year 3 Milestones:** Integration and adoption of moonshot pipelines. Integration reports submitted for publication. **Post-moonshot:** Continued integration, measurement of outcomes, communication of successes and failures.

4.1.2. Resources. Our order-of-magnitude assumption is that this effort will primarily be performed by existing vendors, and that the costs will be shared by the beneficiaries (e.g., incumbent MDMs), government (e.g., DARPA, ARPA-H), and academic and scientific participation (e.g., NASEM). Many of the efforts required for the moonshot are already underway. The DIGIHEALS, AIxCC [65], and UPGRADE programs at DARPA and ARPA-H are all related to and intend to product outcomes that achieve some of the moonshot goals. Nudge BG is developing an Automated Insulin Delivery Interoperability Framework (AIDIF) to be qualified as an MDDT [66]. The Nudge BG AIDIF is a simulator of clinical outcomes for type 1 diabetes, and it may be a critical element of connecting safety risk assessment to the moonshot pipeline.

The novelty—and resources needed—in the moonshot is the composition of the solutions into a mission-focused solution that manufacturers of cyber-physical systems could adopt. We expect that the orchestration of the moonshot will cost more than one million and less than ten million dollars, with the assumption that directly related and ongoing existing efforts at ARPA-H, DARPA, FDA, and other related programs are sustained. Post-moonshot sustainability would be achieved through an industry consortium model where manufacturers fund the shared infrastructure.

4.2. Current State and Systemic Bottlenecks

The primary bottlenecks include: (1) Manual verification and validation processes that require human experts to design test cases for each update; (2) Sequential regulatory review where documentation preparation and FDA interaction occur after technical validation is complete; (3) Fragmented testing environments where cybersecurity, clinical safety, and interoperability verification and validation occur in isolation; and (4) Conservative risk assessment frameworks that treat any software change as potentially introducing novel failure modes requiring extensive re-validation.

These bottlenecks compound when addressing supply chain vulnerabilities or newly discovered attack vectors. The 2019 Medtronic MiniMed pump recall [29] exemplified this challenge: despite the manufacturer's ability to detect potentially malicious wireless communications, the inability to rapidly deploy patches left patients and providers with no option beyond manual device replacement—a process requiring months to complete.

4.3. Regulatory Innovation

The moonshot leverages existing FDA frameworks while proposing novel approaches to medical device software regulation. The MDDT pathway [62] provides a mechanism for pre-approving testing protocols and validation tools, enabling manufacturers to achieve regulatory certainty while dramatically reducing time-to-market for critical updates.

We believe that manufacturers adopting this approach could achieve a greater return on investment within five years through reduced regulatory compliance costs, faster time-to-market for new features, and enhanced competitive positioning. More critically, the ability to rapidly respond to cybersecurity threats protects market position and patient trust in an increasingly connected healthcare ecosystem.

The proposed framework also addresses perverse economic incentives identified in Section 3.5 by creating clear competitive advantages for manufacturers who invest in rapid, secure update capabilities. Healthcare providers increasingly prioritize cybersecurity in procurement decisions, creating market demand for demonstrably secure and rapidly updatable devices.

4.4. Technical Implementation and Innovation

The moonshot implementation centers on three technical innovations that address fundamental Cyber Hard Problems while advancing the state of medical device cybersecurity.

Digital Twin Validation Networks: High-fidelity digital models of AID devices and physiological systems will enable exhaustive testing of software updates without requiring physical hardware for every test scenario [67]. Digital twin validation networks employ diverse modeling approaches and independent implementations to minimize common-mode failures, with disagreements triggering enhanced human review rather than automatic approval. When twins disagree or both indicate potential issues, the system defaults to existing manual validation processes, ensuring that automation accelerates only high-confidence scenarios. These digital twins incorporate both device behavior models and patient glucose dynamics, allowing validation of algorithmic changes under thousands of simulated clinical conditions within hours rather than weeks. This aligns with CHP3 (secure composition; Section 3.3) and CHP8 (cyber-physical systems; Section 3.8).

Cryptographic Update Provenance: Every software update is cryptographically signed and timestamped, with metadata including chain-of-custody information, creating an immutable audit trail from code commit to patient device deployment. This addresses CHP4 (supply chain; Section 3.4) and CHP7 (information provenance; Section 3.7) concerns by increasing assurance and transparency together, and later enabling rapid rollback capabilities if issues are discovered post-deployment.

Adaptive Risk Assessment Algorithms: Machine learning (ML) systems continuously analyze update patterns, deployment outcomes, and threat intelligence to refine risk classification models. Updates determined to be low-risk based on historical data and automated testing can proceed through accelerated pathways, while high-risk changes trigger enhanced validation protocols. This aligns with CHP1 (risk assessment; Section 3.1).

4.5. Proposed Architecture: Continuous Verification and Validation

Our proposed pipeline architecture transforms AID updates from an episodic, manual process to a continuous, automated capability. The core innovation lies in pre-validated testing frameworks that receive regulatory approval as Medical Device Development Tools (MDDTs), enabling manufacturers to leverage pre-approved validation protocols for rapid deployment of security and safety updates.

Automated Verification and Validation Engine: The system employs ML/AI-assisted test generation and execution, automatically creating comprehensive test suites for each software modification [65], [68]. Machine models trained on historical AID device behavior patterns can predict potential failure modes and generate targeted test scenarios, reducing manual test design from weeks to hours. This aligns with CHP9 (AI; Section 3.9).

Continuous Regulatory Alignment: Rather than sequential regulatory review, the pipeline maintains continuous FDA alignment through pre-approved change control protocols. Manufacturers submit software update categories and associated testing protocols for upfront approval, enabling automatic clearance for updates that fall within pre-validated parameters. This aligns with CHP5 (economic incentives; Section 3.5).

Integrated Security and Safety Validation: The pipeline incorporates real-time threat intelligence feeds and automated penetration testing specifically designed for AID attack vectors. This enables simultaneous validation of cybersecurity measures and clinical safety requirements, eliminating the traditional separation between security and safety testing. Further developed, this capability can be expanded to auto-discovery of software vulnerabilities and auto-generation of patches. This aligns with CHP9 (AI; Section 3.9).

Interoperability Assurance Framework: Automated testing of device combinations and communication protocols ensures that security updates do not introduce unexpected interaction effects between CGMs, pumps, and mobile applications from different manufacturers. This aligns with CHP3 (secure composition; Section 3.3) and CHP10 (operational security; Section 3.10).

5. Discussion

The temporal mismatch today between cyber threat emergence (hours) and medical device software update cycles (months) represents a fundamental weakness in our healthcare cybersecurity infrastructure. Through systematic analysis of automated insulin delivery (AID) systems against the National Academies' Cyber Hard Problems [1], we have demonstrated that this medical technology domain offers a uniquely advantageous proving ground for solving cybersecurity challenges that extend far beyond diabetes care. AID offers life-critical stakes, engaged user communities, established regulatory frameworks, and continuous device interaction, creating an environment where cybersecurity innovations can be rapidly developed, validated, and refined.

AID systems would demonstrate that life-critical medical devices can achieve consumer-technology-like update velocities without sacrificing safety or regulatory compliance. This capability becomes increasingly essential as medical devices become more and more connected and cyber threats increasingly sophisticated.

While focused on AID technologies for people with insulin-requiring diabetes, this moonshot establishes principles and infrastructure applicable across the medical device ecosystem. Success in this domain would establish replicable principles for rapid, secure software updates across the broader medical device ecosystem and other safety-critical cyber-physical systems, fundamentally reimagining how medical device manufacturers can respond to cyber threats while maintaining safety assurances.

The regulatory innovation pathway through Medical Device Development Tools provides a concrete mechanism for achieving this transformation without compromising patient safety or existing compliance frameworks. Rather than requiring radical shifts in organizational culture or massive new investments, this approach leverages existing FDA processes while introducing automation and intelligence that can dramatically accelerate security response times. The economic incentives align favorably: manufacturers gain competitive advantages through faster feature deployment and enhanced security postures, while healthcare systems benefit from reduced cyber risk exposure and improved device reliability.

5.1. Limitations

Our approach is tailored specifically to medical technology, and does not consider the Cyber Hard Problems for other scenarios. Instead, we use AID technology to offer a concrete and illustrative mission problem that may aid the broad issue of software understanding [8], [9]. We focus on the process of patch/update development and testing but provide only limited discussion of the challenges of the installation; this latter step will also require innovation and research [69]. Finally, our proposals for regulatory alignment focus on the United States; we leave extensions of the argument to other parts of the world for future work.

This paper offers a framework and a starting position meant to spark discussion. In one sense, it is an oversimplification of the totality of challenges in the real world. We aim to offer sufficient framing to elicit active participation from experts with a variety of backgrounds (e.g., total system reliability, formal methods, past cases of successful and failed technology).

5.2. Conclusion

The cybersecurity challenges facing our healthcare infrastructure demand bold, coordinated action that transcends traditional boundaries between industry, academia, and government. The automated insulin delivery (AID) ecosystem offers an immediate opportunity to demonstrate that life-critical medical devices can achieve consumer-technology-like update velocities without sacrificing safety or regulatory compliance. We call upon cybersecurity researchers, medical device manufacturers, regulatory bodies, and the T1D community to embrace this moonshot as a path toward a more resilient and responsive healthcare cybersecurity posture. The time for incremental improvements has passed: our digital health infrastructure requires the transformational change that only ambitious, focused collaboration can deliver.

Acknowledgments

We thank Matt McMahon and Liam Pender for sharing their experiences and insights. We also thank those who provided input but wish to remain anonymous. Finally, we are grateful to the anonymous reviewers for their helpful suggestions which significantly improved the paper.

References

[1] National Academies of Sciences, Engineering, and Medicine, *Cyber Hard Problems: Focused Steps Toward a Resilient Digital Future*. Washington, DC: The National Academies Press, 2025. [Online]. Available: https://nap.nationalacademies.org/catalog/29056/cyber-hard-problems-focused-steps-toward-a-resilient-digital-future

[2] M. Fang, D. Wang, and E. Selvin, "Prevalence of type 1 diabetes among US children and adults by age, sex, race, and ethnicity," *JAMA*, vol. 331, no. 16, pp. 1411–1413, 2024.

[3] I. B. Hirsch, T. Battelino, A. L. Peters, J. J. Chamberlain, G. Aleppo, and R. M. Bergenstal, *Role of continuous glucose monitoring in diabetes treatment*. American Diabetes Association, 2018, vol. 6.

[4] I. Frankovic, M. Kuhn, and S. Wrzaczek, "Medical innovation and its diffusion: Implications for economic performance and welfare," *Journal of Macroeconomics*, vol. 66, 2020.

[5] S. W. Ruggles, J. R. Perl, Z. Sexton, K. Schulman, and J. Makower, "The need for accelerated Medicare coverage of innovative technologies: Impact on patient access and the innovation ecosystem," *Health Management, Policy and Innovation (HMPI)*, vol. 7, no. 1, January 2022.

[6] SIS International Research, "Digital health technology trends: A survey of digital healthcare decision-makers," https://jabil.com/dam/jcr:2e0e9986-d8d8-4c37-af81-1dbe540d306d/2021-Digital-Healthcare-Survey-Report_Jabil.pdf, 2021.

[7] Taking Control Of Your Diabetes, "Automated insulin delivery systems: Which one is right for you?" https://tcoyd.org/2025/04/automated-insulin-delivery-systems-comparison/, April 2025.

[8] D. Ghormley, T. Amon, C. Harrison, and T. Loffredo, "The national need for software understanding: The present crisis, technical capability gaps, and path forward," Sandia National Laboratories, Tech. Rep. SAND2025-03576R, March 2025. [Online]. Available: https://www.sandia.gov/app/uploads/sites/224/2025/03/The-National-Need-for-Software-Understanding-v1.0.pdf

[9] CISA, DARPA, OUSD(R&E), and NSA, "Closing the software understanding gap," https://www.cisa.gov/resources-tools/resources/closing-software-understanding-gap, Cybersecurity and Infrastructure Security Agency, Defense Advanced Research Projects Agency, Office of the Under Secretary of Defense for Research and Engineering, National Security Agency, Tech. Rep., January 2025. [Online]. Available: https://www.cisa.gov/sites/default/files/2025-01/joint-guidance-closing-the-software-understanding-gap-508c.pdf

[10] U.S. Food and Drug Administration, "Cybersecurity in medical devices: Quality system considerations and content of premarket submissions," https://www.fda.gov/media/119933/download, 2025.

[11] S. Lantzy, Personal communication, June 2025.

[12] U.S. Food and Drug Administration, "Postmarket management of cybersecurity in medical devices," https://www.fda.gov/media/95862/download, 2016.

[13] L. Reverberi and D. Oswald, "Breaking (and fixing) a widely used continuous glucose monitoring system," in *11th USENIX Workshop on Offensive Technologies (WOOT '17)*, 2017.

[14] J. Radcliffe, "Hacking medical devices for fun and insulin: Breaking the human SCADA system," 2011. [Online]. Available: https://www.youtube.com/watch?v=avf5XF8yS60

[15] D. C. Klonoff, "Cybersecurity for connected diabetes devices," *Journal of Diabetes Science and Technology*, vol. 9, no. 5, pp. 1143–1147, 2015.

[16] B. Nazzal, M. H. Alalfi, and J. R. Cordy, "A survey on the verification and validation of artificial pancreas software systems," in *2024 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*. IEEE, 2024, pp. 1–9.

[17] Y. Niu and S.-K. Lam, "Securing automated insulin delivery systems: A review of security threats and protectives strategies," *arXiv preprint arXiv:2503.14006*, 2025.

[18] P. Chang, V. Krish, and A. Rahmati, "Security analysis of RL-based artificial pancreas systems," in *Workshop on Cybersecurity in Healthcare*, 2023, pp. 69–76.

[19] ——, "Security analysis of RL-based artificial pancreas systems," in *Workshop on Cybersecurity in Healthcare*, 2024, pp. 69–76.

[20] D. Lewis, "History and perspective on DIY closed looping," *Journal of Diabetes Science and Technology*, vol. 13, no. 4, pp. 790–793, 2019.

[21] Z. Abou Khalil, E. Constantinou, T. Mens, and L. Duchien, "On the impact of release policies on bug handling activity: A case study of Eclipse," *Journal of Systems and Software*, vol. 173, 2021.

[22] F. Khomh, B. Adams, T. Dhaliwal, and Y. Zou, "Understanding the impact of rapid releases on software quality: The case of Firefox," *Empirical Software Engineering*, vol. 20, no. 2, pp. 336–373, 2015.

[23] R. Bielawski, R. Gaynier, D. Ma, S. Lauzon, and A. Weimerskirch, "Cybersecurity of firmware updates," National Highway Traffic Safety Association, Tech. Rep. DOT HS 812 807, October 2020. [Online]. Available: https://www.nhtsa.gov/sites/nhtsa.gov/files/documents/cybersecurity_of_firmware_updates_oct2020.pdf

[24] M. L. Tanenbaum, E. Iturralde, S. J. Hanes, S. C. Suttiratana, J. M. Ambrosino, T. T. Ly, D. M. Maahs, D. Naranjo, N. Walders-Abramson, S. A. Weinzimer, B. A. Buckingham, and K. K. Hood, "Trust in hybrid closed loop among people with diabetes: Perspectives of experienced system users," *Journal of Health Psychology*, vol. 25, no. 4, pp. 429–438, 2020.

[25] H. Venugopalan, S. M. A. Vijayanand, C. Stanford, S. Crossen, and S. T. King, "GlucOS: Security, correctness, and simplicity for automated insulin delivery," *arXiv preprint arXiv:2406.18262*, 2024.

[26] J. Carlson, "St. Jude stock tumbles as report questions company's cybersecurity," *The Minnesota Star Tribune*, August 2016. [Online]. Available: https://www.startribune.com/st-jude-stock-tumbles-as-report-questions-company-s-cybersecurity/391310831

[27] R. E. Thompson, M. McLaughlin, C. Powers, and D. Votipka, ""There are rabbit holes I want to go down that I'm not allowed to go down": An investigation of security expert threat modeling practices for medical devices," in *USENIX Security Symposium*, 2024, pp. 4909–4926.

[28] U.S. Food and Drug Administration, "Evaluation of automatic class III designation for Control-IQ technology: Decision memorandum," https://www.accessdata.fda.gov/cdrh_docs/reviews/DEN190034.pdf, 2019.

[29] ——, "FDA warns patients and health care providers about potential cybersecurity concerns with certain Medtronic insulin pumps," https://www.prnewswire.com/news-releases/fda-warns-patients-and-health-care-providers-about-potential-cybersecurity-concerns-with-certain-medtronic-insulin-pumps-300876422.html, June 2019.

[30] Tandem Diabetes Care, Inc., "Tandem Diabetes Care announces FDA approval and launch of t:slim X2 insulin pump with Dexcom G5 mobile CGM integration," https://investor.tandemdiabetes.com/news-releases/news-release-details/tandem-diabetes-care-announces-fda-approval-and-launch-tslim-x2, August 2017.

[31] Association for the Advancement of Medical Instrumentation (AAMI), "AAMI TIR57:2016/(R)2023; Principles for medical device security—Risk management," Association for the Advancement of Medical Instrumentation (AAMI), Tech. Rep., 2023.

[32] ——, "AAMI TIR97:2019/(R)2023; Principles for medical device security—Postmarket risk management for device manufacturers," Association for the Advancement of Medical Instrumentation (AAMI), Tech. Rep., 2023.

[33] Tidepool Project, "Tidepool Loop," https://www.tidepool.org/tidepool-loop, 2025, accessed 7/30/2025.

[34] C. T. Hager and S. F. Midkiff, "An analysis of Bluetooth security vulnerabilities," in *IEEE Wireless Communications and Networking (WCNC)*, vol. 3, 2003, pp. 1825–1831.

[35] M. E. Garbelini, C. Wang, S. Chattopadhyay, S. Sumei, and E. Kurniawan, "SweynTooth: Unleashing mayhem over Bluetooth Low Energy," in *USENIX Annual Technical Conference (USENIX ATC)*, July 2020, pp. 911–925.

[36] D. Lewis, "OpenAPS," https://openaps.org/, 2025, accessed 7/30/2025.

[37] J. W. Lum, R. J. Bailey, V. Barnes-Lomen, D. Naranjo, K. K. Hood, R. A. Lal, B. Arbiter, A. S. Brown, D. J. DeSalvo, J. Pettus, P. Calhoun, and R. R. Beck, "A real-world prospective study of the safety and effectiveness of the loop open source automated insulin delivery system," *Diabetes technology & therapeutics*, vol. 23, no. 5, pp. 367–375, 2021.

[38] C. M. Mejía-Granda, J. L. Fernández-Alemán, J. M. Carrillo-de Gea, and J. A. García-Berná, "Security vulnerabilities in healthcare: An analysis of medical devices and software," *Medical & Biological Engineering & Computing*, vol. 62, no. 1, pp. 257–273, 2024.

[39] Multiple contributors, "Omnipod app on Android failure today," https://forum.breakthrought1d.org/t/omnipod-app-on-android-failure-today/74786, May 2025, accessed 10/30/2025.

[40] G. A. Akerlof, "The market for "lemons": Quality uncertainty and the market mechanism," *The Quarterly Journal of Economics*, vol. 84, no. 3, pp. 488–500, August 1970.

[41] 21 U.S. Code § 360n-2, "Ensuring cybersecurity of devices," https://www.law.cornell.edu/uscode/text/21/360n-2, 2022.

[42] U.S. Department of Health and Human Services, Office for Civil Rights, "HIPAA Security Rule Notice of Proposed Rulemaking to Strengthen Cybersecurity for Electronic Protected Health Information: Fact Sheet," December 2024, notice of Proposed Rulemaking issued December 27, 2024. [Online]. Available: https://www.hhs.gov/hipaa/for-professionals/security/hipaa-security-rule-nprm/factsheet/

[43] S. O'Donnell, "My diabetes data is a matter of life and death," *Zócalo Public Square*, 2025. [Online]. Available: https://www.zocalopublicsquare.org/diabetes-data-matter-of-life-death/

[44] P. Edmiston, R. A. Lal, and J. K. Wagner, "HIPAA is obsolete," 2024, submitted to the New England Journal of Medicine.

[45] RunSafe, "2025 medical device cybersecurity index," https://runsafesecurity.com/whitepaper/medical-device-index-2025/, 2025.

[46] Reuters, "FDA warns of cybersecurity risk with certain Medtronic insulin pumps," *Reuters*, September 2022. [Online]. Available: https://www.reuters.com/business/healthcare-pharmaceuticals/fda-warns-cybersecurity-risk-with-certain-medtronic-insulin-pumps-2022-09-20/

[47] A. Whitten and J. D. Tygar, "Why Johnny can't encrypt: A usability evaluation of PGP 5.0," in *USENIX Security Symposium*, vol. 348, 1999, pp. 169–184.

[48] A. Huhndt, Y. Chen, S. O'Donnell, D. Cooper, H. Ballhausen, K. A. Gajewska, T. Froment, M. Wäldchen, D. M. Lewis, K. Raile, T. C. Skinner, and K. Braune, "Barriers to uptake of open-source automated insulin delivery systems: Analysis of socioeconomic factors and perceived challenges of caregivers of children and adolescents with type 1 diabetes from the OPEN survey," *Frontiers in Clinical Diabetes and Healthcare*, vol. 3, 2022.

[49] J. Dykstra and E. H. Spafford, "The case for disappearing cyber security," *Communications of the ACM*, vol. 61, no. 7, pp. 40–42, 2018.

[50] L.-S. Priesterroth, N. Hermanns, B. Kulzer, T. Haak, and D. Ehrmann, "Counting the minutes: Perceived diabetes mental load and its associations with technology use and mental disorders," *Journal of Diabetes Science and Technology*, vol. 19, no. 3, pp. 830–835, 2025.

[51] S. Vasistha, A. Kanchibhatla, J. E. Blanchette, J. Rieke, and A. S. Hughes, "The sugar-coated truth: The quality of diabetes health information on TikTok," *Clinical Diabetes*, vol. 43, no. 1, pp. 53–58, 2025.

[52] T. G. Kirner and A. M. Davis, "Nonfunctional requirements of real-time systems," in *Advances in Computers*, M. V. Zelkowitz, Ed. Elsevier, 1996, vol. 42, pp. 1–37.

[53] P. G. Jacobs, P. Herrero, A. Facchinetti, J. Vehi, B. Kovatchev, M. D. Breton, A. Cinar, K. S. Nikita, F. J. Doyle, J. Bondia, T. Battelino, J. R. Castle, K. Zarkogianni, R. Narayan, and C. Mosquera-Lopez, "Artificial intelligence and machine learning for improving glycemic control in diabetes: Best practices, pitfalls, and opportunities," *IEEE Reviews in Biomedical Engineering*, vol. 17, pp. 19–41, 2023.

[54] A. Das and P. Rad, "Opportunities and challenges in explainable artificial intelligence (XAI): A survey," *arXiv preprint arXiv:2006.11371*, 2020.

[55] U.S. Food and Drug Administration, "Artificial intelligence-enabled medical devices," https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-enabled-medical-devices/, 2025, accessed 7/30/2025.

[56] P. Randine, M. Kopperstad Wolff, M. Pocs, I. R. Connell, J. A. Cafazzo, and E. Årsand, "Unlocking real-time data access in diabetes management: Toward an interoperability model," *Journal of Diabetes Science and Technology*, 2025.

[57] J. Soikkeli, L. Muñoz González, and E. Lupu, "Efficient attack countermeasure selection accounting for recovery and action costs," in *International Conference on Availability, Reliability and Security (ARES)*, 2019.

[58] Health Sector Coordinating Council Cybersecurity Working Group, "Medical device and health IT joint security plan version 2," Health Sector Coordinating Council, Tech. Rep., March 2024. [Online]. Available: https://healthsectorcouncil.org/wp-content/uploads/2024/03/Medical-Technology-and-Health-IT-Joint-Security-Plan-v2.pdf

[59] A. Thierer, "Deep technologies & moonshots: Should we dare to dream?" https://www.mercatus.org/economic-insights/expert-commentary/deep-technologies-moonshots-should-we-dare-dream, September 2018.

[60] Defense Advanced Research Projects Agency (DARPA), "The Heilmeier Catechism," accessed 7/30/2025. [Online]. Available: https://www.darpa.mil/about/heilmeier-catechism

[61] Advanced Research Projects Agency for Health (ARPA-H), "The hidden questions behind the Heilmeier questions," https://arpa-h.gov/sites/default/files/2023-10/Qs_behind_the_HQs.pdf, 2023, accessed 7/30/2025.

[62] U.S. Food and Drug Administration, "Medical Device Development Tools (MDDT)," https://www.fda.gov/medical-devices/medical-device-development-tools-mddt, November 2024, accessed 8/1/2025.

[63] R. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, 3rd ed. Wiley, 2020.

[64] B. Woods, A. Coravos, and J. D. Corman, "The case for a Hippocratic oath for connected medical devices," *Journal of Medical Internet Research*, vol. 21, no. 3, 2019.

[65] Defense Advanced Research Projects Agency (DARPA), "DARPA AI Cyber Challenge (AIxCC)," https://aicyberchallenge.com/, 2025, accessed 10/20/2025.

[66] The Leona M. and Harry B. Helmsley Charitable Trust, "Nudge BG, Inc. – Helmsley Charitable Trust," https://helmsleytrust.org/grants/nudge-bg-inc-22549441/, accessed 8/1/2025.

[67] D. Wagg, K. Worden, R. Barthorpe, and P. Gardner, "Digital twins: State-of-the-art and future directions for modeling and simulation in engineering dynamics applications," *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, vol. 6, no. 3, 2020.

[68] Defense Advanced Research Projects Agency (DARPA), "Cyber Grand Challenge (CGC)," https://www.darpa.mil/research/programs/cyber-grand-challenge, accessed 10/20/2025.

[69] L. Kustosch, C. Gañán, M. van Eeten, and S. Parkin, "Patching up: Stakeholder experiences of security updates for connected medical devices," in *USENIX Security Symposium*, 2025.