

Medical Device Security Research

# Finding, Scoring, and Disclosing a Vulnerability in a Clinical Device

A Methodology and Disclosure Handbook

Drawn from research on a large-volume infusion pump

---

**Kalyan Gopalam**

M.S. Cybersecurity, Northeastern University

Research conducted under Prof. Kevin Fu and Axel Wirth

[gopalam.k@northeastern.edu](mailto:gopalam.k@northeastern.edu) | [kalyangopalam.com](http://kalyangopalam.com)

Version 1.0 | August 2026

## What this document withholds, and why

Read this page before any other. It defines the document.

The device this research examined is a large-volume infusion pump that remains in clinical use, in hospitals, delivering medication to patients today. That single fact governs everything about how this document is written.

### The rule this document follows

Nothing in these pages tells a reader how to compromise a specific device. No credential values. No interface or protocol path to reach the finding. No reproduction steps. No firmware extraction procedure. No configuration sequences. Those details exist, they are documented, and they live in the coordinated disclosure package held with the manufacturer, where they belong.

What this document contains instead is the part that transfers: **how to think about a clinical device, how to model its threats, how to score a finding so that the score means something to a regulator, how to run a disclosure without harming anybody, and how to decide what to publish.** That methodology is reusable across devices. The specifics of one pump are not, and publishing them would put patients at risk for no research benefit.

I want to be direct about why this is not a limitation of the document but the point of it. In medical device research, restraint is a professional credential. The field's credibility was built by researchers who found serious problems in devices keeping people alive and chose to be careful about how they said so. Working under people who hold that position is where I learned that the discipline is not a constraint imposed on the research from outside; it is part of the research.

A reader who wants the withheld detail has one legitimate route, which is to be the manufacturer or a coordinating body and to ask through the disclosure channel. A reader who wants it for any other reason is the reason it is withheld.

## On the finding itself

The research surfaced an undocumented hardcoded default credential, a weakness of the class catalogued as CWE-798, on a device that had previously been the subject of a published disclosure covering plaintext-stored credentials, CVE-2016-9355, of class CWE-312.

I describe the *weakness class* throughout this document, generically, because CWE-798 is textbook material taught in every application security course and understanding it gives an attacker nothing. I do not describe *this device's instance* of it. The distinction between those two things is the line this document walks, and Chapter 19 is an entire chapter about how I decided where to draw it.

## Status and provenance

- **The methodology, threat model structure, scoring approach, and disclosure reasoning are mine**, developed during research conducted under the supervision of Prof. Kevin Fu and Axel Wirth. These are the load-bearing parts of the document.
- **The regulatory material was verified against current sources in August 2026**. Medical device regulation moves quickly, and several things changed after the research was conducted. Chapter 10 flags each of them. **Verify anything you intend to rely on against the current FDA and CISA publications before you act on it**. Guidance documents are superseded regularly, and a stale citation in a disclosure package damages your credibility with exactly the people you most need to take you seriously.
- **The templates in Part VII are generic**. They contain no device-specific content and are meant to be filled in for whatever device you are working on.
- **Appendix D lists what I have not confirmed**. It is short and honest, and it is at the back where I can find it.

## How to read this document

Eight parts, ordered so that each one answers a question raised by the one before it.

| Part | Title                  | The question it answers                                                             |
|------|------------------------|-------------------------------------------------------------------------------------|
| I    | The Clinical Context   | Why a medical device is not an IT asset, and why that changes everything downstream |
| II   | Threat Modelling       | How to map a device's exposure when patient harm is the impact axis                 |
| III  | Risk Analysis          | How to score a finding so the number means something clinically and to a regulator  |
| IV   | The Regulatory Frame   | What the manufacturer is legally obliged to do, and what that means for you         |
| V    | Coordinated Disclosure | How to report a finding without harming patients or yourself                        |
| VI   | Research Conduct       | How to run the work safely, and how to keep a record that holds up                  |
| VII  | Templates              | Fill-in artefacts for your own research                                             |
| VIII | The Decision Register  | Every judgement call, with what would change my mind                                |

## The three kinds of boxes

### The point

The sentence to remember if you forget everything else on the page. If you are skimming, read only these.

**Where researchers get this wrong**

A place where a well-intentioned researcher causes harm, loses credibility, or exposes themselves legally. Most of these are mistakes I have watched happen or came close to making.

**What else I could have done**

The alternative approach, why it is legitimate, and why I did not take it. Every methodological chapter has one, because a methodology document without rejected options is just a description of what somebody happened to do.

**Conventions**

|                              |                                                                                                                          |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| <b>CWE-<i>nnn</i></b>        | A weakness class from the MITRE Common Weakness Enumeration. A <i>class</i> of mistake, not a specific bug.              |
| <b>CVE-<i>yyyy-nnnnn</i></b> | A specific, publicly disclosed vulnerability in a specific product.                                                      |
| <b>HDO</b>                   | Healthcare delivery organisation. The hospital or clinic operating the device.                                           |
| <b>MDM</b>                   | Medical device manufacturer.                                                                                             |
| <b>CVD</b>                   | Coordinated vulnerability disclosure.                                                                                    |
| <b>PIPS</b>                  | Potential impact to patient safety. Terminology borrowed from the MITRE scoring rubric; see Chapter 8.                   |
| <b>P-<i>nn</i></b>           | A numbered principle from Chapter 3. Appendix C traces each to the chapter that applies it.                              |
| <b>“the device”</b>          | Written generically throughout. Where I mean the specific pump studied, I say so, and I say only what is already public. |

## A note on the writing

I have written this the way I would explain it to somebody sitting next to me: short sentences, ordinary words, and no longer term where a shorter one is exact. Where a piece of jargon is load-bearing I define it once and then use it, because that is how you actually learn a field.

Two things I have avoided on purpose.

The first is passive voice used to hide who decided what. When I write “I chose,” I chose. When I write “the regulation requires,” somebody else decided and I am working inside it. In a document partly about professional judgement, that difference matters to a reader trying to assess the judgement.

The second is confidence I have not earned. Where I am uncertain, the sentence says so rather than a footnote nobody reads. There are several of these, particularly around regulation that shifted after the research was done. They are not weaknesses in the work. Pretending to certainty about a superseded FDA guidance document would be the actual weakness.

*Kalyan Gopalam  
Boston, Massachusetts  
August 2026*



# Contents

|            |                                                        |           |
|------------|--------------------------------------------------------|-----------|
| <b>I</b>   | <b>The Clinical Context</b>                            | <b>1</b>  |
| <b>1</b>   | <b>What This Research Was</b>                          | <b>2</b>  |
| 1.1        | The shape of the work . . . . .                        | 2         |
| 1.2        | What the research produced . . . . .                   | 3         |
| 1.3        | What this document is for . . . . .                    | 3         |
| <b>2</b>   | <b>Why a Clinical Device Is Not an IT Asset</b>        | <b>4</b>  |
| 2.1        | The impact axis changes . . . . .                      | 4         |
| 2.2        | The lifecycle changes . . . . .                        | 5         |
| 2.3        | The people change . . . . .                            | 5         |
| <b>3</b>   | <b>Principles</b>                                      | <b>6</b>  |
| 3.1        | Why numbered principles . . . . .                      | 6         |
| 3.2        | Patient safety . . . . .                               | 6         |
| 3.3        | Honesty about the finding . . . . .                    | 6         |
| 3.4        | Restraint in publication . . . . .                     | 6         |
| 3.5        | Professional conduct . . . . .                         | 7         |
| <b>II</b>  | <b>Threat Modelling</b>                                | <b>8</b>  |
| <b>4</b>   | <b>STRIDE in a Clinical Context</b>                    | <b>9</b>  |
| 4.1        | Why threat model at all . . . . .                      | 9         |
| 4.2        | STRIDE, and the clinical translation . . . . .         | 9         |
| 4.3        | Where STRIDE alone is not enough . . . . .             | 10        |
| <b>5</b>   | <b>Data Flow Diagrams and Trust Boundaries</b>         | <b>12</b> |
| 5.1        | The diagram is the analysis . . . . .                  | 12        |
| 5.2        | A generic infusion pump data flow . . . . .            | 12        |
| 5.3        | Building the diagram without internal access . . . . . | 13        |
| <b>6</b>   | <b>From Model to Attack Surface</b>                    | <b>15</b> |
| 6.1        | Enumeration . . . . .                                  | 15        |
| 6.2        | Reachability before severity . . . . .                 | 15        |
| 6.3        | The weakness class: CWE-798 . . . . .                  | 16        |
| <b>III</b> | <b>Risk Analysis</b>                                   | <b>18</b> |
| <b>7</b>   | <b>Why CVSS Mis-Ranks Medical Findings</b>             | <b>19</b> |
| 7.1        | What CVSS is for . . . . .                             | 19        |
| 7.2        | The specific mismatches . . . . .                      | 19        |

|           |                                                                         |           |
|-----------|-------------------------------------------------------------------------|-----------|
| <b>8</b>  | <b>Scoring a Finding So the Number Means Something</b>                  | <b>21</b> |
| 8.1       | The MITRE rubric . . . . .                                              | 21        |
| 8.2       | A change since this research was conducted . . . . .                    | 21        |
| 8.3       | The three-layer report . . . . .                                        | 22        |
| <b>9</b>  | <b>Safety Risk and Security Risk</b>                                    | <b>23</b> |
| 9.1       | Two disciplines that had to merge . . . . .                             | 23        |
| 9.2       | The standards that bridge the gap . . . . .                             | 23        |
| 9.3       | What happens inside the manufacturer when your report arrives . . . . . | 24        |
| 9.4       | Writing the risk statement . . . . .                                    | 24        |
| <b>IV</b> | <b>The Regulatory Frame</b>                                             | <b>26</b> |
| <b>10</b> | <b>What the Manufacturer Is Obligated to Do</b>                         | <b>27</b> |
| 10.1      | Why a researcher needs to know this . . . . .                           | 27        |
| 10.2      | The statutory floor: section 524B . . . . .                             | 27        |
| 10.3      | The guidance layer, and what changed . . . . .                          | 28        |
| 10.4      | The Secure Product Development Framework . . . . .                      | 28        |
| 10.5      | Postmarket obligations . . . . .                                        | 29        |
| 10.6      | What this means for disclosure timing . . . . .                         | 29        |
| <b>V</b>  | <b>Coordinated Disclosure</b>                                           | <b>31</b> |
| <b>11</b> | <b>The Disclosure Process</b>                                           | <b>32</b> |
| 11.1      | What coordinated disclosure is . . . . .                                | 32        |
| 11.2      | The process, end to end . . . . .                                       | 32        |
| 11.3      | When it goes wrong . . . . .                                            | 34        |
| <b>12</b> | <b>Who to Contact, and CVE Assignment</b>                               | <b>36</b> |
| 12.1      | The ecosystem . . . . .                                                 | 36        |
| 12.2      | CVE assignment . . . . .                                                | 36        |
| 12.3      | Credit . . . . .                                                        | 37        |
| <b>13</b> | <b>Legal and Ethical Footing</b>                                        | <b>38</b> |
| 13.1      | The three questions to settle before testing . . . . .                  | 38        |
| 13.2      | Safe harbour, and its limits . . . . .                                  | 39        |
| 13.3      | Ethics beyond legality . . . . .                                        | 39        |
| <b>VI</b> | <b>Research Conduct</b>                                                 | <b>40</b> |
| <b>14</b> | <b>Lab Safety and Device Sourcing</b>                                   | <b>41</b> |
| 14.1      | The absolute rule . . . . .                                             | 41        |
| 14.2      | Sourcing a device . . . . .                                             | 41        |
| 14.3      | The bench environment . . . . .                                         | 42        |

|                                                                  |           |
|------------------------------------------------------------------|-----------|
| <b>15 The Research Record</b>                                    | <b>43</b> |
| 15.1 Why documentation is a security control . . . . .           | 43        |
| 15.2 What to record . . . . .                                    | 43        |
| 15.3 Handling the sensitive part . . . . .                       | 43        |
| <b>16 Deciding What to Publish</b>                               | <b>45</b> |
| 16.1 The chapter this document is really about . . . . .         | 45        |
| 16.2 The test . . . . .                                          | 45        |
| 16.3 Applying it . . . . .                                       | 45        |
| 16.4 Obscuring is not withholding . . . . .                      | 46        |
| 16.5 The argument against withholding, taken seriously . . . . . | 46        |
| 16.6 What to publish instead . . . . .                           | 46        |
| <b>VII Templates</b>                                             | <b>48</b> |
| <b>17 Threat Model Template</b>                                  | <b>49</b> |
| 17.1 How to use these . . . . .                                  | 49        |
| 17.2 Section 1: device characterisation . . . . .                | 49        |
| 17.3 Section 2: data flow diagram . . . . .                      | 49        |
| 17.4 Section 3: entry point inventory . . . . .                  | 50        |
| 17.5 Section 4: STRIDE analysis . . . . .                        | 50        |
| 17.6 Section 5: coverage statement . . . . .                     | 50        |
| <b>18 Disclosure Report Template</b>                             | <b>51</b> |
| 18.1 Structure . . . . .                                         | 51        |
| 18.2 Main report . . . . .                                       | 51        |
| 18.3 First contact message . . . . .                             | 52        |
| <b>19 Risk Register and Finding Log</b>                          | <b>53</b> |
| 19.1 Finding log . . . . .                                       | 53        |
| 19.2 Risk register . . . . .                                     | 53        |
| <b>VIII The Decision Register</b>                                | <b>54</b> |
| <b>20 Every Judgement Call, In One Place</b>                     | <b>55</b> |
| 20.1 Why collect them . . . . .                                  | 55        |
| 20.2 The register . . . . .                                      | 55        |
| 20.3 The three I would revisit first . . . . .                   | 57        |
| <b>A Standards and Regulatory Reference</b>                      | <b>58</b> |
| A.1 United States regulatory . . . . .                           | 58        |
| A.2 Consensus standards . . . . .                                | 58        |
| A.3 Disclosure and coordination . . . . .                        | 59        |
| <b>B Glossary</b>                                                | <b>60</b> |
| <b>C Principle Traceability</b>                                  | <b>62</b> |

---

|          |                                           |           |
|----------|-------------------------------------------|-----------|
| <b>D</b> | <b>Open Items</b>                         | <b>63</b> |
| D.1      | What I have not confirmed . . . . .       | 63        |
| D.2      | What I would add with more time . . . . . | 63        |

## **PART I**

# **The Clinical Context**

## What This Research Was

### 1.1 The shape of the work

The project was a security assessment of a large-volume infusion pump, conducted in an academic setting under the supervision of two people who work in medical device security professionally. It produced one new finding, of a weakness class already well known to the field, on a device that had been publicly examined before.

Stated that plainly, it sounds modest. I want to explain why the modest description is the accurate one and why the project still mattered, because the gap between those two things is the most useful lesson in the whole document.

#### 1.1.1 What an infusion pump does

An infusion pump delivers fluid, usually medication, into a patient's body at a controlled rate over a controlled period. A large-volume pump is the kind you see on a pole beside a hospital bed. A clinician programmes a dose, a rate, and a duration; the pump delivers it and alarms if something goes wrong.

Modern pumps are networked. They pull drug libraries from a hospital server so that dose limits reflect that hospital's formulary. They report status to nursing stations. They integrate with electronic health records so that what was ordered and what was delivered can be reconciled. Each of those connections is a clinical benefit, and each is also an interface.

##### The point

The features that make a modern infusion pump safer clinically are the same features that give it an attack surface. Networked drug libraries prevent dosing errors, which is a real and substantial safety gain. They also mean the pump trusts something on a network. You cannot analyse the security of a clinical device without understanding that its connectivity exists for a reason, and that reason is usually patient safety.

#### 1.1.2 Why this device, and why a device somebody had already looked at

The pump had a published vulnerability against it from an earlier disclosure cycle, covering credentials stored in plaintext, catalogued as CVE-2016-9355 under weakness class CWE-312.

Choosing a device that had already been through disclosure looks like choosing the easy target. It is closer to the opposite, and the reasoning is worth setting out because it applies to any research programme.

##### **The obvious problems are already gone.**

Anything a previous researcher found and reported has, presumably, been addressed. What remains is what the previous pass missed, which is by definition the harder material.

##### **There is a public baseline to reason against.**

A prior disclosure tells you what the manufacturer's remediation looked like, how they responded, and what their disclosure posture is. That is intelligence you do not have on a device nobody has examined.

##### **Finding something new is a stronger result.**

Reproducing a known CVE demonstrates that you can follow a writeup. Finding an undocu-

mented weakness alongside it demonstrates that you can assess a device independently. Those are different claims about your ability and reviewers can tell them apart.

### It says something about the remediation.

When a device is disclosed against, fixed, and still carries a related weakness of a different class, that is itself a finding about the completeness of the fix. Chapter 3 develops this, because it is the most interesting part of the result.

## 1.2 What the research produced

An undocumented hardcoded default credential: a secret value fixed in the product rather than generated per-device or set by the operator. Weakness class CWE-798.

That sits beside the earlier CWE-312 finding in an instructive way. Both are credential-handling weaknesses. They are different classes: one is about how credentials are *stored*, the other about where credentials *come from*. A remediation programme aimed at the first does not necessarily catch the second, because they live in different parts of a codebase and different parts of a design review.

### The point

The result was not “this device has a hardcoded credential.” The result was “this device had a credential-handling problem, it was disclosed and addressed, and a related credential-handling problem of a different class remained.” The second framing is a finding about the manufacturer’s remediation process, not just about the device, and it is considerably more useful to the field.

## 1.3 What this document is for

Three audiences, and the document tries to serve all three without pretending they want the same thing.

| Reader                                     | What they should take away                                                                                                                                                  |
|--------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| A student starting medical device research | A complete working method: how to model, how to score, how to disclose, and what not to do. Read Parts I to VI in order.                                                    |
| A reviewer assessing my work               | The reasoning behind each judgement call, especially the disclosure decisions in Chapter 19 and the register in Chapter 23.                                                 |
| A practitioner in another domain           | The transferable parts. Threat modelling under a safety constraint, scoring when the standard metric does not fit, and disclosure ethics generalise well beyond healthcare. |

## Why a Clinical Device Is Not an IT Asset

### 2.1 The impact axis changes

In enterprise security, the question behind every risk decision is what happens to data. Confidentiality, integrity, availability, mapped to business consequences: breach cost, regulatory penalty, downtime, reputational damage.

In clinical security, the question is what happens to a patient. That is not a rhetorical elevation of the stakes; it is a different measurement problem, and it changes which findings matter.

| Property        | Enterprise IT                       | Clinical device                               |
|-----------------|-------------------------------------|-----------------------------------------------|
| Worst outcome   | Data breach, financial loss, outage | Patient injury or death                       |
| Availability    | Costed in downtime                  | May itself be the safety-critical property    |
| Integrity       | Corrupted records                   | Incorrect therapy delivered to a person       |
| Confidentiality | Regulatory exposure                 | Real, but often the least severe of the three |
| Patch cadence   | Days to weeks                       | Months to years, and requires validation      |
| Service life    | 3–5 years                           | 10–20 years                                   |
| Reboot to fix   | Usually acceptable                  | May interrupt therapy in progress             |
| Who operates it | Trained IT staff                    | Clinicians, under time pressure, mid-shift    |

#### 2.1.1 Availability inversion

In enterprise systems, availability is generally the least severe of the CIA triad. Losing a service costs money and annoys people. In clinical systems it can invert completely: a pump that stops delivering a vasopressor to an unstable patient is an emergency in a way that a leaked configuration file is not.

This inversion breaks the intuition of experienced enterprise practitioners, and it breaks scoring tools built on that intuition, which is Chapter 7.

##### Where researchers get this wrong

The instinct to demonstrate impact by causing a visible failure is the most dangerous instinct you can bring into this field. In enterprise testing, crashing a service in a lab is a legitimate proof. On a clinical device, “I made it stop” is a demonstration you should think very hard about performing even on a bench unit, and never in any environment where a device could be returned to service without full revalidation. The demonstration proves little that reasoning would not establish, and the downside is unbounded.

## 2.2 The lifecycle changes

An enterprise server is replaced every few years and patched every month. An infusion pump may be in service for fifteen years or more, running software validated at manufacture, on hardware chosen a decade before that.

Three consequences for the researcher:

**Old components are the norm, not negligence.** Finding a decade-old library version in a clinical device is expected. Whether it constitutes a finding depends entirely on whether the specific weakness is reachable and consequential in that device's architecture. Reporting a stale SBOM entry as a vulnerability without establishing reachability wastes the manufacturer's time and teaches them to discount your reports.

**Patching is not a decision the manufacturer makes alone.** A change to software that affects device function may require verification, validation, and in some cases regulatory notification before it can ship. Then the hospital has to schedule the update across a fleet, often device by device, without interrupting care. A ninety-day disclosure deadline that is reasonable for a web application is not obviously reasonable here, and Chapter 14 deals with that tension honestly.

**Legacy devices may be unfixable.** Some devices in service predate the manufacturer's current security programme, or run on hardware that cannot support the fix. The realistic outcome for those is compensating controls at the hospital: network segmentation, access restriction, monitoring. Your report is more useful if it acknowledges this and offers mitigations the HDO can apply without the manufacturer.

### The point

Write your findings so a hospital can act on them even if the manufacturer cannot ship a fix for two years. Compensating controls that an HDO can deploy this month are frequently worth more in real patient-risk reduction than the eventual patch.

## 2.3 The people change

The operator is a nurse or a clinician, working a twelve-hour shift, under time pressure, with multiple patients. They are not a system administrator and it is not their job to be.

This matters for security analysis in a way that is easy to miss. A control requiring the operator to do something extra during an emergency will be bypassed, and the bypass will be correct clinical practice. Any recommendation you make that adds friction to an urgent workflow will fail in the field regardless of its technical merit.

### What else I could have done

I could have assessed the device purely technically, ignoring the clinical workflow, and produced a findings list ranked by technical severity. It is faster, it needs no clinical knowledge, and it is what a generic penetration test would produce. I did not, because a finding ranked without workflow context is regularly ranked wrong: weaknesses reachable only through a workflow nobody performs get overrated, and weaknesses reachable through the routine setup a nurse does forty times a shift get underrated. The cost of the workflow-aware approach is that you have to learn how the device is actually used, which takes time and some humility about what you do not know.

## Principles

### 3.1 Why numbered principles

Everything after this chapter is method. Method without stated principles becomes procedure that people follow when convenient and abandon when inconvenient. These are the commitments the rest of the document is accountable to, and Appendix C traces each to where it is applied.

### 3.2 Patient safety

**P-01 No research activity may create risk to a patient.**

No testing on devices in clinical service. No testing on a clinical network. No device returned to service without the operator knowing it was used for research and revalidating it. *Applied in:* Chapter 17.

**P-02 Patient safety impact is the primary severity axis.**

Where a technical severity score and a clinical harm assessment disagree, the clinical assessment governs the recommendation. *Applied in:* Chapters 7–9.

**P-03 Disclosure timing serves patient safety, not the researcher’s schedule.**

Publication timing is decided by when patients are safer with the information public than private, not by conference deadlines or portfolio needs. *Applied in:* Chapters 14 and 19.

### 3.3 Honesty about the finding

**P-04 A finding is described at the severity it actually has.**

No inflation to make the research look better. No deflation to make the manufacturer comfortable. *Applied in:* Chapters 7–9.

**P-05 Reachability is established before a weakness is called a vulnerability.**

A weak component that no realistic path reaches is a hygiene observation, not a vulnerability, and calling it one damages your credibility on the reports that matter. *Applied in:* Chapter 6.

**P-06 Prior work is credited precisely.**

Where a finding extends published research, say what was published, by whom, and what is new here. *Applied in:* Chapters 1 and 20.

### 3.4 Restraint in publication

**P-07 Published material must not materially assist an attacker against a device in service.**

The test is whether a reader is meaningfully closer to compromising a real device than before. *Applied in:* Chapter 19.

**P-08 Withheld detail is withheld, not merely obscured.**

Vague phrasing that a determined reader can reconstruct is not restraint; it is the appearance of restraint, which is worse because it fools the author. *Applied in:* Chapter 19.

**P-09 The existence of a finding may be disclosed even when detail is not.**

Hospitals and regulators need to know a class of problem exists in order to act. Silence is not the safe default. *Applied in:* Chapters 14 and 19.

### 3.5 Professional conduct

**P-10 The manufacturer is contacted first and given a genuine opportunity to remediate.**

*Applied in:* Chapters 13–14.

**P-11 The research record is complete enough to reconstruct any claim.**

If you cannot show how you reached a conclusion, you cannot defend it when challenged, and you will be challenged. *Applied in:* Chapter 18.

**P-12 Legal footing is established before testing begins, not after.**

*Applied in:* Chapter 16.

#### The point

Twelve principles. The one that does the most work is **P-05**: establish reachability before you call something a vulnerability. Most of the friction between security researchers and medical device manufacturers traces to reports of weaknesses that were never reachable, and every such report makes the manufacturer slower to act on the next one.

## **PART II**

# **Threat Modelling**

## STRIDE in a Clinical Context

### 4.1 Why threat model at all

A threat model is a structured answer to four questions: what are we building, what can go wrong, what are we going to do about it, and did we do a good job. On a device you did not build, the first question becomes “what is this, actually,” and answering it properly is most of the work.

The alternative to threat modelling is poking at interfaces until something breaks. That finds things. It finds them unsystematically, it gives you no way to know what you missed, and it produces a findings list you cannot defend as complete. A threat model gives you coverage you can argue for.

#### The point

The value of a threat model to a researcher is not that it finds vulnerabilities. It is that it tells you where you have and have not looked, which is the difference between “I found three things” and “I examined eleven trust boundaries and found three things at these two.”

### 4.2 STRIDE, and the clinical translation

STRIDE enumerates six threat categories. It is well suited to this work because it is exhaustive over a useful space and because it is one of the approaches the FDA’s premarket cybersecurity guidance names as a common threat modelling method, which means a report built on it lands in a familiar shape for a regulatory reader.

The standard definitions are enterprise-flavoured. The right-hand column is the translation that matters clinically.

| Category               | Standard meaning              | What it means on a clinical device                                                                                      |
|------------------------|-------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| Spoofing               | Pretending to be someone else | A rogue endpoint impersonating the drug-library server, so the pump accepts dose limits from the wrong source           |
| Tampering              | Unauthorised modification     | Alteration of therapy parameters, drug library entries, or delivered dose                                               |
| Repudiation            | Denying an action occurred    | Inability to reconstruct what dose was actually delivered, which is a clinical and legal problem after an adverse event |
| Information disclosure | Unauthorised read             | Patient data exposure, but also configuration and credential exposure that enables the other five                       |
| Denial of service      | Loss of availability          | Therapy interruption, alarm suppression, or the device made unavailable during an emergency                             |

| Category               | Standard meaning                | What it means on a clinical device                                                                |
|------------------------|---------------------------------|---------------------------------------------------------------------------------------------------|
| Elevation of privilege | Gaining unauthorised capability | Moving from an operator context to a service or engineering context that can change safety limits |


#### 4.2.1 The two categories that behave differently here

**Repudiation matters more than it does in enterprise IT.** When a patient is harmed, there is a clinical investigation and frequently a legal one. The device’s record of what it delivered is evidence. A weakness that lets that record be altered or that leaves it ambiguous is serious in a way that has no clean enterprise analogue, and it is routinely underweighted by researchers coming from application security.

**Denial of service can be the most severe category, not the least.** Covered in Chapter 2, and it is worth repeating because the instinct is strong and wrong.

#### 4.3 Where STRIDE alone is not enough

STRIDE asks what an attacker can do. It does not ask what happens to the patient when they do it. That second question is a safety hazard analysis, and the two have to be run together.



The output you want is not “an attacker can tamper with parameter X.” It is “an attacker who can tamper with parameter X can cause the device to deliver a dose outside the intended range, and the clinical consequence of that for the drugs this pump commonly delivers is Y.” The first sentence is a security finding. The second is something a manufacturer’s risk management process and a regulator can act on.

##### Where researchers get this wrong

Do not invent the clinical half. If you are not clinically trained, you cannot assess what a dosing deviation does to a patient, and guessing produces confident nonsense that a clinical reviewer will spot instantly. The correct move is to state the technical finding precisely, state the therapy-relevant parameter it affects, and explicitly flag that clinical impact assessment is required from a qualified reviewer. Naming the boundary of your competence is a strength in a report. Stepping over it silently is how researchers lose the room.

##### What else I could have done

**Attack trees** start from an attacker goal and decompose it into paths, which is better than STRIDE at showing how several small weaknesses chain into one serious outcome. I would use them alongside STRIDE on a device where chaining is the main concern. **PASTA**

is a heavier, business-impact-driven process, appropriate for a manufacturer's internal programme, oversized for an external assessment. **Kill-chain models** are attacker-lifecycle oriented and fit intrusion analysis better than device assessment. STRIDE won here because it is systematic, teachable, produces defensible coverage claims, and is a method the regulatory audience recognises.

## Data Flow Diagrams and Trust Boundaries

### 5.1 The diagram is the analysis

A data flow diagram is not documentation of the analysis. Drawing it *is* the analysis, because you cannot draw one without deciding what the components are, what talks to what, and where trust changes. Those decisions are where the findings come from.

Four element types, and that is all:

| Element         | What it is                                                                                                                    |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------|
| External entity | Something outside your analysis boundary that interacts with the system: a clinician, a hospital server, a service technician |
| Process         | Something that transforms data: a firmware component, a control loop, a parser                                                |
| Data store      | Something that holds data at rest: configuration memory, an event log, a drug library                                         |
| Data flow       | Movement of data between the above                                                                                            |


And one annotation that carries all the weight: the **trust boundary**, drawn wherever data crosses from one level of trust to another.

#### The point

Every finding you will make lives on a trust boundary. Data that crosses a boundary without being validated, authenticated, or authorised on the far side is the general shape of every vulnerability in this document and most others. Draw the boundaries first and the analysis largely tells you where to look.

### 5.2 A generic infusion pump data flow

The diagram below is deliberately generic. It reflects the architecture class that networked large-volume pumps share, at the abstraction level found in published literature and manufacturer marketing material. It is not a map of any specific product's internals and it identifies no specific interface.



Two boundaries, and they are not equally interesting.

**Boundary A** separates the outside world from the device. Everything crossing it is attacker-influenced in some scenario. This is where most assessment effort goes and where most published medical device findings live.

**Boundary B** separates the device's external-facing components from the therapy-critical core. This is the boundary that determines whether a finding is embarrassing or dangerous. A weakness that lets an attacker reach boundary A is a problem. A weakness that lets them cross boundary B is a patient-safety event. **The single most valuable question in a medical device assessment is what enforces boundary B.**

#### The point

Ask what separates the connected, complex, attacker-reachable part of the device from the part that decides how much drug to deliver. If the answer is a software check in the same execution context, that is a finding in itself, independent of whether you can currently reach it. If the answer is a separate processor with a narrow, validated interface, the device was designed by people who thought about this.

### 5.3 Building the diagram without internal access

You will usually not have design documentation. The diagram is reconstructed from what is observable, and every element should be traceable to evidence.

| Source                              | What it gives you                                                                           |
|-------------------------------------|---------------------------------------------------------------------------------------------|
| Operator manual                     | Workflows, user roles, configuration options, what the device is meant to do                |
| Service manual                      | Maintenance procedures, service roles, and the existence of privileged modes                |
| Regulatory submission summaries     | Public clearance documents describe intended use and sometimes architecture at a high level |
| Manufacturer security documentation | MDS2 forms and security white papers state what controls the manufacturer claims            |
| Published SBOM                      | Component inventory, where available under current expectations                             |
| Prior disclosures                   | What has been found before, and by implication what part of the device it lived in          |
| Standards the device claims         | Conformance claims tell you what the design was supposed to achieve                         |

#### Where researchers get this wrong

Mark inferred elements as inferred. A diagram that mixes documented architecture with your guesses, without distinguishing them, will eventually put you in a meeting asserting something about a manufacturer's product that is not true. I annotate every element with its evidence source and use a distinct style for inferred components. When a manufacturer corrects an inference, that is a productive conversation. When they catch you presenting an inference as fact, the conversation is over.

## From Model to Attack Surface

### 6.1 Enumeration

With boundaries drawn, enumerate every place data crosses one. For the research this document draws on, that enumeration produced six or more distinct entry points, which is typical for a networked clinical device of this generation.

For each entry point, record the same fields. Consistency is what makes the list analysable rather than merely long.

| Field               | Content                                                                      |
|---------------------|------------------------------------------------------------------------------|
| Identifier          | Stable reference used everywhere else                                        |
| Boundary            | Which trust boundary it crosses                                              |
| Reachability        | Physical, adjacent-network, network, or requires prior access                |
| Actor               | Who legitimately uses it: clinician, biomedical engineer, server, technician |
| Authentication      | What is required to use it, as documented and as observed                    |
| Authorisation       | What the actor may do once authenticated                                     |
| Reaches boundary B? | Whether this path can influence therapy-critical state                       |
| Evidence            | How you know all of the above                                                |

The “reaches boundary B” column is the one that drives prioritisation. An entry point that cannot influence therapy-critical state, however weakly protected, is a lower priority than a well-protected one that can.

### 6.2 Reachability before severity

**P-05** says a weakness is not a vulnerability until reachability is established. This is the discipline that separates a useful report from noise, and it is worth being precise about what “reachable” means.

| Level         | Meaning                                           | How to report it                                     |
|---------------|---------------------------------------------------|------------------------------------------------------|
| Theoretical   | The weakness exists in a component                | Hygiene observation. Do not call it a vulnerability. |
| Reachable     | A path exists from an entry point to the weakness | A finding. Describe the path class, not the recipe.  |
| Exploitable   | The path can be traversed with realistic effort   | A vulnerability. Severity assessment applies.        |
| Consequential | Traversing it produces clinical impact            | The finding that gets prioritised.                   |

**The point**

Most researcher-manufacturer friction is a disagreement about which row a finding is in. The researcher reports at “exploitable”; the manufacturer assesses at “theoretical” and dismisses it. Doing this classification yourself, explicitly, in the report, and stating your evidence for the level you claim, removes most of that argument before it starts.

### 6.3 The weakness class: CWE-798

The finding was a hardcoded credential. Here is the class, generically, because understanding it is standard security education and confers no attack capability against any particular device.

#### 6.3.1 What it is

A credential fixed in the product rather than generated per-device or established by the operator. It ships identical in every unit. The value may live in firmware, in a configuration default, or in documentation.

The security property that fails is **uniqueness**. A per-device credential means compromising one device tells you nothing about the next. A hardcoded credential means the opposite: one disclosure compromises the entire fleet simultaneously, and the fleet may be tens of thousands of devices across hundreds of hospitals.

#### 6.3.2 Why it persists in embedded and clinical products

Not because engineers are careless. Because the alternatives are genuinely harder, and understanding why is what lets you write a useful remediation recommendation rather than a scolding.

**Per-device credentials require provisioning infrastructure.**

Something must generate, install, record, and later retrieve them. That is a manufacturing system, a secure database, and a support process, not a code change.

**Service access must work when everything else does not.**

A technician repairing a device with a failed network and a corrupted configuration still needs to get in. That requirement pushes designs toward a credential that always works, which is exactly the weakness.

**Long service life outlasts the design assumptions.**

A credential scheme chosen when the device was air-gapped becomes indefensible when the same device is networked a decade later. The design was not wrong when made; the environment changed underneath it.

**Support economics.**

Per-device credentials generate support calls, and lost-credential recovery is itself a security-sensitive process that has to be designed.

**The point**

Write remediation recommendations that engage with why the weakness exists. “Do not hardcode credentials” is not advice; the engineers know. “Per-device credentials generated at manufacture, recorded in the service database, with a documented recovery procedure for field failure” is advice, because it addresses the requirement that produced the weakness.

### 6.3.3 The relationship to CWE-312

CWE-312 is cleartext storage of sensitive information: the credential is stored without protection. CWE-798 is a hardcoded credential: the credential is fixed in the product.

They are frequently confused and they are not the same. Consider the four combinations:

|            | Stored protected | Stored in cleartext |
|------------|------------------|---------------------|
| Per-device | The target state | CWE-312 only        |
| Hardcoded  | CWE-798 only     | CWE-798 and CWE-312 |

A remediation programme aimed at cleartext storage, encrypting credentials at rest, moves you from the bottom-right cell to the bottom-left. It closes CWE-312 completely and does nothing whatsoever about CWE-798, because the credential is still the same on every device.

#### The point

This is the structural insight the research produced, and it generalises well beyond one pump. **When a device has been disclosed against and remediated, check the neighbouring weakness classes.** A fix scoped to the reported class will often leave a related class untouched, because remediation is usually scoped to the report rather than to the underlying design decision. Reviewing what a fix did *not* cover is a genuinely productive research strategy and it is underused.

#### What else I could have done

I could have pursued breadth, enumerating as many weaknesses as possible across the whole device. It produces a longer findings list and looks more impressive on a slide. I went narrow instead, on credential handling specifically, informed by the prior disclosure. Narrow won because depth in one area let me establish reachability properly, and one finding at “consequential” is worth more to a manufacturer than fifteen at “theoretical.” If I were assessing a device with no disclosure history and no prior to anchor on, I would start broad and narrow after the first pass told me where the weak design thinking was.

### PART III

## Risk Analysis

## Why CVSS Mis-Ranks Medical Findings

### 7.1 What CVSS is for

The Common Vulnerability Scoring System, maintained by FIRST, produces a number from 0.0 to 10.0 describing the severity of a vulnerability. It is built from metric groups: Base, which describes intrinsic properties; Threat or Temporal, which describes the current exploitation situation; and Environmental, which adjusts for a specific deployment.

It is genuinely useful. It gives a shared vocabulary across researchers, vendors, and defenders, and it makes prioritisation comparable across products. The problem is not that CVSS is bad. The problem is that it was designed for enterprise information technology and a clinical device is not that.

#### The point

CVSS measures impact to a system. Medical device severity is about impact to a person. Those correlate loosely and diverge exactly where it matters most, which is why a raw CVSS score on a clinical device should never be the last word in your report.

### 7.2 The specific mismatches

#### **The impact metrics measure the wrong things.**

Confidentiality, Integrity, and Availability are scored against the *system*. There is no metric for patient harm. A vulnerability that alters delivered therapy and one that corrupts a log file can both score Integrity: High, and they are not remotely comparable in consequence.

#### **Availability is systematically underweighted.**

CVSS treats availability as generally the least severe impact, reflecting enterprise reality. On a device delivering time-critical therapy this can invert entirely, as Chapter 2 set out.

#### **Attack Vector does not describe clinical reachability.**

“Physical” access sounds like a strong mitigating factor and scores accordingly. In a hospital, a pump sits in a corridor, in a patient room, in an unlocked storage bay. Physical proximity is not the barrier the metric assumes. Conversely, “Network” may overstate exposure on a properly segmented clinical network.

#### **Privileges Required does not map onto clinical roles.**

The clinician, biomedical engineer, and service technician roles do not sort cleanly into none, low, and high, and the mapping choice materially moves the score.

#### **Scope does not capture fleet effects.**

A hardcoded credential compromises every device of that model simultaneously. CVSS Scope describes impact beyond the vulnerable component’s security authority, which is a related but distinctly different idea, and it does not express “and this is true of forty thousand units.”

#### 7.2.1 The empirical evidence that this is a real problem

This is not a theoretical complaint. Comparisons of medical device vulnerability scoring between the National Vulnerability Database and ICS-CERT advisories have found that only a minority of vector assignments matched, with the largest disagreements in precisely the confidentiality,

integrity, and availability metrics. When two competent bodies score the same medical device vulnerabilities and disagree on most of the vectors, the instrument is not fit for the material without additional guidance.

#### Where researchers get this wrong

Reporting a bare CVSS number for a medical device finding, with no clinical framing, invites two failure modes. A high score with no patient-harm reasoning reads as inflation, and the manufacturer discounts it. A low score, because the metrics did not capture the clinical reality, gets the finding deprioritised into a backlog it never leaves. Either way the number did your argument harm. Score it, then explain what the score does not capture.

## Scoring a Finding So the Number Means Something

### 8.1 The MITRE rubric

The FDA recognised this problem and contracted MITRE to address it. The result is the *Rubric for Applying CVSS to Medical Devices*, developed with a working group spanning FDA, manufacturers, healthcare delivery organisations, and safety and risk experts.

It does not replace CVSS. It structures how you apply it, as a series of questions and decision trees at each point where you must assign a metric value. That structure buys three things:

1. **Repeatability.** Two analysts working the same finding through the rubric are much more likely to produce the same vector, which is the property that was demonstrably missing.
2. **Clinical framing.** The questions surface medical-device-specific considerations that raw CVSS does not prompt for.
3. **Regulatory standing.** FDA qualified the rubric as a Medical Device Development Tool. A score produced through it stands on a process the agency has reviewed and accepted, rather than on your individual judgement.

#### The point

That last property is why this matters to you as a researcher, not just to manufacturers. A finding scored through an FDA-qualified process is much harder to wave away than one scored by a student with a web calculator. Use the rubric, say you used it, and record your answers to the decision points.

#### 8.1.1 The PIPS flag

The rubric carries a concept worth adopting even if you use nothing else from it: **potential impact to patient safety**. Where an answer suggests the vulnerability might adversely affect patient safety, the rubric explicitly flags that a safety-oriented hazard analysis may be required, and that the issue may fall under postmarket reporting obligations.

This is the hinge between security analysis and safety analysis, and it is the single most important annotation in a medical device finding. A finding with a PIPS flag is on a different track from one without: different urgency, different reviewers, potentially different regulatory consequences for the manufacturer.

### 8.2 A change since this research was conducted

#### Verify this before you rely on it

The research this document draws on used CVSS v3.1, which was correct practice at the time. The situation has since moved: FDA recognised CVSS v4.0 in late 2023, with a transition period for v3.1 that runs into 2027. **If you are scoring a finding today, check which version the manufacturer and the relevant regulator currently expect, and consider providing both vectors during the transition.** Providing both costs you ten minutes and removes an entire category of argument about whether your severity assessment is current.

CVSS v4.0 addresses several of the complaints in Chapter 7 directly. It introduces supplemental metrics including a Safety metric, which is a substantial improvement for this domain, and it restructures the impact metrics to separate the vulnerable system from subsequent systems. It does not eliminate the need for clinical framing, but it narrows the gap.

### 8.3 The three-layer report

My recommendation, and what I would do on any future finding:

| Layer                        | What it contains                                                                                                                                                      |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Technical severity           | CVSS vector and score, produced through the MITRE rubric, with your decision-point answers recorded. Provide both v3.1 and v4.0 vectors during the transition period. |
| Clinical framing             | What the vulnerability affects in therapy terms, the PIPS determination, and an explicit statement of where clinical review by a qualified person is required.        |
| Fleet and deployment context | How many devices share the weakness, whether compromise of one implies compromise of others, and what compensating controls an HDO could apply now.                   |

The third layer is the one researchers most often omit and the one hospitals most need. A hospital cannot patch the device. They can segment the network, restrict physical access, change a procedure, or add monitoring. Telling them what would help is frequently the highest-value paragraph in the entire report.

#### What else I could have done

**Skip CVSS and write a narrative severity assessment.** More honest in some ways, since it does not launder judgement through a number. It fails because manufacturers and regulators run on comparable scores; a narrative cannot be sorted into a queue. **Use a proprietary healthcare scoring system.** Several vendors offer them and some are good. They are not recognised by the regulator and not portable across organisations. **Use CVSS environmental metrics to encode clinical context.** Partially works, and I do use them, but the environmental group describes a specific deployment and you usually do not know the hospital's configuration. CVSS through the rubric, plus explicit clinical framing, gives comparability where it exists and honesty where it does not.

## Safety Risk and Security Risk

### 9.1 Two disciplines that had to merge

Medical device engineering has a mature, decades-old risk management discipline in ISO 14971, which governs risk management for medical devices generally. It thinks in terms of hazards, probability of occurrence, severity of harm, and risk control.

Security risk does not fit that frame cleanly, and the reason is precise:

#### The point

Safety risk analysis asks how *likely* a hazard is to occur, and answers with something like an actuarial estimate. Security risk analysis cannot, because an attacker is not a random process. The right question is not how probable the failure is but how *exploitable* the weakness is. This shift, from probability to exploitability, is the central conceptual move in applying risk management to device security, and current FDA guidance reflects it explicitly.

An adversary chooses the worst moment. A probability estimate built on component failure rates says nothing useful about a person who is deliberately trying to make the component fail.

### 9.2 The standards that bridge the gap

| Standard       | What it covers and why it matters to you                                                                                                                            |
|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ISO 14971      | Risk management for medical devices. The foundational safety risk framework the manufacturer already runs. Your finding has to land in this process to be actioned. |
| IEC 62304      | Medical device software lifecycle processes. Governs how software is developed, maintained, and changed. Explains why a patch is not a quick fix.                   |
| AAMI TIR57     | Principles for medical device security risk management. The bridging document that explains how security risk relates to ISO 14971 safety risk.                     |
| ANSI/AAMI SW96 | Security risk management for device manufacturers, building on TIR57.                                                                                               |
| IEC 81001-5-1  | Security in the health software lifecycle. Increasingly the reference for secure development in this sector.                                                        |
| ISO 13485      | Quality management systems for medical devices, now incorporated by reference into the FDA's quality system regulation. Chapter 10.                                 |

You do not need to have read all of these to do research. You need to know they exist, because when you report a finding, the manufacturer's response is shaped by them, and a report that anticipates that shape gets actioned faster.

### 9.3 What happens inside the manufacturer when your report arrives

Understanding this changes how you write.

1. **Intake and triage.** Someone decides whether this is real and whether it is theirs. A report that clearly identifies the product, version, and weakness class survives this step; a vague one dies here.
2. **Reproduction.** They try to confirm it. Anything you can give that helps them reproduce it *internally* shortens this dramatically. This is where the detail you withheld from publication goes.
3. **Security risk assessment.** Scored, usually with the rubric. Your score is an input, not the answer.
4. **Safety impact assessment.** If there is a PIPS flag, this runs as a hazard analysis under their ISO 14971 process, and it may involve clinical staff.
5. **Remediation decision.** Fix, mitigate, or accept, and on what timeline. Constrained by IEC 62304 change control, validation requirements, and possibly regulatory notification.
6. **Communication.** Advisory to customers, coordination with CISA and FDA where warranted, and eventually a public advisory.

#### The point

Steps 3 and 4 are where your report either succeeds or stalls. A report that hands them a defensible severity assessment and a clear statement of what therapy-relevant behaviour is affected lets those steps run fast. A report that says “critical, fix immediately” without that material makes them redo your analysis from scratch, which takes months and makes them less inclined to trust your next report.

### 9.4 Writing the risk statement

The single most useful artefact you can produce is one paragraph in a fixed shape:

| Clause                | Content                                                                               |
|-----------------------|---------------------------------------------------------------------------------------|
| Weakness              | The class, with CWE reference                                                         |
| Location              | Which component or function, at the abstraction level you can support                 |
| Reachability          | Which entry point and trust boundary, and at what level from Chapter 6's table        |
| Precondition          | What the attacker needs first: proximity, network position, credentials, prior access |
| Capability gained     | What they can then do, technically                                                    |
| Therapy relevance     | Which therapy-affecting parameter or function is implicated, or explicitly none       |
| Clinical impact       | Stated, or explicitly flagged as requiring qualified clinical review                  |
| Fleet scope           | Whether this is per-device or model-wide                                              |
| Compensating controls | What an HDO can do before a fix exists                                                |

Nine clauses. Written properly it is one paragraph and it answers every question the manufacturer's triage process is going to ask, in the order they ask it.

#### What else I could have done

**Write findings in a generic vulnerability report format**, the way you would for a bug bounty. Faster and familiar. It omits therapy relevance, fleet scope, and compensating controls, which are the three fields that make a medical device report useful to the people receiving it. **Write to a formal standard template** such as those used in regulatory submissions. More rigorous, and wrong for an external researcher: you are not making a submission, and adopting that formality suggests a role you do not have. The nine-clause paragraph is the smallest thing that carries everything the recipient needs.

## **PART IV**

# **The Regulatory Frame**

## What the Manufacturer Is Obligated to Do

### 10.1 Why a researcher needs to know this

You are not the regulated party. The manufacturer is. So why spend a chapter on regulation? Three reasons, and they are practical rather than academic.

1. **It tells you what leverage exists.** A manufacturer who is slow to respond may have obligations they are not meeting. Knowing which ones, precisely, changes the conversation and changes who else you can reasonably involve.
2. **It tells you what is genuinely hard for them.** Not every delay is indifference. Some remediation timelines are long because validation and change control legitimately take that long. Knowing the difference keeps you from burning a relationship over a constraint that is real.
3. **It makes your report land in the right shape.** A report that speaks the language of the manufacturer's obligations gets routed correctly on the first pass.

#### Verify before relying on any of this

This chapter was checked against current sources in August 2026 and the landscape moved substantially between when the research was conducted and now. Guidance documents get superseded, statutory requirements get interpreted, and effective dates arrive. **Treat everything below as a map of the terrain rather than a current legal citation, and verify against FDA and CISA publications before you put any of it in a document that goes to a manufacturer.** A researcher citing a superseded guidance version signals they have not kept up, and that impression is expensive.

### 10.2 The statutory floor: section 524B

The binding legal requirement, as distinct from guidance, comes from section 524B of the Federal Food, Drug, and Cosmetic Act, added by the Consolidated Appropriations Act of 2023 and effective from March 2023. It applies to “cyber devices” and it obliges the manufacturer, in a premarket submission, to:

- Submit a plan to monitor, identify, and address postmarket cybersecurity vulnerabilities and exploits, including through coordinated vulnerability disclosure.
- Design, develop, and maintain processes and procedures providing reasonable assurance the device and related systems are cybersecure, and make available postmarket updates and patches.
- Provide a software bill of materials, including commercial, open source, and off-the-shelf components.
- Comply with such other requirements as the Secretary may require.

#### The point

The first bullet is the one that matters to you directly. **Manufacturers of cyber devices**

**are statutorily required to have a coordinated vulnerability disclosure process.** If you cannot find a way to report a vulnerability to a manufacturer whose device was cleared under this requirement, that is not merely inconvenient; it is a gap in something they were obliged to establish, and it is legitimate to say so and to escalate to a coordinating body.

Noncompliance with 524B is treated as a prohibited act under the FD&C Act, and FDA has authority to refuse to accept submissions that do not meet it.

### 10.3 The guidance layer, and what changed

Above the statute sits FDA guidance, which is non-binding but describes what the agency expects and what a submission will be reviewed against.

The relevant document is the premarket cybersecurity guidance. Its history matters because citing the wrong version dates you:

| Version        | Status                                                                                                                                                                                                                                                                                                                                              |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| September 2023 | Final guidance, superseded                                                                                                                                                                                                                                                                                                                          |
| June 2025      | Final guidance, superseded                                                                                                                                                                                                                                                                                                                          |
| Early 2026     | Current at the time of writing. Retitled to <i>Quality Management System Considerations and Content of Premarket Submissions</i> , aligning terminology with the Quality Management System Regulation which took effect in February 2026 and incorporates ISO 13485:2016 by reference into 21 CFR Part 820. Adds a section addressing section 524B. |

Two substantive shifts in the current version are worth understanding:

**Cybersecurity is embedded in the quality system, not adjacent to it.** The retitling is not cosmetic. The expectation is that security lives inside risk management, design controls, validation, and postmarket surveillance, rather than as a separate deliverable produced for the submission. For you, this means a manufacturer with a mature quality system has somewhere structured to put your report, and one without will struggle regardless of goodwill.

**Exploitability replaces probability.** As Chapter 9 set out, the framework asks how likely a vulnerability is to be exploited rather than how likely a hazard is to occur. This is the conceptual change that lets security risk sit inside a safety risk framework without distorting either.

### 10.4 The Secure Product Development Framework

FDA guidance recommends manufacturers adopt an SPDF, defined as a set of processes that reduces the number and severity of vulnerabilities across the device lifecycle. It covers threat modelling, security architecture, testing, and postmarket vulnerability management.

Threat modelling is named explicitly, with STRIDE among the recognised approaches. This is a useful fact for a researcher: **the threat model you built in Part II is the same artefact the manufacturer is expected to have built.** If your model identifies a boundary theirs missed, that is a finding about their SPDF, not just about the device, and it is a more consequential observation than a single vulnerability.

## 10.5 Postmarket obligations

Once a device is on the market, the manufacturer has ongoing duties: monitoring for vulnerabilities, assessing them, remediating them, and communicating. Where a vulnerability presents a risk of patient harm, reporting obligations may be triggered.

| Obligation               | What it means when your report arrives                                                       |
|--------------------------|----------------------------------------------------------------------------------------------|
| Vulnerability monitoring | They are expected to be watching for reports like yours, including from external researchers |
| Risk assessment          | Your finding gets scored and assessed for patient safety impact                              |
| Remediation              | Fix or mitigate, on a timeline driven by assessed risk                                       |
| Customer communication   | Advisories to HDOs operating the device                                                      |
| Regulatory reporting     | Where patient harm is implicated, reporting pathways may apply                               |
| Coordination             | Engagement with CISA and FDA on advisories where warranted                                   |

### The point

The manufacturer's obligations create a channel that exists whether or not they want to talk to you. A researcher who reports carefully, documents the attempt, and escalates through coordinating bodies when a manufacturer is unresponsive is operating inside a system designed to handle exactly that situation. You are not dependent on the manufacturer's goodwill, and knowing that changes how you handle a stonewall.

## 10.6 What this means for disclosure timing

The regulatory frame is the honest answer to why medical device disclosure timelines are longer than software timelines.

| Step                        | Why it takes time                                                                  |
|-----------------------------|------------------------------------------------------------------------------------|
| Reproduction and triage     | Requires the specific hardware and software revision, which may need sourcing      |
| Risk and safety assessment  | Runs through a formal process, may need clinical input                             |
| Development of the fix      | Governed by IEC 62304 change control                                               |
| Verification and validation | The change must be shown not to break device function; this is where the months go |
| Regulatory consideration    | Some changes require notification or submission                                    |
| Customer deployment         | Hospitals update fleets device by device, scheduled around care                    |

### Where researchers get this wrong

Two opposite errors, and both are common. The first is treating a long timeline as evidence

of bad faith and publishing early to force action, which can leave hospitals exposed with no fix available. The second is accepting any timeline the manufacturer proposes, indefinitely, which leaves patients at risk while a fix that is not actually being worked on remains perpetually forthcoming. The way through is to ask for a specific plan with dates and milestones, and to judge good faith by whether milestones are met rather than by how long the plan is. A manufacturer who commits to eighteen months and reports progress every month is acting in good faith. One who says “we are looking into it” for six months is not, regardless of the nominal timeline.

#### What else I could have done

**Ignore regulation entirely and treat this as a technical exercise.** Faster, and it is what most security research does. It costs you the ability to judge whether a manufacturer’s response is reasonable, which is the single judgement that most affects whether a disclosure goes well. **Go deep into regulatory analysis and write a compliance assessment.** Beyond my competence and beyond my role: I am not a regulatory affairs professional and a researcher who presents as one will be caught. The middle position, know the terrain well enough to navigate it and be explicit about not being a lawyer or a regulatory specialist, is the honest and useful one.

## **PART V**

# **Coordinated Disclosure**

## The Disclosure Process

### 11.1 What coordinated disclosure is


Coordinated vulnerability disclosure is an agreement, usually informal, between the person who found a problem and the people who can fix it: the finder gives the fixer time and information; the fixer works the problem and communicates; and public disclosure happens when it does more good than harm.

It is not the same as a bug bounty. A bounty adds money and eligibility rules on top. Plenty of credible disclosure programmes pay nothing, and in medical devices most pay nothing.

#### The point

The purpose of coordination is not politeness to the manufacturer. It is that patients are safer when a fix exists before the problem is public. Every decision in this part follows from that, and where manufacturer convenience and patient safety diverge, patient safety wins, including when that means publishing over a manufacturer's objection.

### 11.2 The process, end to end



#### 11.2.1 Step 1: prepare before you contact anybody

Do the analysis fully first. A report sent early, on an incomplete understanding, wastes the manufacturer's triage capacity and costs you credibility you will want later.

Have ready: the nine-clause risk statement from Chapter 9, the severity assessment from Chapter 8, internal reproduction detail sufficient for their engineers, your evidence record, and a decision about what you intend to publish eventually.

#### Where researchers get this wrong

Deciding your publication plan *after* the manufacturer relationship starts is a mistake. Tell them at first contact what you intend to publish and roughly when. A manufacturer who learns at month five that you plan a conference talk will react badly and reasonably. One who knew from day one can plan around it, and the conversation about timing happens when nobody is under pressure.

### 11.2.2 Step 2: finding the channel

In order of preference:

| Channel                       | Notes                                                                                                                              |
|-------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| Product security team / PSIRT | The right answer where it exists. Larger manufacturers have one.                                                                   |
| <code>security.txt</code>     | Increasingly present; check the manufacturer's domain.                                                                             |
| Published VDP page            | States scope, expectations, and often safe harbour language. Read it fully before reporting.                                       |
| CISA as coordinator           | Appropriate when no channel exists, when the manufacturer is unresponsive, or when you want a third party involved from the start. |
| Sector ISAC                   | Health-sector information sharing bodies can route reports and reach HDOs.                                                         |
| General contact address       | Last resort. Expect it to be handled by someone who does not know what to do with it.                                              |

#### The point

For medical devices and industrial control systems, CISA operates as a coordinating body and as the root CVE Numbering Authority for that space. That is the escalation path when direct contact fails, and it is a legitimate first contact if you would rather not approach the manufacturer alone. As a student researcher, involving a coordinator from the outset is a reasonable choice and I would not treat it as a failure of nerve.

### 11.2.3 Steps 3 to 5: reporting and coordinating

Send the full package. Ask for acknowledgement within a stated period, commonly five business days. State your intended publication approach.

Then be genuinely useful. Answer questions quickly. Offer to test a candidate fix if you can do so safely. If their severity assessment differs from yours, engage with the reasoning rather than restating your score; they know things about the device that you do not, and sometimes they are right.

### Where researchers get this wrong

Do not go quiet and reappear at the deadline. Coordination that consists of a report, six months of silence, and a publication is not coordination. Check in on an agreed cadence. Ask for status. Document each exchange. If the relationship later goes wrong, that record is what shows you behaved reasonably, and if it goes right, it is what makes the next report easier.

#### 11.2.4 Steps 6 to 8: remediation and publication

Remediation is theirs. Your job is to stay available, track whether committed milestones are met, and be honest with yourself about whether progress is real.

Publication timing follows **P-03**: publish when patients are safer with the information public than private. In practice that usually means after a fix is available, or after mitigations exist that hospitals can apply, or after it has become clear that no fix is coming and hospitals need to know so they can protect themselves.

### 11.3 When it goes wrong

| Situation                      | What to do                                                                                                                                      |
|--------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| No acknowledgement             | Resend once, through a second channel. Then go to a coordinator. Document both attempts.                                                        |
| Acknowledged, then silence     | Ask for a named contact and a milestone plan. If neither materialises, escalate to a coordinator.                                               |
| Severity disputed              | Engage on the technical merits. If unresolved, a coordinator can provide an independent assessment, which is often what breaks the deadlock.    |
| Legal threat                   | Stop corresponding directly. Get counsel; your university has some. Notify the coordinator. Do not negotiate alone and do not respond in anger. |
| Fix promised but not delivered | Track it. Missed milestones without explanation are the signal that matters, not elapsed time.                                                  |
| No fix possible                | Shift entirely to compensating controls hospitals can apply, and publish those. This is a legitimate and valuable outcome.                      |

### What else I could have done

**Full disclosure**, publishing everything immediately, is defended by serious people on the grounds that it forces action and that defenders need the information too. In medical devices I think it is wrong, and the reason is specific rather than general: the people exposed by early publication are patients who have no ability to act on the information, cannot patch anything, and often do not know the device exists. **Non-disclosure**, reporting privately and never publishing, keeps everyone comfortable and removes the pressure that makes remediation happen, while denying the field the ability to learn. Coordinated disclosure is the middle position and it is the one the field has converged on for good

reasons.

## Who to Contact, and CVE Assignment

### 12.1 The ecosystem

| Body                | Role                                                                                                                                                                        |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Manufacturer PSIRT  | Fixes the problem. First contact where a channel exists.                                                                                                                    |
| CISA                | Coordinates disclosure for critical infrastructure including medical devices; publishes advisories; acts as root CNA for ICS and medical devices and as CNA of last resort. |
| FDA (CDRH)          | Regulates the device. Engages where patient safety is implicated. Not a vulnerability reporting service, but relevant when safety impact is significant.                    |
| Health sector ISAC  | Shares threat information across healthcare; can reach HDOs directly.                                                                                                       |
| MITRE / CVE Program | Maintains CVE and CWE. CVE identifiers are assigned by CNAs.                                                                                                                |
| Your institution    | For academic work, the university may have disclosure policy, legal support, and an interest in being told before you publish.                                              |

#### 12.1.1 Recent development worth knowing

In July 2026, CISA, NSA, and international partners including JPCERT/CC, NCSC-NL, and NCSC-UK published joint guidance on establishing coordinated vulnerability disclosure programmes to work with security researchers. It covers the full chain: receiving reports, communicating with researchers, triage, fix development, coordinated release, and CVE assignment. It also addresses using coordinators as intermediaries where a manufacturer's own programme is immature.

That document is written for manufacturers, which makes it useful to researchers in a particular way: **it tells you what good looks like**. Among the practices it recommends are safe harbour language assuring researchers that good-faith work is authorised, avoiding blanket non-disclosure agreements and silent fixes, and publishing advisories openly rather than behind a paywall. When a manufacturer's process departs from those, you now have an authoritative reference for saying so.

### 12.2 CVE assignment

A CVE identifier gives everyone a common name for the vulnerability. Without one, hospital asset systems, vulnerability scanners, and vendor advisories all describe the same problem differently and nobody can correlate them.

| Route                 | When to use it                                                                                                 |
|-----------------------|----------------------------------------------------------------------------------------------------------------|
| Manufacturer is a CNA | They assign. Simplest path; ask during coordination.                                                           |
| CISA                  | For ICS and medical devices where the manufacturer is not a CNA, or where the manufacturer declines to assign. |
| CNA of last resort    | Where no other CNA has scope.                                                                                  |

#### Where researchers get this wrong

A manufacturer refusing to assign a CVE for a confirmed vulnerability is a meaningful signal, and it happens. Silent fixes deprive hospitals of the ability to know whether their devices are affected. If a manufacturer confirms a vulnerability, fixes it, and declines a CVE, escalating to CISA as root CNA for this space is appropriate and is exactly the mechanism's purpose. There is a documented appeals process for CNAs that are not complying with the rules.

### 12.3 Credit

Ask for it, in writing, early. Most manufacturers credit researchers as a matter of course and the ones that do not usually say so upfront. For a student, a named credit in a vendor advisory is a durable, verifiable, third-party attestation that you did real work, and it is worth substantially more than a line on a résumé that only you can vouch for.

Be equally careful to credit prior work. **P-06** exists because extending someone else's disclosure without saying so is both dishonest and easily caught.

## Legal and Ethical Footing

### I am not a lawyer

This chapter describes the landscape as I understand it. It is not legal advice and I am not qualified to give any. Laws vary by jurisdiction, change over time, and turn on facts specific to your situation. **If you are doing this work, talk to counsel before you start.** If you are at a university, your institution has lawyers and a research compliance office, and they would much rather hear from you beforehand than afterwards.

### 13.1 The three questions to settle before testing

#### 13.1.1 Do you have the right to possess and examine this device?

Buy it, or have it lawfully provided. A decommissioned device acquired through legitimate secondary channels, or one supplied by an institution for research, is a different position from one borrowed from a hospital without a clear agreement.

Get the provenance in writing. It matters later.

#### 13.1.2 Does examining it violate anything you agreed to?

Check the terms attached to the device, its software, and any service agreement. Some contain clauses restricting reverse engineering. Whether such clauses are enforceable against security research varies, and that variation is precisely why this is a question for counsel and not for a researcher's judgement.

#### 13.1.3 Does your method implicate anti-circumvention or computer misuse law?

Two areas commonly relevant in the United States.

**Anti-circumvention.** The Digital Millennium Copyright Act restricts circumventing technological protection measures. There is an exemption for good-faith security research, established and renewed through the Copyright Office's triennial rulemaking, and medical devices have featured in those proceedings. Exemptions have conditions attached, and the conditions are the part that matters. **Check the current exemption text and its conditions; do not rely on a summary, including this one.**

**Computer misuse.** Statutes addressing unauthorised access apply to accessing systems you are not authorised to access. Research on a device you own, on an isolated bench network, is a very different posture from touching a hospital network. This is one of several reasons **P-01** forbids the latter absolutely.

### The point

The strongest legal position is the one where the questions barely arise: a device you lawfully own, on an isolated network you control, with no connection to any clinical system, no patient data, and documented provenance. Almost every legal complication in security research comes from testing something you do not own or touching a network you were not invited onto. Set the work up so those are simply not true.

## 13.2 Safe harbour, and its limits

Some manufacturers publish safe harbour language committing not to pursue legal action against good-faith researchers who follow their policy. The recent CISA-led joint guidance recommends this practice explicitly.

Read any such language carefully:

- **Scope.** Which products and which activities. Testing outside the stated scope is outside the protection.
- **Conditions.** Usually includes reporting promptly, not accessing data you do not need, and not disrupting service.
- **Who it binds.** The manufacturer. It does not bind a prosecutor, and it does not bind third parties.
- **What it is not.** A promise not to sue is not authorisation under every applicable statute.

## 13.3 Ethics beyond legality

Legal and right are not the same, and in this field the gap is wide enough to matter.

**Do not test on devices in service.**

**P-01.** There is no research question important enough to justify the risk, and no framing under which it becomes acceptable.

**Do not access patient data.**

If you encounter it, stop, do not copy it, and report the exposure through the appropriate channel. Handling it “carefully” is not a defence and not a plan.

**Consider the hospital, not only the manufacturer.**

The HDO operating the device is affected by your disclosure and usually has no say in it. Compensating controls they can apply are the part of your report that serves them.

**Do not overstate.**

**P-04.** Inflating severity to get attention works once and costs you the ability to be believed afterwards.

**Think about who reads your publication.**

**P-07.** The audience includes people who would use it. Chapter 19 is entirely about this.

### What else I could have done

**Work only within a formal bug bounty with explicit authorisation.** Cleanest legal position available, and it restricts you to manufacturers who run programmes, which in medical devices is a small subset and skews toward the ones already taking security seriously. The devices most likely to have serious problems are the ones nobody is looking at. **Work only through an institution with existing manufacturer relationships.** Safer and slower, and it is genuinely the right answer for a student starting out; the supervision is worth more than the independence. That is the position I was in, and I would recommend it to anyone in the same position.

## **PART VI**

# **Research Conduct**

## Lab Safety and Device Sourcing

### 14.1 The absolute rule

#### The point

Never test on a device that is in clinical service, and never on a network that touches clinical systems. There is no exception, no supervision level, and no research question that changes this. **P-01** is the one principle in this document with no balancing considerations.

The reason is not liability. It is that a device you have been probing may behave differently afterwards in ways you did not intend and cannot fully verify, and the person who finds out is a patient.

### 14.2 Sourcing a device

| Source                          | Position                      | Cautions                                                                                          |
|---------------------------------|-------------------------------|---------------------------------------------------------------------------------------------------|
| Secondary market                | Common for academic work      | Verify it is genuinely decommissioned; keep purchase records; assume it may contain residual data |
| Manufacturer loan               | Best position where available | Usually comes with an agreement; read the scope and the publication clauses                       |
| Institutional inventory         | Good, if formally documented  | Get written confirmation it is out of service and will not return without revalidation            |
| Hospital biomedical engineering | Possible via partnership      | Requires a formal agreement; never informal borrowing                                             |
| Clinical environment            | Never                         | See above                                                                                         |

#### Where researchers get this wrong

Devices bought on the secondary market frequently arrive with residual data from their prior clinical life: configuration, event logs, and sometimes information relating to patients. **Plan for this before the device arrives.** Decide in advance what you will do if you encounter patient data, which is to stop, not copy it, document that you found it, and report the exposure appropriately. Discovering it with no plan, mid-analysis, is how well-meaning researchers end up handling data they should never have had.

### 14.3 The bench environment

| Requirement                 | Why                                                                                                         |
|-----------------------------|-------------------------------------------------------------------------------------------------------------|
| Physically isolated network | No path to any production or clinical network, and no accidental bridge through a laptop's second interface |
| No patient data, ever       | Test data only, clearly synthetic                                                                           |
| No real medication          | Water or a simulant. There is no analysis that requires actual drugs.                                       |
| Device permanently retired  | Marked, recorded, and never returned to service                                                             |
| Controlled physical access  | The device is now a known-vulnerable unit; treat it accordingly                                             |
| Documented configuration    | So that any result can be tied to a known device state                                                      |

The isolation requirement deserves emphasis because the common failure is accidental rather than deliberate: a research laptop connected to both the bench network and the institutional network becomes a bridge. Use separate machines, or verify isolation deliberately rather than assuming it.

## The Research Record

### 15.1 Why documentation is a security control

**P-11** requires the record to be complete enough to reconstruct any claim. This is not administrative hygiene. Three concrete reasons:

1. **You will be challenged.** A manufacturer may dispute a finding. Your ability to show exactly what you did, on which device revision, in what state, is the difference between a technical discussion and an argument about memory.
2. **Reproduction is the manufacturer's bottleneck.** Chapter 9 covers this. Detailed records are what let them confirm quickly.
3. **Findings change shape.** What looks like one thing early often turns out to be another. Without a record you cannot revisit the evidence when your interpretation shifts.

### 15.2 What to record

| Category        | Content                                                                          |
|-----------------|----------------------------------------------------------------------------------|
| Device identity | Model, hardware revision, software version, serial, and how you established each |
| Provenance      | Where it came from, when, and under what agreement                               |
| Baseline state  | Configuration as received, before you changed anything                           |
| Environment     | Network topology, tools and versions, bench setup                                |
| Session log     | Dated entries: what you did, why, what happened, what you concluded              |
| Evidence        | Captures and observations, each tied to a session entry                          |
| Inferences      | Marked explicitly as inference, with the reasoning                               |
| Dead ends       | What you tried that did not work. Saves you and others repeating it.             |
| Correspondence  | Every exchange with the manufacturer or a coordinator, dated                     |

#### The point

Record dead ends. They are the most commonly discarded and most quietly valuable part of a research record. Six weeks later you will consider an approach that feels novel and is not, and the note saying you tried it and why it failed will save you a week. When a second researcher picks up the work, that section is worth more than the findings.

### 15.3 Handling the sensitive part

The research record contains exactly the material this document withholds: exact reproduction steps, specific values, precise paths. It is the most sensitive artefact you will produce.

- Encrypted at rest, on institutional storage where available.
- Access limited to the research team and supervisors.
- Never in a public repository, never in a personal cloud account, never in an email attachment to somebody you have not verified.
- Retained after publication, because questions arrive years later.
- With a plan for what happens to it when the project ends and when you leave the institution.

#### Where researchers get this wrong

The most common way sensitive research detail leaks is not a breach. It is a screenshot in a slide deck, a photograph in a portfolio, or a repository that was private during the work and got flipped public during a cleanup two years later. **Audit your own published material for detail you did not intend to include**, including images, and set a reminder to re-audit when you change jobs or refresh a portfolio.

## Deciding What to Publish

### 16.1 The chapter this document is really about

Everything else here is method that could be learned from a textbook. This chapter is the judgement, and it is the part I would want a reviewer to read if they read nothing else.

### 16.2 The test

#### The test I apply

**Does this material make a reader meaningfully closer to compromising a device that is currently keeping someone alive?** If yes, it does not go in. If no, it can. Everything else is applying that question honestly to specific sentences.

Two words carry the weight. **Meaningfully**: general security education does not count, because a reader who learns what a hardcoded credential is has gained understanding, not capability. **Currently**: a device that has been remediated, withdrawn, or is no longer deployed sits differently, though even then be careful, since fleets persist far longer than manufacturers expect.

### 16.3 Applying it

| Material                                          | Publish?     | Reasoning                                                                                       |
|---------------------------------------------------|--------------|-------------------------------------------------------------------------------------------------|
| Weakness class and CWE                            | Yes          | Standard security education; confers no capability                                              |
| Why the class occurs in embedded medical products | Yes          | Design insight; helps engineers avoid it                                                        |
| Threat modelling method                           | Yes          | Fully transferable, device-independent                                                          |
| Scoring approach                                  | Yes          | Methodology                                                                                     |
| Disclosure process                                | Yes          | Methodology, and the field needs more of it published                                           |
| That a finding of class X was made on device Y    | Usually      | Hospitals and regulators need to know a problem exists; coordinate timing with the manufacturer |
| The affected component, named precisely           | Case by case | Depends on whether naming it materially narrows an attacker's search                            |
| Interface or protocol path to reach it            | No           | Directly reduces attacker effort                                                                |
| Reproduction steps                                | No           | This is the recipe                                                                              |
| Credential values or equivalents                  | No           | Immediate, universal compromise of the fleet                                                    |
| Firmware extraction procedure                     | No           | Enables independent rediscovery of everything else                                              |

| Material                      | Publish? | Reasoning           |
|-------------------------------|----------|---------------------|
| Tooling built for this device | No       | Packaged capability |

## 16.4 Obscuring is not withholding

**P-08** exists because of a specific and seductive failure.

The temptation is to describe the finding in vague terms that feel responsible while conveying enough that a knowledgeable reader can reconstruct it. It reads as restraint. It functions as disclosure with extra steps, and it is worse than either publishing or withholding, because the author believes they were careful.

### Where researchers get this wrong

Test your own text adversarially. Give the draft to someone with relevant skills who has not seen the work, and ask them what they could do with it. If they can describe the path, you have published the path regardless of how carefully you phrased it. I would rather cut a paragraph I was proud of than discover afterwards that a reader reconstructed the finding from it.

## 16.5 The argument against withholding, taken seriously

There is a real position on the other side and it deserves better than a strawman.

Defenders need detail. A hospital security team assessing exposure benefits from knowing specifically what to look for. Withheld detail also protects manufacturers from scrutiny, and manufacturers have historically used safety framing to suppress research that embarrassed them. And there is an argument that attackers capable of harming patients through a networked medical device are sophisticated enough to find these problems themselves, so withholding mostly disadvantages defenders.

I take that seriously, particularly the third point. My answer is that the asymmetry runs the wrong way here. **The defender's action does not usually require the detail.** A hospital segments the network, restricts access, and monitors; those are the same actions whether or not they know the exact credential. The attacker's action does require it. So publishing the detail helps the attacker more than the defender, which is not true in every domain but is true in this one.

### The point

The reason to withhold is not that the information is dangerous in the abstract. It is that in this specific domain, the defender's realistic response does not depend on the detail and the attacker's does. Where that asymmetry reverses, and in some domains it does, the calculus should reverse with it.

## 16.6 What to publish instead

Withholding leaves a gap. Fill it with the things that are genuinely useful and genuinely safe:

- **The methodology**, in full. Somebody else can apply it to a different device.

- **The design lesson.** Chapter 6's observation about remediation scoped to a reported class rather than to the underlying design decision is more valuable to the field than the specific finding was.
- **Compensating controls,** stated concretely enough for an HDO to act on.
- **The disclosure experience.** What worked, what did not, how long it took. The field is short of this and it helps the next researcher.
- **The reasoning about restraint itself.** Which is what this chapter is.

#### What else I could have done

**Publish nothing.** Zero risk of harm and zero contribution; the field learns nothing and the incentive to fix things weakens. **Publish fully after a fixed embargo,** the ninety-day model. It works well for software and fits medical devices badly, for the lifecycle reasons in Chapter 2: a fix may exist and still not be deployed across a fleet for years, so the embargo expiring tells you very little about whether patients are protected. **Publish detail only to a vetted group,** such as an ISAC. Genuinely attractive, and it is the model I would use for material too sensitive to publish but too useful to withhold entirely. Its weakness is that vetting is imperfect and it excludes independent researchers who have no institutional membership. I chose method-public, detail-private because it maximises what transfers while keeping the part that would cause harm in the one place it needs to be, which is with the people who can fix it.

## PART VII

# Templates

## Threat Model Template

### 17.1 How to use these

The three templates in this part are generic and contain no device-specific content. Copy them, fill them in, delete the guidance text. They encode the method from Parts II and III so you do not have to rederive it each time.

### 17.2 Section 1: device characterisation

| Field                     | Fill in                                                        |
|---------------------------|----------------------------------------------------------------|
| Device                    | Model, hardware revision, software version                     |
| Intended use              | From the manufacturer's own labelling, quoted, not paraphrased |
| Clinical setting          | Where it is used and by whom                                   |
| Therapy-critical function | The thing that must not go wrong. State it in one sentence.    |
| Connectivity              | Every documented interface, with its purpose                   |
| User roles                | Every documented role and what it can do                       |
| Service life              | Expected years in service; typical age of deployed fleet       |
| Update mechanism          | How software changes reach a deployed device                   |
| Prior disclosures         | CVEs and published research, with references                   |
| Manufacturer claims       | Security controls the manufacturer states, with source         |
| Evidence sources          | Every document you used, with dates and versions               |

### 17.3 Section 2: data flow diagram

Draw it. Include external entities, processes, data stores, flows, and trust boundaries. Then produce the element inventory:

| ID    | Type       | Description | Evidence / inferred? |
|-------|------------|-------------|----------------------|
| E-01  | External   |             |                      |
| P-01  | Process    |             |                      |
| D-01  | Data store |             |                      |
| F-01  | Flow       |             |                      |
| TB-01 | Boundary   |             |                      |

Mark every inferred element. Chapter 5 explains why this matters more than it looks like it should.

### 17.4 Section 3: entry point inventory

One row per place data crosses a trust boundary. Fields from Chapter 6:

| ID    | Boundary | Reachability | Actor | AuthN | Reaches B? |
|-------|----------|--------------|-------|-------|------------|
| EP-01 |          |              |       |       |            |
| EP-02 |          |              |       |       |            |

### 17.5 Section 4: STRIDE analysis

For each entry point, work all six categories. Recording “not applicable, because” is as valuable as recording a threat, because it shows you considered it.

| Entry | STRIDE | Threat statement | Existing control / N/A because |
|-------|--------|------------------|--------------------------------|
| EP-01 | S      |                  |                                |
| EP-01 | T      |                  |                                |
| EP-01 | R      |                  |                                |
| EP-01 | I      |                  |                                |
| EP-01 | D      |                  |                                |
| EP-01 | E      |                  |                                |

### 17.6 Section 5: coverage statement

Write, explicitly: what you examined, what you did not, and why. This is what turns a findings list into a defensible assessment.

| Statement             | Fill in                                                                |
|-----------------------|------------------------------------------------------------------------|
| Entry points examined | Of the N enumerated, which and how thoroughly                          |
| Not examined          | Which, and why: no access, out of scope, insufficient time             |
| Depth achieved        | Per entry point: documentation review, observation, or active analysis |
| Confidence            | Where your conclusions are firm and where they are provisional         |

## Disclosure Report Template

### 18.1 Structure

Two documents, and keeping them separate is the whole point.

| Document        | Contents and handling                                                                                                                              |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Main report     | Summary, risk statement, severity, recommendations, compensating controls. Shareable within the manufacturer's organisation and with coordinators. |
| Technical annex | Reproduction detail, exact values, precise paths, tooling. Encrypted, restricted distribution, explicitly marked. <b>Never published.</b>          |

### 18.2 Main report

| Section                 | Contents                                                                                                                      |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Header                  | Report ID, date, researcher, affiliation, contact, and how you want to be credited                                            |
| Handling                | Classification, intended recipients, and your publication intent with an indicative timeline                                  |
| Summary                 | Three to five sentences. Assume this is the only part a senior person reads.                                                  |
| Affected products       | Model, versions, and how you determined the affected range                                                                    |
| Risk statement          | The nine clauses from Chapter 9, as one paragraph                                                                             |
| Weakness classification | CWE reference and why that class                                                                                              |
| Severity                | CVSS vector and score via the MITRE rubric, both v3.1 and v4.0 during the transition, with your rubric decision-point answers |
| Patient safety          | PIPS determination, and an explicit statement of where qualified clinical review is required                                  |
| Fleet scope             | Per-device or model-wide; estimated deployment if you can support it                                                          |
| Reachability            | Which level from Chapter 6's table, and your evidence for that level                                                          |
| Recommendations         | Remediation that engages with why the weakness exists, per Chapter 6                                                          |
| Compensating controls   | What an HDO can do now, before any fix                                                                                        |
| Prior work              | Related disclosures, credited precisely                                                                                       |
| Annex reference         | That a technical annex exists and how to request it                                                                           |

| Section             | Contents                                   |
|---------------------|--------------------------------------------|
| Disclosure timeline | Dates of every contact and response so far |

#### The point

The handling section goes second, immediately after the header, before the finding. The recipient needs to know what they are holding and what you intend to do with it before they read the details. Burying publication intent at the end of a long report is how a manufacturer ends up feeling ambushed by something you did disclose.

### 18.3 First contact message

Short. The report is attached; the message exists to get it to the right person.

| Element               | Content                                                                                 |
|-----------------------|-----------------------------------------------------------------------------------------|
| Subject               | Security vulnerability report, product name, and your reference                         |
| Who you are           | One sentence, including institutional affiliation if you have one                       |
| What you have         | One sentence: weakness class, product, and that it is confirmed                         |
| What you are asking   | Acknowledgement within a stated period, and a named technical contact                   |
| What you are offering | Full technical detail, reproduction support, and a fix-validation offer                 |
| Your intent           | That you intend coordinated disclosure and will not publish detail without coordination |
| Secure channel        | PGP key or an alternative for the annex                                                 |

#### Where researchers get this wrong

Do not put the technical detail in the first email. You do not yet know who reads that inbox or where it gets forwarded. Send the summary and the offer, establish a contact and a secure channel, then send the annex. This costs a day and prevents the annex sitting in a shared mailbox indefinitely.

## Risk Register and Finding Log

### 19.1 Finding log

One row per finding, maintained across the project.

| ID   | Weakness | Entry pt | Reachability  | PIPS   | Status       |
|------|----------|----------|---------------|--------|--------------|
| F-01 | CWE-...  | EP-...   | Consequential | Yes/No | Reported     |
| F-02 |          |          | Theoretical   |        | Not reported |

Keep the ones you decided not to report, with the reason. Chapter 6’s reachability discipline means most observations do not become reports, and recording why protects you from rediscovering and re-dismissing the same thing.

### 19.2 Risk register

For findings that reach the reporting threshold:

| Field                 | Content                                                |
|-----------------------|--------------------------------------------------------|
| ID                    | Matches the finding log                                |
| Risk statement        | Nine clauses, Chapter 9                                |
| Severity              | Vector, score, method, date                            |
| PIPS                  | Determination and reasoning                            |
| Fleet scope           | Per-device or model-wide                               |
| Manufacturer position | Their assessment, where it differs from yours, and why |
| Remediation plan      | What they committed to, with dates                     |
| Milestones            | Each committed date, and whether it was met            |
| Compensating controls | What HDOs can do now                                   |
| Disclosure status     | Where in the Chapter 13 process this sits              |
| CVE                   | Identifier, or the reason there is none                |

#### The point

The milestone row is the one that does real work. Judging a manufacturer’s good faith by elapsed time is unfair, because legitimate timelines are long. Judging it by whether they meet dates they themselves chose is fair, and it gives you an objective basis for deciding whether to escalate. Record every committed date at the moment it is committed, not retrospectively.

## **PART VIII**

# **The Decision Register**

## Every Judgement Call, In One Place

### 20.1 Why collect them

Each chapter has its own alternatives box, which helps while you are reading that chapter and is useless when somebody asks you a question three months later. This collects them so the whole set can be reviewed at once.

The right-hand column is what the individual boxes do not have: **what would change my mind**. A decision without a stated reversal condition is a preference wearing a decision's clothes.

### 20.2 The register

| #    | Chose                               | Over                            | Reason, and what would change my mind                                                                                                                                                                                                                                 |
|------|-------------------------------------|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| D-01 | A previously disclosed device       | An unexamined device            | <i>Reason:</i> the easy findings are gone, there is a public baseline to reason against, and a new finding is a stronger claim about independent capability. <i>Reversal:</i> a research goal about breadth of the device population rather than depth on one device. |
| D-02 | Narrow depth on credential handling | Broad enumeration               | <i>Reason:</i> depth let me establish reachability properly; one finding at “consequential” beats fifteen at “theoretical.” <i>Reversal:</i> no disclosure history to anchor on. Then go broad first, narrow after the first pass.                                    |
| D-03 | STRIDE                              | Attack trees, PASTA, kill chain | <i>Reason:</i> systematic, defensible coverage claims, and recognised by the regulatory audience. <i>Reversal:</i> a device where chaining several small weaknesses is the main concern. Then add attack trees alongside.                                             |
| D-04 | Workflow-aware analysis             | Purely technical assessment     | <i>Reason:</i> findings ranked without clinical workflow context get ranked wrong in both directions. <i>Reversal:</i> none, though it costs real time to learn the workflow.                                                                                         |
| D-05 | Reachability before severity        | Report everything found         | <i>Reason:</i> <b>P-05</b> . Reports of unreachable weaknesses are the main source of manufacturer distrust. <i>Reversal:</i> none.                                                                                                                                   |
| D-06 | CVSS via the MITRE rubric           | Bare CVSS                       | <i>Reason:</i> repeatability, clinical framing, and an FDA-qualified process behind the number. <i>Reversal:</i> none for medical devices.                                                                                                                            |
| D-07 | CVSS at all                         | Narrative severity only         | <i>Reason:</i> manufacturers and regulators run on comparable scores; narratives cannot be queued. <i>Reversal:</i> none, but always pair the score with what it does not capture.                                                                                    |
| D-08 | Three-layer severity report         | Single score                    | <i>Reason:</i> technical, clinical, and fleet context answer different readers' questions. <i>Reversal:</i> none.                                                                                                                                                     |

| #    | Chose                          | Over                          | Reason, and what would change my mind                                                                                                                                                                                                                             |
|------|--------------------------------|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| D-09 | Nine-clause risk statement     | Generic bug report format     | <i>Reason:</i> it answers the manufacturer's triage questions in the order they ask them. <i>Reversal:</i> a manufacturer with a mandated report format. Use theirs.                                                                                              |
| D-10 | Learn the regulatory frame     | Treat it as purely technical  | <i>Reason:</i> it is the only way to judge whether a manufacturer's response is reasonable. <i>Reversal:</i> none.                                                                                                                                                |
| D-11 | Coordinated disclosure         | Full disclosure               | <i>Reason:</i> the people exposed by early publication are patients who cannot act on the information. <i>Reversal:</i> an active, exploited-in-the-wild situation where hospitals need to defend immediately. Then coordinate with CISA on emergency disclosure. |
| D-12 | Coordinated disclosure         | Non-disclosure                | <i>Reason:</i> silence removes the pressure that makes remediation happen and denies the field the ability to learn. <i>Reversal:</i> none.                                                                                                                       |
| D-13 | Judge good faith by milestones | Judge by elapsed time         | <i>Reason:</i> legitimate medical device timelines are genuinely long; met commitments are the honest signal. <i>Reversal:</i> none.                                                                                                                              |
| D-14 | Method public, detail private  | Full publication              | <i>Reason:</i> in this domain the defender's realistic response does not need the detail and the attacker's does. <i>Reversal:</i> a domain where that asymmetry inverts, which does exist elsewhere.                                                             |
| D-15 | Method public, detail private  | Publish nothing               | <i>Reason:</i> the methodology and the design lesson transfer, and the field is short of published disclosure experience. <i>Reversal:</i> none.                                                                                                                  |
| D-16 | Method public, detail private  | Vetted-group distribution     | <i>Reason:</i> vetting is imperfect and excludes independent researchers with no institutional membership. <i>Reversal:</i> material too useful to withhold and too sensitive to publish. Then an ISAC is the right venue.                                        |
| D-17 | Withhold rather than obscure   | Vague description             | <i>Reason:</i> <b>P-08</b> . Obscured detail a reader can reconstruct is disclosure that fooled its own author. <i>Reversal:</i> none.                                                                                                                            |
| D-18 | Institutional supervision      | Independent research          | <i>Reason:</i> for a student, supervision is worth more than independence, legally and methodologically. <i>Reversal:</i> experience, an established manufacturer relationship, and access to counsel.                                                            |
| D-19 | Bench device, isolated network | Any device in service         | <i>Reason:</i> <b>P-01</b> . No balancing considerations. <i>Reversal:</i> none, ever.                                                                                                                                                                            |
| D-20 | State the competence boundary  | Assess clinical impact myself | <i>Reason:</i> I am not clinically trained, and guessing produces confident nonsense a clinical reviewer spots instantly. <i>Reversal:</i> a clinically qualified collaborator on the team. Then they write that section.                                         |

### 20.3 The three I would revisit first

1. **D-20, the clinical gap.** The largest genuine weakness in the work. A finding without a clinical impact assessment is incomplete, and the fix is not more study on my part; it is a clinical collaborator. If I ran this again, recruiting one would be the first thing I did, before touching the device.
2. **D-02, narrow versus broad.** I still think narrow was right given the prior disclosure. But I chose it partly for the reason it happened to be defensible rather than after weighing it properly, and I want to be honest that the reasoning was cleaner in retrospect than at the time.
3. **D-14, the publication line.** I am confident in where I drew it. I am less confident that I would draw it in the same place on a device with no fix available and a hospital population actively exposed, and I have not had to make that call yet.

#### The point

Twenty decisions, and the honest summary is that one of them, D-20, is a real limitation rather than a trade-off I would defend. Being able to say which of your choices were judgment and which were constraint, out loud and unprompted, is most of what distinguishes a research writeup from a marketing document.

## Standards and Regulatory Reference

### Check currency before citing

Everything in this appendix was verified against current sources in August 2026. Guidance documents are superseded regularly and standards are revised. **Confirm the current version before citing any of it in a document that goes to a manufacturer or a regulator.**

### A.1 United States regulatory

| Instrument                            | What it does                                                                                                                                                                                                                                                                                              |
|---------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FD&C Act §524B                        | Statutory cybersecurity requirements for “cyber devices” in premarket submissions: postmarket vulnerability management plan including coordinated disclosure, secure design processes, patch availability, and a software bill of materials. Effective March 2023. Noncompliance is a prohibited act.     |
| FDA premarket cybersecurity guidance  | Agency expectations for device design, labelling, and submission documentation. Retitled in the early-2026 revision to reference quality management system considerations, aligning with the QMSR. Supersedes the June 2025 version, which superseded the September 2023 version. Docket FDA-2021-D-1158. |
| QMSR (21 CFR Part 820)                | Quality system regulation, effective February 2026, incorporating ISO 13485:2016 by reference.                                                                                                                                                                                                            |
| FDA postmarket cybersecurity guidance | Expectations for monitoring, assessing, and remediating vulnerabilities in marketed devices.                                                                                                                                                                                                              |
| MITRE CVSS rubric for medical devices | Structured questions and decision trees for applying CVSS to medical devices. Qualified by FDA as a Medical Device Development Tool.                                                                                                                                                                      |

### A.2 Consensus standards

| Standard  | Scope                                                                                                              |
|-----------|--------------------------------------------------------------------------------------------------------------------|
| ISO 14971 | Risk management for medical devices. The safety risk framework everything else attaches to.                        |
| IEC 62304 | Medical device software lifecycle processes. Governs change control, which is why patches take as long as they do. |

| Standard       | Scope                                                                                                                                                                |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ISO 13485      | Quality management systems for medical devices.                                                                                                                      |
| AAMI TIR57     | Principles for medical device security risk management; bridges security risk to ISO 14971 safety risk.                                                              |
| ANSI/AAMI SW96 | Security risk management for manufacturers, building on TIR57.                                                                                                       |
| IEC 81001-5-1  | Security in the health software and health IT systems lifecycle.                                                                                                     |
| IEC 60601-1    | Basic safety and essential performance of medical electrical equipment. Defines “essential performance,” which is the concept your impact analysis should attach to. |

### A.3 Disclosure and coordination

| Resource                       | Role                                                                                                                                                                                                                                                                                  |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CISA CVD programme             | Coordinates disclosure for critical infrastructure including medical devices; publishes advisories.                                                                                                                                                                                   |
| CISA as root CNA               | Root CVE Numbering Authority for industrial control systems and medical devices, with a documented appeals process where a CNA is not complying with the rules.                                                                                                                       |
| Joint CVD guidance (July 2026) | CISA, NSA, JPCERT/CC, NCSC-NL, and NCSC-UK guidance on establishing coordinated vulnerability disclosure programmes. Written for manufacturers; useful to researchers as a statement of what good practice looks like, including safe harbour language and avoidance of silent fixes. |
| ISO/IEC 29147                  | Vulnerability disclosure.                                                                                                                                                                                                                                                             |
| ISO/IEC 30111                  | Vulnerability handling processes.                                                                                                                                                                                                                                                     |
| Health sector ISAC             | Threat information sharing across healthcare; can reach HDOs.                                                                                                                                                                                                                         |
| FIRST CVSS                     | Specification and calculators. Note the v3.1 to v4.0 transition.                                                                                                                                                                                                                      |
| MITRE CWE                      | Weakness classification.                                                                                                                                                                                                                                                              |

## Glossary

**Adverse event**

An undesirable clinical outcome associated with device use. Triggers investigation and potentially regulatory reporting.

**CNA**

CVE Numbering Authority. An organisation authorised to assign CVE identifiers within a defined scope.

**Compensating control**

A measure that reduces risk without fixing the underlying weakness. Frequently the only thing available to a hospital.

**CVD**

Coordinated vulnerability disclosure. Chapter 13.

**CVE**

A specific, publicly disclosed vulnerability in a specific product.

**CVSS**

Common Vulnerability Scoring System. Chapters 7 and 8.

**CWE**

Common Weakness Enumeration. A *class* of mistake, not a specific instance.

**Drug library**

A hospital-specific set of medications and dose limits loaded onto an infusion pump, used to catch programming errors.

**Essential performance**

From IEC 60601-1: the performance necessary to achieve freedom from unacceptable risk. The concept your impact analysis should attach to.

**HDO**

Healthcare delivery organisation. The hospital or clinic operating the device.

**Large-volume pump**

An infusion pump for delivering fluid volumes over extended periods, typically pole-mounted at the bedside.

**MDM**

Medical device manufacturer.

**MDS2**

Manufacturer Disclosure Statement for Medical Device Security. A standardised form stating a device's security characteristics. Useful reconnaissance for threat modelling.

**PIPS**

Potential impact to patient safety. Chapter 8.

**PSIRT**

Product security incident response team. The manufacturer function that receives your report.

**Reachability**

Whether a path exists from an entry point to a weakness. Chapter 6.

**SBOM**

Software bill of materials. A component inventory, required for cyber devices under §524B.

**SPDF**

Secure product development framework. Processes reducing vulnerability number and severity across the lifecycle.

**STRIDE**

Threat classification: spoofing, tampering, repudiation, information disclosure, denial of service, elevation of privilege. Chapter 4.

**Trust boundary**

A point where data crosses between differing levels of trust. Where the findings live. Chapter 5.

**VDP**

Vulnerability disclosure policy. A published statement of how an organisation wants to receive reports.

## Principle Traceability

Every principle from Chapter 3, where it is applied, and how you would tell whether it was followed.

| ID   | Principle                                 | Applied in | How to verify                                            |
|------|-------------------------------------------|------------|----------------------------------------------------------|
| P-01 | No risk to patients                       | Ch. 17     | Device provenance records; network isolation evidence    |
| P-02 | Safety is the primary severity axis       | Ch. 7–9    | PIPS determination present in every finding              |
| P-03 | Disclosure timing serves safety           | Ch. 14, 19 | Publication decision documented with reasoning           |
| P-04 | Severity stated accurately                | Ch. 7–9    | Scoring method recorded with decision-point answers      |
| P-05 | Reachability before “vulnerability”       | Ch. 6      | Reachability level stated with evidence in every report  |
| P-06 | Prior work credited precisely             | Ch. 1, 20  | Prior disclosures section present and specific           |
| P-07 | No material attacker assistance           | Ch. 19     | Adversarial read by a third party                        |
| P-08 | Withhold, do not obscure                  | Ch. 19     | Same adversarial read; can a skilled reader reconstruct? |
| P-09 | Existence may be disclosed without detail | Ch. 14, 19 | Advisory issued even where detail withheld               |
| P-10 | Manufacturer contacted first              | Ch. 13–14  | Dated disclosure timeline                                |
| P-11 | Complete research record                  | Ch. 18     | Any claim traceable to a session entry                   |
| P-12 | Legal footing before testing              | Ch. 16     | Counsel consulted; provenance documented                 |

## Open Items

### D.1 What I have not confirmed

Short, and here because writing down what I do not know is cheaper than having a reviewer find it.

#### D-1. Clinical impact assessment. (Most significant.)

The finding has no clinical impact assessment by a qualified clinician. I stated the technical finding and the therapy-relevant parameter and flagged that clinical review was required, which is the correct handling of a boundary I cannot cross. It remains a gap in the work rather than a resolved question. **Action: on any future device work, recruit a clinical collaborator before beginning.**

#### D-2. Current regulatory citations.

Chapter 10 and Appendix A were checked in August 2026, after the research was conducted. Several documents changed in between, notably the premarket guidance revision and the QMSR taking effect. Anything cited to a manufacturer should be re-verified at the time of citation.

#### D-3. CVSS version.

The research scored with v3.1, correct at the time. FDA has since recognised v4.0 with a transition period. **Action: rescore with v4.0 and provide both vectors if this finding is discussed further.**

#### D-4. Fleet scope.

I have not established the deployed population of the affected model. A fleet-scope claim in a report should be supportable, and mine currently is not beyond “widely deployed,” which is weaker than I would like.

#### D-5. Disclosure outcome.

This document describes the disclosure method. It does not report a completed disclosure timeline with manufacturer milestones, because that record belongs in the coordination file rather than in a public methodology document, and because the process detail is not mine alone to publish.

### D.2 What I would add with more time

1. A clinical collaborator, and a proper hazard analysis attached to the finding (D-1, D-20 in the register).
2. Rescoring under CVSS v4.0, using the supplemental Safety metric that v4.0 introduces.
3. A compensating-controls annex written specifically for hospital biomedical engineering teams, in their language rather than a researcher’s.
4. A structured comparison of what the prior remediation covered against what it did not, which is the generalisable research contribution and deserves a paper of its own.
5. Applying this same methodology to a second device, to test whether the method transfers as cleanly as I have claimed it does.

End of document. Version 1.0, August 2026.

No device-specific exploitation detail is contained in these pages.

Technical annex held under coordinated disclosure.