Detection Is Not a Device Property: A Sensing-Envelope and Responder-Conditioned Framework for Risk-Control Verification in Life-Sustaining Medical Devices

Jherrod Thomas

Independent Researcher — The Lion of Functional Safety
jherrodthomas.com

***Abstract*—Two United States Food and Drug Administration field actions published within thirteen days of each other in July 2026 falsify two distinct and rarely stated assumptions embedded in ISO 14971 risk files for life-sustaining devices. In the first, the FDA classified as Class I an Insulet field action covering approximately seven million Omnipod insulin pods whose cannula may tear just above the skin, venting insulin outside the body while every internal signal reports normal delivery; the agency's own communication instructs users not to rely on pod alerts, and the substitute detection channels it offers are the smell of insulin and wetness on the adhesive. In the second, the FDA issued an Early Alert covering ResMed Astral 100 and 150 life-support ventilators in which a leaking supercapacitor can corrode the main printed-circuit-board assembly and drive the device into a fail-safe state that stops therapy — a state that is safe only if a competent responder is present within the patient's apnea tolerance. We argue these are the same defect class expressed twice: a risk control was credited without bounding the two external quantities on which its efficacy actually depends. We propose the Detection–Responder Contract, a six-step framework that (i) classifies every credited control by observation channel against a five-tier taxonomy with per-tier crediting ceilings, (ii) factorizes detectability into physical manifestation, transduction, and decision-threshold terms, (iii) classifies each safe state as benign or life-withholding, and (iv) conditions residual harm on responder availability within a physiologic tolerance time. Three numbered relations make the previously implicit assumptions computable. A worked example on both 2026 cases yields eight derived requirements and shows that under a one-sided detection architecture the creditable detection for an open-path fault is bounded near zero irrespective of alarm-system conformance.**

***Index Terms*—ISO 14971, IEC 60601-1-8, ISO 80601-2-72, IEC 60601-2-24, Risk Control, Detectability, Safe State, Insulin Pump, Home Ventilation, Alarm System**

A. Introduction

A risk file is a ledger of credits. ISO 14971 requires the manufacturer to identify hazards, estimate and evaluate the associated risks, implement risk control measures, and then evaluate the residual risk that remains after those measures are in place [3]. The arithmetic of that last step is where the ledger is settled: a hazardous situation is assigned some probability of occurrence, a risk control is credited with reducing it, and the reduced figure is compared against the manufacturer's acceptability criteria. Two of the most common credits in the ledger of a life-sustaining device are *detection* — the device will notice the fault and annunciate it — and *safe state* — the device will withdraw to a defined condition in which it cannot do harm.

Both credits contain an unstated dependency on the world outside the device. The detection credit assumes the fault produces a signal the device's transducers can actually see. The safe-state credit assumes that ceasing to act is harmless, which for a device that is sustaining a physiologic function it is not: a stopped ventilator on a ventilator-dependent patient is a hazard, not a mitigation, unless a responder arrives in time. Neither dependency is a property of the device. Both are properties of the device *in a context*, and neither is forced onto the page by any artifact in the standard risk-management workflow.

Two July 2026 FDA actions make the omission concrete. On July 2, 2026, the FDA classified as Class I an Insulet Medical Device Correction covering Omnipod 5, Omnipod DASH, and Omnipod Eros pods in which a small tear can develop in the cannula just above the skin, between the pod and the point of entry into the body, allowing insulin to leak outside the patient. The scope is approximately seven million pods, with twenty-four reports of serious adverse events involving hospitalization and diabetic ketoacidosis (DKA) and no deaths; a separate and earlier Insulet correction for internal tubing tears in the same platform, classified Class I on April 29, 2026, had accumulated 476 serious injuries [1]. The FDA's user communication is unusually direct about what the device cannot do: the issue "may happen without triggering any alert," and users are told, in bold, not to rely only on pod alerts to know if something is wrong. The alternative cues the agency offers are wetness on the skin or adhesive and the smell of insulin [1].

On July 15, 2026, the FDA published an Early Alert covering ResMed Astral 100 and 150 portable life-support ventilators built prior to October 2024, plus spare main boards distributed as service parts. A supercapacitor may leak over time; the leakage may damage circuitry on the main printed-circuit-board assembly (PCBA); and the damage "can result

in the ventilator inadvertently entering a fail-safe state." If this occurs during therapy, therapy stops, a maximum-volume high-priority alarm sounds, and alternative means of ventilation must be provided. Five serious injuries and no deaths were reported as of June 23, 2026, and the correction is rationed: PCBA availability is "significantly constrained," so ResMed is triaging patients into clinical risk tiers and correcting the highest-risk users first [2].

The conventional reading treats these as unrelated — one a manufacturing escape in a disposable fluid path, the other a component wear-out on a durable board. We read them as one defect class. In the first, a detection credit was taken on a fault that lies outside the device's sensing envelope: the pod's fault architecture is built around occlusion, which raises back-pressure, and a tear is the mirror image of an occlusion — the path is more open, not less, so the pressure signature moves in the direction the detector reads as healthy. In the second, a safe-state credit was taken on a state that withholds a life-sustaining function, so its harmlessness is a claim about the caregiver, not about the ventilator. Neither risk file was required to state the assumption it depended on, and therefore neither was required to verify it.

This paper contributes:

- A five-tier **observation-channel taxonomy** (Table I) that classifies every credited detection by the physical channel that carries the fault signature, and assigns each tier a ceiling on the risk-control credit that may lawfully be taken, together with the verification evidence required to reach that ceiling.
- A **detectability factorization** (1) that decomposes the probability of detecting a failure mode into physical-manifestation, transduction, and decision-threshold terms, exposing which term is the binding constraint and hence which design change can move it.
- A **safe-state classification** distinguishing benign safe states from *life-withholding* safe states, and a **responder-conditioned residual harm** relation (2) that makes the external responder an explicit, quantified parameter of the risk file rather than an implicit assumption, together with an inversion (3) that yields the responder availability a stated residual-risk target requires.
- A **six-step audit method**, the Detection–Responder Contract, producing a matrix suitable for design review, notified-body assessment, and post-market corrective-action review under 21 CFR Part 820.
- A **worked example** on both 2026 field actions, yielding eight derived requirements and a bound showing that a one-sided detection architecture caps creditable detection of an open-path fault near zero regardless of alarm-system conformance to IEC 60601-1-8 [4].

Section II reviews background and related work. Section III presents the framework. Section IV applies it. Section V discusses limitations and threats to validity. Section VI concludes.

B. II. Background and Related Work

1) A. Risk Management Standards and the Detectability Question: ISO 14971:2019 requires that risk control options be selected in a defined order of priority — inherently safe design, protective measures including alarms, and information for safety — and that the effectiveness of each implemented control be verified [3]. It does not, however, prescribe how the *effectiveness* of a detection-based protective measure is to be bounded when the fault of interest may not manifest in any quantity the device measures. The standard's companion guidance treats detectability as an input to risk estimation rather than as a claim requiring its own evidence.

The critique is not new in principle. Nolan and McDermott survey the use and limitations of failure mode and effects analysis in medical-device risk management and identify the detection rating as among the most problematic elements: it is assigned subjectively, it is treated as commensurate with severity and occurrence in the risk-priority-number product, and the method is structurally limited to single-fault conditions and silent on normal-condition hazards [24]. Their critique addresses the *rating* of detectability. Our concern is one level earlier: whether the fault is observable at all through the channels the device possesses, which is a question about physics and instrumentation rather than about scoring discipline.

Two device-specific standards constrain the annunciation side. IEC 60601-1-8 governs alarm systems, assigning alarm priorities and requiring that alarm conditions be annunciated with prescribed auditory and visual characteristics [4]. IEC 60601-2-24 sets particular requirements for infusion pumps, including occlusion alarm behavior [6]. ISO 80601-2-72 sets particular requirements for home-healthcare ventilators for ventilator-dependent patients [5]. Each standard is exacting about what happens *after* an alarm condition is declared. None of them constrains the antecedent: whether the alarm condition can be declared at all for a given physical fault. A device may be fully conformant to IEC 60601-1-8 and still be blind to the failure mode that harms its user, because conformance is evaluated against the alarm conditions the manufacturer has defined.

2) B. Fault Detection in Insulin Delivery: The limits of pressure-based occlusion detection in insulin pumps are well documented. Freckmann et al. measured occlusion detection time at two basal rates and found detection latency strongly dependent on delivery rate, with low basal rates producing long delays before an alarm threshold is crossed [12]. This is a Tier-2 limitation in the taxonomy of Section III: the fault is physically manifest and transducible, but the decision threshold is slow relative to the harm. The literature has responded with richer instrumentation and better estimators. Blanco et al. developed a real-time force-based algorithm for infusion-failure detection, reporting a bagging classifier with 96% accuracy, 94% sensitivity, and 98% specificity against traditional occlusion alarm pressure thresholds in preclinical work [11].

A second thread moves detection off the device entirely and

onto the physiologic feedback loop. Howsmon et al. demonstrated real-time detection of infusion-site failures inside a closed-loop artificial pancreas, achieving 88.0% sensitivity at 0.22 false positives per day under zone model predictive control, against 73.3% sensitivity in a sensor-augmented pump arm [13]. Meneghetti et al. developed unsupervised methods to detect pump malfunction from continuous-glucose-monitor and insulin-delivery traces [14], and later machine-learning anomaly-detection algorithms to alert sensor-augmented-pump users to infusion-site failures [15]. Turksoy et al. addressed the reciprocal problem — detecting faults in the glucose sensor itself — with a real-time model-based method [16].

This body of work is precisely the evidence base for a Tier-3 credit in our taxonomy: detection derived from a physiologic surrogate rather than from the device's own fluid path. It is also, critically, evidence that such a credit is bounded well below unity and is conditional on the closed loop being active. Klonoff et al. characterized the time course from an insulin-pump occlusion to moderate and severe hyperglycemia and ketonemia, establishing that the physiologic tolerance window is measured in hours rather than minutes but is not indefinite [17]. Registry-scale work by Kalscheuer et al. across 46,966 adults with type 1 diabetes quantifies DKA event rates and risk factors, giving a population baseline against which a delivery-interruption hazard must be evaluated [18]. Orbell et al. analyzed severe insulin-pump-related adverse events and their potential root causes [19], and Gao et al. performed a hazard analysis of Class I recalls of infusion pumps, establishing that detection and alarm deficiencies recur across the device class rather than being idiosyncratic to one manufacturer [20].

3) C. Safe States That Are Not Safe: Outside medicine, the recognition that a fail-safe state can itself be hazardous drove the fail-operational architectures now standard in automated driving and brake-by-wire. Sinha analyzed a fail-operational brake-by-wire architecture from an ISO 26262 perspective, motivated explicitly by the fact that de-energizing a by-wire brake is not a safe state [22]. Fu et al. formally verified a fail-operational safety concept for automated driving [23]. The common premise is that when the function being performed *is* the safety function, ceasing to perform it cannot be the mitigation.

Medicine has the same premise and less of the architecture. Life-support ventilation is the paradigm case: the device is performing a physiologic function the patient cannot perform, and the fail-safe state returns that burden to a patient who by definition cannot carry it. The empirical literature on home ventilator failure is thin but consistent. Srinivasan et al. documented the frequency, causes, and outcomes of home ventilator failure [27]; Chatwin et al. analyzed home support and ventilator malfunction across 1,211 ventilator-dependent patients, characterizing both the malfunction rate and the support infrastructure that must absorb it [28]. Both establish that the responder — family caregiver, home-care provider, or clinician — is the actual mitigation, and that the responder's availability is variable, unverified by the manufacturer, and in the home setting frequently a single fatigued person.

The alarm literature completes the picture, because the responder is reached through an alarm. Cho et al. documented alarm fatigue and the perceived obstacles to alarm management among intensive-care nurses [25] — in a setting with far denser staffing than the home. An alarm system that is conformant, loud, and correctly prioritized still yields a responder-availability term strictly less than one, and the standards do not require the manufacturer to estimate it.

4) D. Wear-Out, Diagnostic Coverage, and the Missing Prognostic: The Astral failure mechanism is a wear-out, and wear-out is the failure class the reliability literature is best equipped to bound. Wang and Blaabjerg surveyed capacitor reliability for DC-link applications, including electrolytic wear-out mechanisms and their operating-condition dependence [29]. Rigamonti et al. developed particle-filter-based prognostics for an electrolytic capacitor operating under variable conditions, demonstrating that remaining-useful-life estimation is tractable in situ [30]. Zhao et al. surveyed condition-monitoring techniques for capacitors in DC-link applications [21]. Ayadi et al. measured the impact of thermal cycling on supercapacitor performance during calendar ageing, showing that ageing rate is a strong function of thermal history — which for a portable ventilator carried between hospital, vehicle, and home is exactly the uncontrolled variable [26].

The point is not that ResMed should have foreseen this specific leak. It is that the analytical tools to convert a wear-out mechanism into a bounded exposure interval — the medical-device analogue of a proof-test interval — are mature and were not applied. The process-safety community formalized this decades ago: Jin and Rausand analyzed the reliability of safety-instrumented systems subject to partial testing and common-cause failures, making the test interval an explicit design parameter [31]. Alemzadeh et al. analyzed safety-critical computer failures in medical devices from regulatory databases and found recurring classes of failure that were, in retrospect, analytically foreseeable [8]. Avizienis et al. supply the vocabulary that unifies all of it: a fault becomes an error becomes a failure only through activation and propagation, and detection is a property of the propagation path, not of the fault [10].

5) E. What Is Missing: No published framework, to our knowledge, requires a medical-device risk file to state (i) which physical channel carries a given fault's signature and whether the device instruments that channel, or (ii) whether the safe state credited for a hazard withholds a life-sustaining function and, if so, what responder availability the residual-risk claim assumes. Section III supplies both.

C. III. Approach: The Detection–Responder Contract

1) A. Observation Channels and Crediting Ceilings: Let f be a failure mode of a life-sustaining device and h the hazardous situation it produces. A credited detection d for f is delivered over an *observation channel*: the physical pathway along which the fault's signature travels to something that can declare an alarm condition. We define five tiers.

C0 — Direct measurement of the regulated quantity. The device measures the quantity it is controlling at the point of delivery: for an infusion device, the mass of drug actually entering the patient; for a ventilator, the volume actually entering the lung. C0 is the only tier for which a near-unity detection credit is defensible, because the fault cannot occur without moving the measured quantity.

C1 — Surrogate internal measurement. The device measures a proxy inside its own boundary: line pressure, motor force, plunger position, valve current. C1 credits are directional: a surrogate detects only the faults that move it in the direction the detector interrogates. This is the tier at which the Omnipod cannula tear escapes.

C2 — Internal state and self-test. The device tests its own hardware and software: memory tests, watchdogs, board-level integrity checks. C2 detects faults of the device, not faults of the therapy delivery.

C3 — Derived physiologic surrogate. Detection is inferred from the patient's response, as in continuous-glucose-monitor-based infusion-failure detection [13]–[15]. C3 credits are bounded by published algorithm sensitivity, are conditional on the surrogate sensor being present and trusted, and carry a latency set by physiology.

C4 — Unaided human sensory channel. Detection depends on the user smelling, seeing, or feeling the fault. C4 has no requirement, no verification, and no availability guarantee during sleep, and under ISO 14971's priority ordering it is information for safety, the weakest control class [3].

Table I gives the tiers with their verification evidence and crediting ceilings. The ceilings are normative proposals, not measured constants; their purpose is to force the manufacturer either to supply tier-appropriate evidence or to accept the lower credit.

2) B. Detectability Factorization: For a failure mode f detected over channel c, the probability that the device declares the correct alarm condition before the harm threshold is reached factorizes as

$$D(f, c) = P_{man}(f, c) \cdot P_{trans}(c) \cdot P_{dec}(f, c, T_{tol}) \qquad (1)$$

where P_{man} is the probability that f produces a signature on channel c at all — a question of physics; P_{trans} is the probability that the transducer and signal chain reproduce that signature with sufficient fidelity — a question of instrumentation; and P_{dec} is the probability that the decision logic crosses its threshold within the physiologic tolerance time T_{tol} — a question of algorithm and threshold placement. The factorization is deliberately multiplicative and deliberately ordered, because it identifies the binding constraint and hence the only design change that can help. If $P_{\text{man}} \approx 0$, no amount of algorithm work, alarm-priority escalation, or IEC 60601-1-8 conformance moves D: the fault is outside the sensing envelope, and the remedy is a new channel, not a better threshold. If P_{man} is high but P_{dec} is low, the remedy is threshold and estimator work of the kind Blanco et al. and Freckmann et al. describe [11], [12].

Fig. 1 renders the envelope: the set of failure modes that physically manifest on some channel, the transducible subset the device instruments, and the subset whose signature crosses a declared threshold in time. Faults outside the innermost region are undetectable *by construction*, and a risk file that credits detection for them is not conservative but incorrect.

3) C. Safe-State Classification and Responder Conditioning: Let σ be the safe state a device enters on fault detection. We classify σ as:

Benign. Entering σ removes the hazard and imposes no new one; the patient is no worse off than before therapy began. An infusion pump that stops delivering a non-life-sustaining drug is typically benign.

Life-withholding. Entering σ suspends a function the patient depends on for physiologic viability. A life-support ventilator entering a fail-safe state is life-withholding, and so is a closed-loop insulin system suspending basal delivery in a patient with no endogenous production over a long enough interval.

For life-withholding safe states, the residual harm is not a property of the device. Let P_{e} be the probability of entering σ over the exposure interval of interest, T_{tol} the physiologic tolerance time for the withheld function, and $A_{\text{r}}(T_{\text{tol}})$ the probability that a competent responder perceives the alarm, obtains the alternative means, and restores the function within T_{tol}. Then the residual probability of harm is

$$P_H = P_e \cdot [1 - A_r(T_{tol})] \qquad (2)$$

and inverting against an acceptability target $P_{\text{H,tgt}}$ yields the responder availability the risk file is implicitly claiming:

$$A_{r,req} \geq 1 - P_{H,tgt} / P_e \qquad (3)$$

Relation (3) is the working instrument of the framework. It converts an unstated assumption into a number the manufacturer must either substantiate with human-factors and infrastructure evidence or disclaim. If (3) demands a responder availability above what the alarm and home-support literature supports [25], [27], [28], the risk file has two honest options: reduce P_{e} by design — bound the wear-out, add redundancy, make the state fail-operational rather than fail-safe — or restrict the intended use to settings where A_{r} can be assured. The composite human-error literature, including integrated THERP and HEART treatments [32], supplies the estimation machinery for A_{r}; the point is that it must be applied rather than assumed.

T_{tol} is the parameter that separates the two cases in Section IV. For a ventilator-dependent patient, T_{tol} is on the order of minutes and may be shorter than the time to retrieve and connect a resuscitation bag from another room. For interrupted subcutaneous insulin delivery, T_{tol} is on the order of hours before moderate ketonemia [17] — which is why an insulin-delivery interruption can lawfully be treated as detection-deferrable while a ventilation interruption cannot.

4) D. The Six-Step Method: **Step 1 — Enumerate credited controls.** For every hazardous situation in the ISO 14971 risk file, list each implemented risk control and the credit taken for it in the residual-risk evaluation [3].

Step 2 — Assign observation channels. For each detection-based control, assign the tier from Table I. Where a control is claimed to work over multiple channels, decompose

TABLE I

OBSERVATION-CHANNEL TAXONOMY, REQUIRED VERIFICATION EVIDENCE, AND PROPOSED CEILINGS ON THE DETECTION CREDIT THAT MAY BE TAKEN IN AN ISO 14971 RESIDUAL-RISK EVALUATION.

C0	Direct measurement of regulated quantity at delivery point	Bench characterization across the full fault space, including open-path and closed-path faults; accuracy traceable to the therapy specification	≤ 0.99
C1	Surrogate internal measurement (pressure, force, position)	Two-sided fault injection: faults that raise *and* lower the surrogate; detection-time distribution at the lowest clinical delivery rate	≤ 0.90, and 0 for faults moving the surrogate toward the healthy direction
C2	Device self-test and internal integrity monitoring	Fault injection at the monitored boundary; coverage claim scoped to device faults only	≤ 0.95 for device faults; 0 for therapy-path faults
C3	Derived physiologic surrogate (e.g., CGM-based inference)	Clinical or in-silico sensitivity and false-positive rate at a stated operating point; availability of the surrogate sensor; latency vs. tolerance time	$\leq$ published sensitivity, and 0 when the loop is inactive
C4	Unaided human sensory detection	Human-factors validation under realistic use, including night-time and unattended use	0 for residual-risk arithmetic; usable only as information for safety

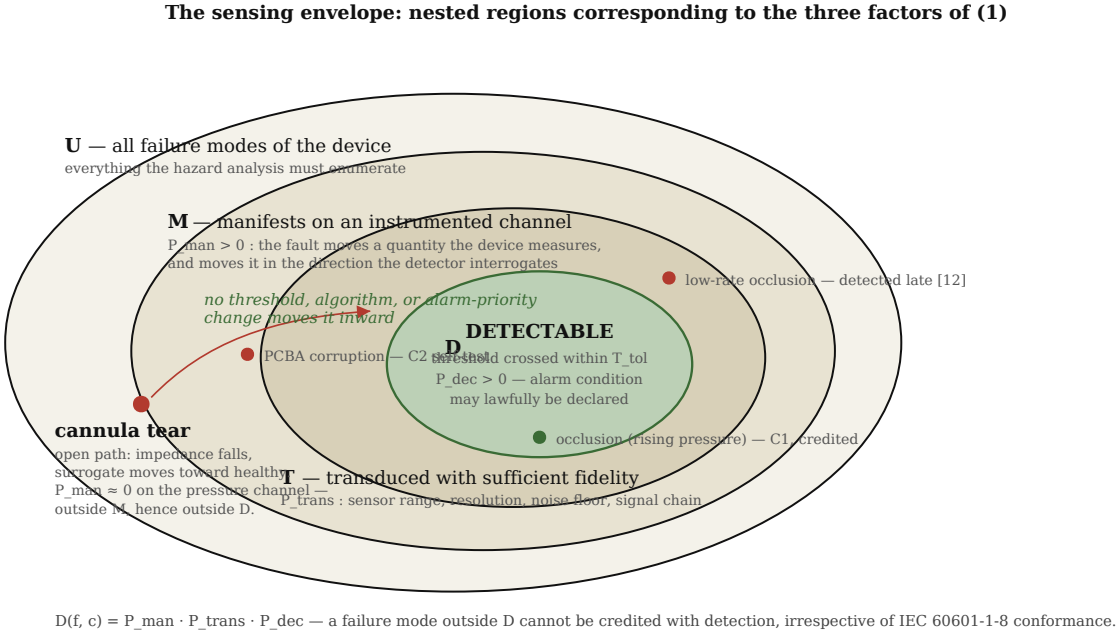

Fig. 1. The sensing envelope. Each nested region corresponds to one factor of (1). A failure mode outside the innermost region cannot be credited with detection regardless of alarm-system conformance. The Omnipod cannula tear (Section IV-A) sits outside the manifestation region for the pressure channel because it moves the surrogate toward the healthy direction.

it: each channel is credited separately and the channels are combined only if their failure modes are demonstrably independent.

Step 3 — Evaluate (1) per failure mode. For each failure mode the control is credited against, evaluate the three factors of (1), and in particular test P_{man} for *directionality*: does the fault move the surrogate toward or away from the healthy reading? Two-sided fault injection is the required evidence, and its absence is itself a finding.

Step 4 — Classify safe states. Mark each safe state benign or life-withholding. Any life-withholding safe state entered as a *reaction to a device fault* — as distinct from a clinically commanded stop — is escalated to a hazardous situation in its own right and returned to Step 1.

Step 5 — Condition on the responder. For every life-withholding safe state, evaluate (2) and invert (3). Record T_{tol}, the assumed A_{r}, and the evidence for it. An unsubstantiated A_{r} is treated as a gap, not as unity.

Step 6 — Derive requirements and verification. Where a ceiling in Table I, a factor in (1), or the inversion in (3) is unmet, write a derived requirement with an explicit verification method, and route it into design controls and, post-market, into corrective and preventive action.

D. IV. Worked Example

We apply the method to both July 2026 field actions. All inputs are drawn from public regulatory record and open literature; numerical values marked *illustrative* are used to show the arithmetic and are not attributed to the manufacturers.

1) A. Case A — Omnipod Cannula Tear (FDA Class I, July 2, 2026): **Steps 1–2.** The hazardous situation is under-delivery of insulin leading to hyperglycemia and DKA. The credited controls, as inferable from the platform's documented behavior and the FDA communication, are: (a) an occlusion alarm on the fluid path — Tier C1; (b) automated-mode glucose-trend logic that can raise an Automated Delivery Restriction alert — Tier C3; and (c) user observation of wetness and insulin odor — Tier C4 [1].

Step 3. For control (a), the failure mode is an open-path leak. Occlusion detection interrogates the pressure channel in the *rising* direction; a tear lowers path impedance, so the pressure signature moves toward the healthy reading. Hence $P_{\text{man}} \approx 0$ for this failure mode on this channel, and by (1) the detection credit is approximately zero irrespective of P_{trans} and P_{dec}. This is the analytical statement of the FDA's plain-language instruction not to rely on pod alerts [1]. Note that the platform's own occlusion path is additionally rate-limited: Freckmann et al. showed detection latency at low basal rates extends well beyond what a rising-glucose trajectory tolerates [12], so even the faults the channel *can* see are credited generously.

For control (c), Table I assigns a credit of zero in the residual-risk arithmetic. Odor and wetness are not annunciated, not verified, and unavailable during sleep — and DKA arising from an overnight interruption is precisely the scenario the registry literature identifies as consequential [18].

Control (b) is the only channel with a defensible non-zero credit, and it is Tier C3: bounded by algorithm sensitivity, conditional on the closed loop being active, and latency-limited by physiology. The published operating points give the honest ceiling. Howsmon et al. report 88.0% sensitivity at 0.22 false positives per day under zone model predictive control and 73.3% in the sensor-augmented-pump arm [13]; unsupervised and machine-learning approaches report comparable ranges [14], [15]. Under (1), taking $P_{\text{man}} \approx 1$ (a real leak does move glucose), $P_{\text{trans}} \approx 0.95$ for CGM availability and trust, and $P_{\text{dec}} \approx 0.88$ at the published operating point, $D \approx 0.84$ for users in automated mode — and identically the C4 case, zero, for Omnipod DASH and Eros users who have no closed loop at all. The recall population includes all three product lines [1].

Steps 4–5. The safe state here is benign in the ISO 14971 sense: pod deactivation removes the leak. T_{tol} is hours [17], so the interruption is detection-deferrable and (2)–(3) are not binding. The binding constraint is entirely on the detection side.

Step 6. Derived requirements DR-1 through DR-4 in Table II.

2) B. Case B — ResMed Astral Fail-Safe Entry (FDA Early Alert, July 15, 2026): **Steps 1–2.** Two hazardous situations must be separated. H1 is *corrupted ventilation* caused by PCBA damage. H2 is *cessation of ventilation* caused by the device's own fault reaction. The credited control for H1 is the safety-monitoring subsystem that detects the corruption and commands the fail-safe state — Tier C2, and by all appearances it worked: the FDA record describes inadvertent *entry into* the fail-safe state, not undetected corrupted therapy [2]. The control for H2 is the high-priority alarm plus the instruction that alternative means of ventilation must be provided [2], [4].

Step 3. The C2 credit for H1 is defensible on this evidence. The instructive finding is elsewhere: the initiating mechanism is supercapacitor electrolyte leakage, a *wear-out*, and wear-out is bounded by lifetime and exposure analysis rather than by detection. The reliability literature makes this tractable — condition monitoring for DC-link capacitors [21], particle-filter prognostics for electrolytic capacitors under variable operating conditions [30], and calendar-ageing measurements showing strong thermal-cycling dependence [26], the last directly applicable to a portable ventilator whose thermal history spans hospital, vehicle, and home. The absence of a bounded replacement or inspection interval — the medical analogue of a proof-test interval [31] — is a finding independent of the alarm system.

Step 4. The fail-safe state is **life-withholding**. Therapy stops; the device "is no longer able to deliver therapy"; alternative means must be provided [2]. By Step 4 this state is escalated to hazardous situation H2 and returned to Step 1 — which is exactly the loop most risk files do not close, because "enter safe state" reads as a terminal node.

Step 5. Evaluate (2). T_{tol} for a ventilator-dependent patient is on the order of minutes. $A_{\text{r}}(T_{\text{tol}})$ is the probability that a caregiver hears the alarm, recognizes it, retrieves the alternative means, and restores ventilation within that window. Home-ventilation studies establish that malfunction is not rare and that the support infrastructure absorbing it is variable [27], [28]; alarm-response research establishes that even professionally staffed environments exhibit degraded response under alarm burden [25]. Suppose, illustratively, an acceptability target $P_{\text{H,tgt}} = 10^{-6}$ per patient-year for death or serious injury from unmitigated cessation, and an entry probability $P_{\text{e}} = 10^{-3}$ per patient-year over the affected build population. Then (3) requires $A_{\text{r,req}} \geq 1 - 10^{-3} = 0.999$. A claim that a single home caregiver, potentially asleep, restores ventilation within minutes on 999 occasions out of 1,000 is not supportable on the published evidence — and no such claim was ever required to be written down. The five reported serious injuries [2] are the empirical shadow of that shortfall.

The Astral's platform history sharpens the point. In August 2016 the Australian TGA recorded a Class I recall of the Astral 100 and 150 in which "an internal electrical issue has led to ventilation ceasing without either the low battery alarm or the critically low battery alarm being activated," with degraded internal battery packs replaced fleet-wide [7]. In March 2024, Germany's BfArM published a further urgent field safety notice for the same models concerning a degraded component in the power path [9]. Three power-subsystem field actions in a decade on one life-support platform is the signature of an unbounded wear-out class, not of three independent escapes — and the corrective-action effectiveness obligation under 21 CFR Part 820 attaches to the class, not to each instance.

Step 6. Derived requirements DR-5 through DR-8 in Table

II. DR-8 addresses a fact from the Early Alert that is a risk control in its own right: the correction is rationed by PCBA availability, with patients triaged into clinical risk tiers [2]. A triage scheme is a risk control measure under ISO 14971 and inherits the standard's verification obligation, but it is being executed post-market, under supply constraint, without the design-control evidence a pre-market control would require.

3) C. Derived Requirements:

E. Discussion

1) A. Limitations: The crediting ceilings in Table I are normative proposals rather than empirically derived constants. They are calibrated to the qualitative structure of the evidence available at each tier — a C0 channel can be characterized exhaustively on a bench, a C4 channel cannot be characterized at all — but a manufacturer with stronger evidence should be free to argue past them, and the framework's value lies in requiring that argument to be made explicitly rather than in the specific numbers.

The factorization in (1) assumes the three terms are independent. They are not entirely: threshold placement is chosen with knowledge of transducer noise, so P_{trans} and P_{dec} are coupled. The factorization remains useful as a diagnostic — it identifies which term is binding — but should not be used to compose a precise numerical detection probability from independently estimated parts.

Relation (2) treats responder availability as a scalar. In reality A_{r} is a distribution over caregivers, times of day, and clinical settings, and is strongly bimodal between attended institutional use and unattended home use. A more faithful treatment would evaluate (2) separately per use environment and take the worst case admitted by the intended-use statement. We chose the scalar form because the framework's purpose is to force the parameter onto the page; refining its distribution is straightforward once it exists.

2) B. Threats to Validity: Construct validity. We infer the Omnipod's detection architecture from the device class's documented behavior, from the FDA's own statements about what does and does not trigger an alert, and from the published occlusion-detection literature [1], [12]. We have no access to Insulet's risk file, and it is possible that the file already contains the analysis we describe as absent. Our claim is therefore about what the *standards* require to be written and verified, not a finding about a specific manufacturer's internal documentation.

Internal validity. The illustrative arithmetic in Section IV-B uses assumed values for $P_{\text{H,tgt}}$ and P_{e}. Different assumptions yield different required responder availabilities. The qualitative conclusion — that (3) demands an A_{r} close to unity for any plausible acceptability target when T_{tol} is minutes — is robust across the plausible range, because P_{e} for a fleet-wide wear-out mechanism with five reported serious injuries is many orders of magnitude above any acceptable harm target.

External validity. Both cases are single-manufacturer field actions in two device classes. The recurrence evidence for the pattern is stronger in infusion than in ventilation: Gao et al. found detection and alarm deficiencies across Class I infusion-pump recalls generally [20], and Alemzadeh et al. found recurring analytically foreseeable failure classes across medical-device regulatory databases [8]. We do not claim the framework generalizes beyond life-sustaining devices where a tolerance time is well defined.

3) C. Where the Standards Lens Stops: The framework is a lens on *what must be written and verified*. It does not resolve three questions that sit outside it.

First, whether a device should be marketed for unattended home use at all when (3) demands a responder availability the home setting cannot supply is a regulatory and clinical-policy judgment, not an engineering one. The framework can show that the assumption is unsupported; it cannot decide what to do about it.

Second, the framework treats the alternative means of ventilation as an external given. In practice its availability, its condition, and the caregiver's training on it are part of the same system, and a complete treatment would model them jointly.

Third, supply-constrained remediation — the rationing described in the Astral Early Alert [2] — raises an allocation question the risk-management standards do not address. ISO 14971 governs the risk of the device; it is silent on the risk of the *queue*.

F. VI. Conclusion and Future Work

Two FDA field actions thirteen days apart in July 2026 expose the same structural omission from opposite sides. In the Omnipod cannula-tear recall, a detection credit was taken for a fault that lies outside the device's sensing envelope, because the fluid-path surrogate is interrogated in only one direction and an open-path leak moves it the other way; the regulator's own remedy was to tell users to smell for insulin. In the ResMed Astral Early Alert, a safe-state credit was taken for a state that stops life-support ventilation, making its safety a claim about a caregiver rather than about the ventilator, on a platform with three power-subsystem field actions in a decade. Neither assumption was required by any artifact in the standard workflow to be stated, quantified, or verified.

The Detection–Responder Contract makes both computable. Table I forces every credited detection onto a named physical channel with a tier-appropriate evidence burden; (1) identifies which factor binds and therefore which design change can move it; and (2)–(3) convert the previously invisible responder assumption into a number a reviewer can challenge. The eight derived requirements of Table II are the concrete output for these two cases, and the most consequential of them are the cheapest: two-sided fault injection, and a single mandatory field in the risk file classifying every safe state as benign or life-withholding.

Future work should proceed in three directions. First, an empirical calibration of the Table I ceilings against a corpus of Class I recalls, extending the recall-analysis method of Gao et al. [20] with the channel taxonomy as the coding scheme, would replace the proposed ceilings with measured

TABLE II
DERIVED REQUIREMENTS FROM THE DETECTION–RESPONDER CONTRACT APPLIED TO THE TWO JULY 2026 FDA FIELD ACTIONS, WITH VERIFICATION METHOD AND THE FRAMEWORK ELEMENT THAT GENERATED EACH.

DR-1	Fault injection for fluid-path controls shall be two-sided: the campaign shall include faults that *decrease* path impedance as well as faults that increase it, and shall report detection outcome for each.	Bench fault-injection campaign; test report enumerating open-path faults	Step 3, P_{man} directionality
DR-2	No detection credit shall be taken in residual-risk arithmetic for a Tier C4 unaided-sensory channel; such cues may appear only as information for safety.	Risk-file review against Table I ceilings	Table I, tier C4
DR-3	Any C3 physiologic-surrogate credit shall state its operating point, shall be set to zero for product variants without the closed loop, and shall be re-verified whenever the control algorithm changes.	In-silico and clinical sensitivity report; variant matrix	Table I, tier C3; [13]–[15]
DR-4	Where lot-level manufacturing risk is known and the device authenticates its consumable at activation, the device shall alter its behavior for affected lots — escalated alerting, tightened surrogate thresholds, or refusal — rather than relying on out-of-band user communication.	Software requirement and unit verification per IEC 62304; traceability to the lot-identification function	Steps 1, 6
DR-5	Every safe state shall be classified benign or life-withholding in the risk file, and every life-withholding safe state entered as a device fault reaction shall be re-entered into hazard analysis as a hazardous situation.	Risk-file structural review; hazard-list completeness check	Step 4
DR-6	For each life-withholding safe state, the risk file shall state T_{tol}, the assumed responder availability A_{r}, and the evidence supporting it; unsubstantiated A_{r} shall be recorded as a gap and shall not default to unity.	Human-factors validation; comparison against (3)	Step 5, (2)–(3)
DR-7	Components with a known wear-out mechanism in the power path of a life-sustaining device shall carry a quantified life model and a bounded inspection or replacement interval, verified against the device's actual thermal and duty history.	Ageing test data; prognostic model; interval justification analogous to a proof-test interval	Step 3; [21], [26], [30], [31]
DR-8	A post-market remediation triage scheme shall be documented as a risk control measure, with stated tier criteria, an estimate of residual risk for patients in lower tiers during the wait, and a monitoring plan for that population.	CAPA record; residual-risk re-evaluation per ISO 14971 Clause 8	Step 6; [2]

ones. Second, the responder term A_{r} deserves a proper human-reliability treatment for the home-care setting, building on the integrated THERP and HEART machinery [32] and on the home-ventilation infrastructure data [27], [28]. Third, and most directly actionable, the fail-operational architectures developed for automotive by-wire systems [22], [23] have an unexploited analogue in life-support devices: a ventilator that degrades to a reduced but continuing mode on detected board corruption, rather than latching out of therapy, converts a life-withholding safe state into a benign one and removes the responder from the critical path entirely.

G. References

1. U.S. Food and Drug Administration, "Insulet recalls Omnipod insulin pods for risk of cannula tear leading to insulin leakage," Medical Device Recall Database, Class I classification July 2, 2026. [Online]. Available: https://www.fda.gov/medical-devices/medical-device-recalls
2. U.S. Food and Drug Administration, "Early alert: ResMed Astral 100 and Astral 150 ventilators may enter a fail-safe state due to supercapacitor leakage," Jul. 15, 2026. [Online]. Available: https://www.fda.gov/medical-devices/medical-device-recalls-and-early-alerts/early-alert-ventilator-issue-resmed
3. *Medical Devices — Application of Risk Management to Medical Devices*, ISO 14971:2019, International Organization for Standardization, Geneva, Switzerland, 2019. [Online]. Available: https://www.iso.org/standard/72704.html
4. *Medical Electrical Equipment — Part 1-8: General Requirements for Basic Safety and Essential Performance — Collateral Standard: General Requirements, Tests and Guidance for Alarm Systems in Medical Electrical Equipment, Medical Electrical Systems and Medical Software*, IEC 60601-1-8:2006+AMD2:2020, International Electrotechnical Commission, Geneva, Switzerland.
5. *Medical Electrical Equipment — Part 2-72: Particular Requirements for Basic Safety and Essential Performance of Home Healthcare Environment Ventilators for Ventilator-Dependent Patients*, ISO 80601-2-72:2015, International Organization for Standardization, Geneva, Switzerland.
6. *Medical Electrical Equipment — Part 2-24: Particular Requirements for the Basic Safety and Essential Performance of Infusion Pumps and Controllers*, IEC 60601-2-24:2012, International Electrotechnical Commission, Geneva, Switzerland.
7. Australian Therapeutic Goods Administration,

"Recall RC-2016-RN-01074-1: Astral 100 and Astral 150 ventilators — degraded battery pack," Aug. 2016. [Online]. Available: https://medicaldevices.icij.org/events/aus-astral-100-and-astral150-ventilators-degraded-battery-pack

8. H. Alemzadeh, R. K. Iyer, Z. Kalbarczyk, and J. Raman, "Analysis of safety-critical computer failures in medical devices," *IEEE Secur. Privacy*, vol. 11, no. 4, pp. 14–26, Jul.–Aug. 2013, doi: 10.1109/MSP.2013.49.
9. Bundesinstitut für Arzneimittel und Medizinprodukte, "Urgent field safety notice 08339/24: ResMed Astral 100 / Astral 150," Mar. 25, 2024. [Online]. Available: https://www.bfarm.de/EN/Medical-devices/_node.html
10. A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Trans. Dependable Secure Comput.*, vol. 1, no. 1, pp. 11–33, Jan.–Mar. 2004, doi: 10.1109/TDSC.2004.2.
11. L. E. Blanco, J. H. Wilcox, M. S. Hughes, and R. A. Lal, "Development of a real-time force-based algorithm for infusion failure detection," *J. Diabetes Sci. Technol.*, Apr. 2024, doi: 10.1177/19322968241247530.
12. G. Freckmann, U. Kamecke, D. Waldenmaier, C. Haug, and R. Ziegler, "Occlusion detection time in insulin pumps at two different basal rates," *J. Diabetes Sci. Technol.*, vol. 12, no. 3, pp. 608–613, May 2018, doi: 10.1177/1932296817750404.
13. D. P. Howsmon, N. Baysal, B. A. Buckingham, G. P. Forlenza, T. T. Ly, D. M. Maahs, E. Dassau, J. Hahn, and B. W. Bequette, "Real-time detection of infusion site failures in a closed-loop artificial pancreas," *J. Diabetes Sci. Technol.*, vol. 12, no. 3, pp. 599–607, May 2018, doi: 10.1177/1932296818755173.
14. L. Meneghetti, G. A. Susto, and S. Del Favero, "Detection of insulin pump malfunctioning to improve safety in artificial pancreas using unsupervised algorithms," *J. Diabetes Sci. Technol.*, vol. 13, no. 6, pp. 1065–1076, Nov. 2019, doi: 10.1177/1932296819881452.
15. L. Meneghetti, E. Dassau, F. J. Doyle III, and S. Del Favero, "Machine learning-based anomaly detection algorithms to alert patients using sensor augmented pump of infusion site failures," *J. Diabetes Sci. Technol.*, 2022, doi: 10.1177/1932296821997854.
16. K. Turksoy, A. Roy, and A. Cinar, "Real-time model-based fault detection of continuous glucose sensor measurements," *IEEE Trans. Biomed. Eng.*, vol. 64, no. 7, pp. 1437–1445, Jul. 2017, doi: 10.1109/TBME.2016.2535412.
17. D. C. Klonoff, A. T. Ayers, C. N. Ho, C. Fabris, M. F. Villa-Tamayo, E. Cengiz, L. Ekhlaspour, J. C. Wong, L. Heinemann, and M. A. Kohn, "Time to moderate and severe hyperglycemia and ketonemia following an insulin pump occlusion," *J. Diabetes Sci. Technol.*, 2024, doi: 10.1177/19322968241280386.
18. H. Kalscheuer, J. Seufert, S. Lanzinger, J. Rosenbauer, W. Karges, D. Bergis, J. K. Mader, S. Zimny, S. M. Schmid, M. Hummel, W. Kerner, and R. W. Holl, "Event rates and risk factors for the development of diabetic ketoacidosis in adult patients with type 1 diabetes: Analysis from the DPV registry based on 46,966 patients," *Diabetes Care*, vol. 42, no. 3, pp. e34–e36, Mar. 2019, doi: 10.2337/dc18-1160.
19. S. L. Orbell, R. A. Codario, M. F. Zupa, and J. L. Estock, "Severe insulin pump-related adverse events: Potential root causes and impact of the COVID-19 pandemic," *J. Diabetes Sci. Technol.*, 2024, doi: 10.1177/19322968241254521.
20. X. Gao, Q. Wen, X. Duan, W. Jin, X. Tang, L. Zhong, S. Xia, H. Feng, and D. Zhong, "A hazard analysis of Class I recalls of infusion pumps," *JMIR Hum. Factors*, vol. 6, no. 2, p. e10366, May 2019, doi: 10.2196/10366.
21. Z. Zhao, P. Davari, W. Lu, H. Wang, and F. Blaabjerg, "An overview of condition monitoring techniques for capacitors in DC-link applications," *IEEE Trans. Power Electron.*, vol. 36, no. 4, pp. 3692–3716, Apr. 2021, doi: 10.1109/TPEL.2020.3023469.
22. P. Sinha, "Architectural design and reliability analysis of a fail-operational brake-by-wire system from ISO 26262 perspectives," *Rel. Eng. Syst. Saf.*, vol. 96, no. 10, pp. 1349–1359, Oct. 2011, doi: 10.1016/j.ress.2011.03.013.
23. Y. Fu, A. Terechko, J. F. Groote, and A. K. Saberi, "A formally verified fail-operational safety concept for automated driving," *SAE Int. J. Connect. Autom. Veh.*, vol. 5, no. 1, pp. 7–21, Jan. 2022, doi: 10.4271/12-05-01-0002.
24. N. Nolan and O. McDermott, "Failure mode effect analysis use and limitations in medical device risk management," *J. Open Innov. Technol. Market Complexity*, vol. 11, no. 1, p. 100439, Mar. 2025, doi: 10.1016/j.joitmc.2024.100439.
25. O. M. Cho, H. Kim, Y. W. Lee, and I. Cho, "Clinical alarms in intensive care units: Perceived obstacles of alarm management and alarm fatigue in nurses," *Healthc. Inform. Res.*, vol. 22, no. 1, pp. 46–53, Jan. 2016, doi: 10.4258/hir.2016.22.1.46.
26. M. Ayadi, O. Briat, A. Eddahech, R. German, G. Coquery, and J. M. Vinassa, "Thermal cycling impacts on supercapacitor performances during calendar ageing," *Microelectron. Rel.*, vol. 53, no. 9–11, pp. 1628–1631, Sep. 2013, doi: 10.1016/j.microrel.2013.07.079.
27. S. Srinivasan, S. M. Doty, T. R. White, V. H. Segura, M. T. Jansen, and S. L. Davidson Ward, "Frequency, causes, and outcome of home ventilator failure," *Chest*, vol. 114, no. 5, pp. 1363–1367, Nov. 1998, doi: 10.1378/chest.114.5.1363.
28. M. Chatwin, S. Heather, A. Hanak, M. I. Polkey, and A. K. Simonds, "Analysis of home support and ventilator malfunction in 1,211 ventilator-dependent patients," *Eur. Respir. J.*, vol. 35, no. 2, pp. 310–316, Feb. 2010, doi: 10.1183/09031936.00073409.
29. H. Wang and F. Blaabjerg, "Reliability of capac-

itors for DC-link applications in power electronic converters—An overview," *IEEE Trans. Ind. Appl.*, vol. 50, no. 5, pp. 3569–3578, Sep./Oct. 2014, doi: 10.1109/TIA.2014.2308357.

30. M. Rigamonti, P. Baraldi, E. Zio, D. Astigarraga, and A. Galarza, "Particle filter-based prognostics for an electrolytic capacitor working in variable operating conditions," *IEEE Trans. Power Electron.*, vol. 31, no. 2, pp. 1567–1575, Feb. 2016, doi: 10.1109/TPEL.2015.2418198.
31. H. Jin and M. Rausand, "Reliability of safety-instrumented systems subject to partial testing and common-cause failures," *Rel. Eng. Syst. Saf.*, vol. 121, pp. 146–151, Jan. 2014, doi: 10.1016/j.ress.2013.08.006.
32. F. Castiglia, M. Giardina, and E. Tomarchio, "THERP and HEART integrated methodology for human error assessment," *Radiat. Phys. Chem.*, vol. 116, pp. 262–266, Nov. 2015, doi: 10.1016/j.radphyschem.2014.12.012.