When the Credited Mechanism Is Off Duty: A State-Indexed Availability Framework for Safety-Mechanism Claims in ISO 26262 Safety Cases

Jherrod Thomas

Independent Researcher — The Lion of Functional Safety
Certified Functional Safety Expert (ISO 26262) · ISO/SAE 21434 (TÜV SÜD) · SOTIF Advanced · PMP
jherrodthomas.com · July 15, 2026

***Abstract*— Two 2026 United States vehicle-safety recalls expose a shared blind spot in automotive functional-safety argumentation: the safety mechanism credited in the safety case is quantified as if it were continuously on duty, when in fact its availability is gated by the vehicle's operating state. In NHTSA recall 26V402, 741,195 Ford and Lincoln park-by-wire vehicles rely on an electronic-park-brake auto-apply and an instrument-cluster telltale to mitigate rollaway after latent park-system damage — but the Powertrain Control Module that hosts the detection "may not be awake" after the vehicle has been powered down, precisely the state in which a parked vehicle rolls away. In NHTSA recall 26V030, 43,881 Volkswagen ID.4 vehicles lacked the Self-Discharge Detection software whose warning, by the manufacturer's own account, would have preceded at least three known battery-fire incidents: an availability of identically zero. We propose a six-step framework that makes availability a first-class, state-indexed property of every credited safety mechanism. The framework enumerates vehicle operating states, marks the states in which each hazard can manifest, assigns each mechanism a per-state availability, and classifies mechanisms as demand-synchronous or detection-deferrable. Three numbered relations convert the state-indexed availability into an effective diagnostic coverage, a residual violation rate, and an effective latent-fault exposure that diverges as availability approaches zero. A worked example on the Ford recall shows illustratively that a sleep-gated monitor can silently collapse a claimed 99% diagnostic coverage to below 10%, degrading a latent-fault-metric claim that ISO 26262-5 requires to exceed 90% at ASIL D. Five derived requirements close the gap.*

Index Terms*—*ISO 26262, safety mechanism, latent fault metric, operating state, availability, dependent failure analysis, vehicle rollaway, battery management system.

I. INTRODUCTION

Every quantitative claim in an ISO 26262 safety case ultimately rests on a small set of multiplications. A fault rate is multiplied by a diagnostic coverage; a residual rate is multiplied by an exposure time; and the products are compared against the single-point, latent-fault, and probabilistic targets that ISO 26262-5 assigns per Automotive Safety Integrity Level (ASIL) [3]. The diagnostic-coverage factor in those multiplications carries an assumption so ingrained that it is rarely written down: that the safety mechanism supplying the coverage is *on duty* whenever the fault it covers can do harm. For mechanisms hosted on a continuously powered microcontroller monitoring a continuously operating function, the assumption is harmless. But the modern vehicle is not continuously operating. It spends the large majority of its calendar life parked, its electronic control units cycling through shutdown windows, deep-sleep states, and partial wake events under the direction of network-management and power-budget logic that was designed against quiescent-current targets, not against the hazard-manifestation profile of the safety case.

Two 2026 recalls show what happens in the gap. The first is NHTSA recall 26V402 (filed June 24, 2026), covering 741,195 Ford and Lincoln trucks and sport-utility vehicles with 10R60 or 10R80MHT park-by-wire transmissions [1]. A valve-body separator plate can restrict flow to the park valve, causing momentary park-pawl engagement during certain shifts; the repeated momentary engagements damage the park system over time, so that the transmission may eventually fail to hold the vehicle in Park. The damage is a textbook latent fault, and the vehicle carries two credited mitigations: a wrench telltale in the instrument cluster, and an electronic park brake (EPB) that auto-applies when the transmission range sensor does not confirm Park. But Ford's own Part 573 report states the limit of the second mitigation: the Powertrain Control Module (PCM) "may not be awake in some instances after the vehicle has been powered down after a period of time and unable to detect movement" [1]. The rollaway hazard of a parked vehicle manifests, almost by definition, after the vehicle has been powered down for a period of time. The credited mechanism and the hazard occupy disjoint regions of the operating-state space.

The second is NHTSA recall 26V030 (filed January 21, 2026), covering 43,881 model-year 2023–2025 Volkswagen ID.4 vehicles whose high-voltage battery modules may self-discharge and overheat due to supplier manufacturing deviations [2]. The recall population is defined, remarkably, not by the defective module alone but by the *absence of a software function*: vehicles outside the recall carry Self-Discharge Detection (SDD) software; vehicles inside it do not. Volkswagen's report states that SDD "would have triggered a warning in advance of at least three (3) known incidents" [2]. Here the availability of the credited detection is not gated by a sleep state — it is identically zero, because the mechanism, though designed, validated, and deployed in sibling production, was never allocated to these vehicles. A cluster of instrument-panel recalls in the same period supplies a third variant of the pattern: a display component developed at Quality Management (QM) integrity silently carrying ASIL-rated warning telltales, so that the annunciation path credited by upstream safety goals fails without any dedicated integrity argument of its own.

These three variants — asleep, absent, and under-rated —

are conventionally treated as unrelated defects. We argue they are the same defect: the safety case asserted a mechanism's coverage without asserting its *availability*, and no artifact in the standard workflow forced the two claims to be confronted state by state. This paper contributes: (i) a formalization of the *availability contract* — the tuple binding a hazard, its credited mechanism, the set of operating states in which the hazard can manifest, and the mechanism's per-state availability; (ii) a classification of mechanisms as demand-synchronous or detection-deferrable, with distinct crediting criteria; (iii) three numbered relations converting state-indexed availability into effective diagnostic coverage, residual violation rate, and effective latent-fault exposure; (iv) a six-step audit method with an availability audit matrix suitable for confirmation review; and (v) a worked example on recall 26V402 with 26V030 as a boundary case, yielding five derived requirements.

Section II reviews background and related work. Section III presents the framework. Section IV applies it to the 2026 recalls. Section V discusses limitations. Section VI concludes.

II. BACKGROUND AND RELATED WORK

A. Hardware Architectural Metrics and Their Hidden Availability Assumption

ISO 26262-5 evaluates a hardware architecture against the single-point fault metric (SPFM), the latent fault metric (LFM), and the probabilistic metric for random hardware failures (PMHF), with LFM targets of at least 60%, 80%, and 90% at ASIL B, C, and D respectively, and PMHF targets of 10^{-7}/h at ASIL B and C and 10^{-8}/h at ASIL D [3]. All three metrics are functions of diagnostic coverage. The standard's own simplified PMHF formulation, and the more rigorous treatments developed since, quantify the dual-point contribution as a product of a latent-fault term and an exposure interval. Sakurai derived generic PMHF equations from continuous-time Markov chains and point-unavailability functions of periodically inspected items, identifying multiple faults in the standard's Part 10 derivation [6], and subsequently resolved a dilemma between the standard's two published formula families [7]. Das and Taylor showed how quantified fault trees compute the same metrics and made the multiple-point-fault detection interval an explicit fault-tree parameter [8]. In every one of these formulations, the diagnostic coverage attributed to a safety mechanism enters as a scalar. The formulations permit — but no artifact in the standard workflow demands — the question this paper centers: *over which operating states was that scalar averaged, and is the average valid over the states where the hazard lives?*

The process industries answered the analogous question decades ago, because their protective functions are explicitly demand-driven. IEC 61508 distinguishes low-demand from high-demand modes and quantifies the average probability of failure on demand of a safety instrumented function, with the proof-test interval controlling the unavailability window; Jin, Lundteigen, and Rausand unified the low- and high-demand treatments in a single framework [9]. The automotive standard has no equivalent of the "demand mode" concept for its safety mechanisms: a mechanism is described by what it detects and how well, not by *when it is on duty relative to the demand*. The power-systems community likewise learned, from cascading-outage analysis, that protection elements carry "hidden failures" — defects invisible until the protected condition occurs — and that exposing them requires modeling the protection system itself as a failure source [10]. Our framework imports exactly this posture into ISO 26262: the safety mechanism is a component with its own availability profile, not a guaranteed property of the architecture.

B. Verification Practice: Fault Injection Exercises the Mechanism in the Wrong States

ISO 26262 names fault injection as a key technique for demonstrating safety-mechanism effectiveness, and a substantial literature has industrialized it. Pintard *et al.* mapped fault-injection obligations across the ISO 26262 lifecycle [11]; recent virtualized frameworks such as QEFIRA inject permanent and transient faults into digital components at simulation speed and classify outcomes against Part 11 failure modes [12]. But fault-injection campaigns overwhelmingly exercise the system in its powered, operational configuration — the state in which the test bench, the debugger, and the fault injector can all run. A campaign that demonstrates 99% diagnostic coverage in the driving state demonstrates nothing about the shutdown window, the deep-sleep state, or the wake-event state unless those states are explicitly included in the campaign's state matrix. The verification gap mirrors the analysis gap: both inherit the unexamined assumption that the mechanism's duty cycle covers the hazard's.

C. Runtime Monitoring and the Battery-Fire Detection Literature

A parallel research thread builds runtime monitors for automated driving: REDriver enforces signal-temporal-logic specifications over planned trajectories at runtime [13]; Gautham *et al.* derive multilevel runtime monitors from STPA hazard analysis for in-time hazard detection [14]; and the connected dependability cage line extends function and anomaly monitoring across development and operation, explicitly targeting ISO 26262 and ISO/PAS 21448 compliance [15]. This literature is sophisticated about *what* to monitor and *how* to detect, but its deployment context is again the operating vehicle; the parked, unpowered vehicle is out of scope by construction. Yet the human backstop assumed in its place is weak: driver-response studies show takeover and response latencies of many seconds even for alert, in-seat drivers [16] — and the driver of a parked vehicle is, in the bounding case, not present at all.

For the specific hazard of recall 26V030, the battery literature is mature: thermal-runaway mechanisms and their precursors are well characterized [17]; soft internal short circuits can be detected from extracted open-circuit-voltage and capacity

signatures during pack operation [18]; and recent equivalent-circuit models such as BattBee support early, computationally cheap detection of internal-short-induced runaway onset [19]. The engineering content of Self-Discharge Detection was therefore neither novel nor speculative at the time the affected vehicles were built. What failed was not detection science but allocation: a known, implementable mechanism carried an availability of zero on a population that needed it. UN Global Technical Regulation No. 20 requires warning of a condition that may lead to thermal runaway [5]; the availability contract of Section III is, in effect, the artifact that would have forced the allocation question to the surface.

III. APPROACH

A. The Availability Contract

Let h be a hazardous event from the ISO 26262-3 hazard analysis, and let m be a safety mechanism credited in the functional or technical safety concept with mitigating a fault path leading to h. Let $S = \{s_1, \ldots, s_n\}$ be the vehicle's operating states, enumerated at the granularity at which power, network, and ECU wake status actually change — at minimum: driving; shutdown window (ignition off, hosts still awake); deep sleep; and wake events (unlock, door, charging, telematics). Define the *hazard-active set* $S_h \subseteq S$ as the states in which h can manifest, and for each state a *hazard weight* $\pi_h(s)$: the fraction of hazard demands expected to arrive in state s. Define the mechanism's *availability* $a_m(s) \in [0,1]$: the probability that m is powered, executing, and able to complete its detection-plus-reaction within its allotted time when a demand arrives in state s. The *availability contract* is the tuple $\langle h, m, S_h, \pi_h, a_m \rangle$, and the *duty factor* of the mechanism with respect to the hazard is the hazard-weighted mean availability

$$A_{m|h} \;=\; \sum_{s \in S_h} \pi_h(s)\, a_m(s), \tag{1}$$

with π_h normalized over S_h. The Ford PCM sleep gap is $a_m(\text{deep sleep}) = 0$ with $\pi_h(\text{deep sleep})$ large; the Volkswagen SDD gap is $a_m(s) = 0$ for all s; the under-rated cluster is a_m nominally 1 but unsubstantiated at the claimed integrity.

B. Two Crediting Classes

The consequence of an availability gap depends on whether the mechanism must act at the instant of demand. We classify every credited mechanism as one of:

Class I (demand-synchronous). The mechanism's reaction must complete within the fault-tolerant time interval measured from the demand itself — e.g., the EPB auto-apply that must arrest a vehicle that has begun to move. For Class I, availability is binding state by state: the crediting criterion is $a_m(s) \geq a_{\min}$ for *every* $s \in S_h$, because a demand arriving in a single uncovered state defeats the mechanism regardless of how well it performs elsewhere. Averaging is not permissible.

Class D (detection-deferrable). The mechanism detects a latent condition whose harm requires a further coincidence, so detection may lawfully wait until the next state in which the mechanism runs — e.g., a telltale that warns of accumulated park-system damage at the next ignition cycle, or SDD flagging a self-discharging module hours before thermal runaway [17]. For Class D, the hazard-weighted duty factor (1) meaningfully discounts the claimed diagnostic coverage:

$$\mathrm{DC}_{\text{eff}} = A_{m|h}\,\mathrm{DC}_m, \qquad \lambda_{\text{res}} = \lambda_h\,(1 - \mathrm{DC}_{\text{eff}}), \tag{2}$$

where DC_m is the coverage demonstrated by fault injection in the states where the mechanism runs, and λ_{res} is the residual rate of undetected hazard-capable faults that enters the LFM and the dual-point PMHF term. A Class D mechanism additionally carries a timing obligation: deferral is lawful only while the expected time to the next detection opportunity remains inside the multiple-point-fault detection interval. If the mechanism runs only in states entered at mean rate proportional to its duty factor, the effective latent exposure stretches as

$$T_{\text{exp,eff}} \;=\; T_{\text{exp},0} \,/\, A_{m|h}, \tag{3}$$

which diverges as $A_{m|h} \to 0$: a mechanism that is never on duty when the hazard can manifest provides unbounded exposure, which is precisely the Volkswagen boundary case. Equations (1)–(3) are deliberately simple; their role is not numerical precision but *forcing function* — they cannot be evaluated at all until the analyst has produced the state enumeration and per-state availabilities that current practice leaves implicit.

C. The Six-Step Audit

The framework runs as six steps, consistent with the audit series this paper extends. **S1 — evidence ingestion:** collect the credited mechanisms from the functional and technical safety concepts, and any field evidence (recalls, investigations) indicating availability gaps. **S2 — hazard localization:** for each mechanism, identify the hazard(s) it is credited against and the ASIL of the governing safety goal. **S3 — state enumeration:** enumerate operating states at power/network granularity and mark the hazard-active set S_h with hazard weights; for parked-vehicle hazards, weights follow parking-duration statistics rather than driving-cycle statistics. **S4 — availability assignment:** assign $a_m(s)$ from the power-management, network-management, and partial-networking design (which ECUs sleep, what wakes them, what runs during wake), and classify each mechanism Class I or Class D. **S5 — quantification:** evaluate (1)–(3); flag every Class I mechanism with any uncovered hazard-active state and every Class D mechanism whose DC_{eff} falls below the LFM-implied coverage floor for the target ASIL. **S6 — requirement derivation:** for each flag, derive a requirement that either (a) extends the mechanism's availability into the uncovered state, (b) reallocates the function to a host that is available there, (c) substitutes a passive/mechanical mechanism with no duty cycle, or (d) demonstrates by dependent-failure analysis per ISO 26262-9 that the

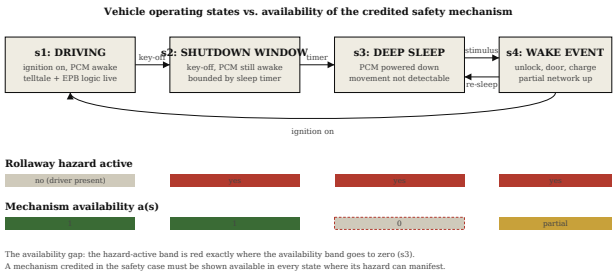

FIG. 1: The availability contract visualized for the rollaway hazard of recall 26V402. Top: operating states at power-management granularity. Middle band: states where the rollaway hazard is active. Bottom band: availability of the PCM-hosted detection-plus-EPB-apply mechanism. The audit flags the intersection of a hazard-active band with a zero availability band (state s_3, deep sleep).

uncovered state cannot carry a demand [3]. Fig. 1 shows the structure of the audit's central artifact; Table 1 instantiates it.

D. Relation to Dependent Failure Analysis

Step S6(d) deserves emphasis. The availability gap is formally a *dependent failure* between the mechanism and the vehicle's power state: the same event (sleep-timer expiry) that removes the mechanism also, in the Ford case, leaves the hazard armed. ISO 26262-9's dependent-failure analysis owns coupling factors of exactly this kind, but its practiced checklists concentrate on shared supplies, shared clocks, and shared silicon — couplings *within* the powered architecture. The operating-state coupling is a shared *temporal* resource, and we propose it be added to DFA coupling-factor catalogs explicitly: any mechanism whose host participates in network-managed sleep shares a coupling factor with every hazard whose active set includes the slept states.

IV. WORKED EXAMPLE: THE 2026 ROLLAWAY AND BATTERY-FIRE RECALLS

A. Ford 26V402 as a Class I / Class D Composite

We instantiate the framework on recall 26V402 [1]. In S1–S2, the latent fault is progressive park-system damage from momentary pawl engagements; the hazardous event is rollaway of a parked vehicle; and a HARA in the style of ISO 26262-3 would plausibly rate the governing safety goal ASIL C or D depending on exposure and controllability assumptions for occupied-versus-unoccupied scenarios — we do not assert Ford's internal rating, only that FMVSS 114's rollaway-prevention purpose [4] and the fatality potential of an unattended moving vehicle place it well above QM. Two mechanisms are credited: m_1, the wrench telltale annunciating detected park-system faults (Class D — it warns the driver to seek service and use the parking brake); and m_2, the EPB auto-apply on failure of the transmission range sensor to confirm Park, extended by movement detection while parked (Class I — it must act when the vehicle begins to move).

In S3, the operating states are those of Fig. 1: s_1 driving, s_2 shutdown window, s_3 deep sleep, s_4 wake events. The rollaway hazard-active set is $\{s_2, s_3, s_4\}$ — the vehicle must be parked — and the hazard weights follow parked-time occupancy. A passenger vehicle is parked roughly 95% of calendar time; if the PCM remains awake for a bounded shutdown window (tens of minutes) and the mean parking event lasts many hours, the fraction of parked time with the PCM awake is small. For illustration only — these are not Ford's numbers — take a 30-minute awake window against a 10-hour mean parking event: $\pi_h(s_2) \approx 0.05$, $\pi_h(s_3) \approx 0.94$, $\pi_h(s_4) \approx 0.01$, assuming rollaway demands arrive approximately uniformly over parked time (settling and thermal effects arguably front-load demands slightly, which is the charitable direction for the mechanism).

In S4–S5, availability assignment uses Ford's own report: $a_{m_2}(s_2) \approx 1$, but in deep sleep the PCM is "not awake … unable to detect movement" [1], so $a_{m_2}(s_3) \approx 0$, and wake events give partial coverage. As a Class I mechanism, m_2 fails the crediting criterion outright: state s_3 is hazard-active and uncovered, and no averaging can repair a demand-synchronous mechanism that is absent when the demand arrives. Had m_2 been (incorrectly) treated as Class D and averaged, (1) gives $A \approx 0.05 \cdot 1 + 0.94 \cdot 0 + 0.01 \cdot 0.5 \approx 0.055$, and by (2) a fault-injection-demonstrated DC_m of 99% collapses to $\mathrm{DC}_{\mathrm{eff}} \approx 5.4\%$ — far below the 90% LFM floor that ISO 26262-5 sets at ASIL D and even the 60% floor at ASIL B [3]. The arithmetic dramatizes the point rather than measures the vehicle: whether the true duty factor is 0.03 or 0.10, the claimed coverage and the delivered coverage differ by an order of magnitude, and the discrepancy is invisible to any artifact that records diagnostic coverage as a scalar. The telltale m_1, by contrast, passes its Class D timing test only if the damage-to-failure progression is slow compared to the driving-cycle cadence — plausible for progressive mechanical damage, but an assumption the safety case must now state and defend rather than inherit silently. Table 1 summarizes the audit; note the remedy in the actual recall (a PCM software update [1]) acts on availability logic, exactly where the framework locates the defect.

B. Volkswagen 26V030 as the Boundary Case

The ID.4 recall exercises the framework's boundary. The latent fault is cell-module self-discharge from supplier manufacturing deviations; the hazard is thermal runaway and fire, whose active set includes parked and charging states [2, 17]. The credited-mechanism question is inverted: SDD existed as engineering (deployed on vehicles outside the recall; the detection science is established [18, 19]) but $a_m(s) = 0$ on the recall population in every state, because the software was never installed. By (1), $A = 0$; by (3), the effective latent exposure is unbounded; and the field record supplies the empirical confirmation — at least three incidents that SDD's warning "would have" preceded [2]. Note what the framework adds over hindsight: the availability contract is an artifact whose *absence* is checkable at design time. A confirmation review holding the

TABLE 1
AVAILABILITY AUDIT MATRIX FOR THE ROLLAWAY HAZARD OF NHTSA RECALL 26V402 (ILLUSTRATIVE)

Mechanism (class)	s_1 driving	s_2 shutdown ($\pi \approx 0.05$)	s_3 deep sleep ($\pi \approx 0.94$)	s_4 wake ($\pi \approx 0.01$)	Audit result
m_1: wrench telltale (D)	1	1	0	0	Pass iff dam vs. drive c must be stat
m_2: EPB auto-apply + movement detect (I)	1 (not hazard-active)	$\approx$1	$\approx$0 (PCM asleep [1])	partial	**Fail:** hazarc ered; Class
m_3: driver perception of movement (D)	1	low	$\approx$0 (driver absent)	$\approx$0	Not credita rollaway [16

tuple $\langle$thermal-runaway hazard, SDD, parked/charging states, $a = 0\rangle$ in hand cannot close, whereas a review holding only a functional-safety-concept sentence crediting "BMS monitoring" can and evidently did. UN GTR No. 20's thermal-event warning requirement [5] gives the contract regulatory teeth; the framework gives it a place to be written down.

C. Derived Requirements

Table 2 lists five derived requirements. DR-1 through DR-3 act on the mechanism; DR-4 and DR-5 act on the process artifacts.

V. DISCUSSION

A. Limitations and Threats to Validity

The numerical values in Section IV are illustrative order-of-magnitude figures chosen to exercise relations (1)–(3); they are not Ford or Volkswagen engineering data, and we do not have access to either manufacturer's HARA, actual sleep-timer values, parking-duration fleet statistics, or the true demand profile of rollaway events over parked time. The uniform-demand assumption for π_h is the largest single modeling risk: if park-system releases cluster in the minutes after shutdown (while hydraulic pressure decays and the PCM is still awake), the delivered availability is better than our illustration; if they cluster after thermal cycling overnight, it is worse. Equation (3) assumes detection opportunities arrive as a rate process scaled by the duty factor, which is coarse for strongly periodic usage patterns. The framework also inherits the hazard analysis it audits: a wrong hazard-active set — for instance, omitting the charging state for a battery hazard — propagates through every downstream term. Finally, our classification of the public recalls into framework categories is an external reconstruction from Part 573 filings [1, 2]; the internal safety cases may contain arguments we cannot see, and the framework's claim is precisely that such arguments should be *visible* in an auditable artifact, not that they were never made.

B. Where the Standards Lens Stops

ISO 26262 does not forbid any of this analysis; Parts 4 and 5 require operating modes to be considered, and Part 9's DFA could in principle catch the sleep coupling. The gap is one of *defaults*: the metrics that gate release (SPFM, LFM, PMHF) accept scalar coverage values, the verification techniques that feed them run on powered benches, and no mandatory work product confronts mechanism duty cycles with hazard-active states. The framework is thus a proposal about artifacts, not about physics. Its own lens stops in three places. It cannot supply the parking-duration and demand-arrival statistics that make (1) quantitative — those require fleet data. It does not address security-motivated availability loss (an attacker suppressing wake events), which couples this analysis to ISO/SAE 21434. And it cannot force the organizational decision that the Volkswagen case turns on: when a mechanism exists in one build and not another, the availability contract makes the discrepancy visible, but only the release process can make it blocking.

VI. CONCLUSION AND FUTURE WORK

Two 2026 recalls — 741,195 Ford park-by-wire vehicles whose rollaway mitigation may be asleep when the rollaway occurs, and 43,881 Volkswagen ID.4 vehicles whose battery-fire warning function was never installed — exhibit the same structural defect: a safety mechanism credited in the safety case without a state-indexed availability claim to back it. We have proposed the availability contract as a first-class safety-case artifact, classified mechanisms as demand-synchronous or detection-deferrable with distinct crediting criteria, and given three relations that collapse claimed coverage into delivered coverage as availability degrades — including the divergence of effective latent exposure as availability approaches zero, which is the Volkswagen boundary case exactly. A worked audit on the Ford recall shows the pattern is detectable at design time from artifacts the developer already possesses: the power-management specification and the hazard analysis, confronted in one matrix.

Future work runs in three directions. First, empirical grounding: parking-duration and demand-arrival distributions from fleet telematics would convert (1)–(3) from forcing functions into measurements, and would let the deferral bound of DR-2 be set with stated confidence. Second, formalization: the availability contract is a natural candidate for machine-checkable safety-case notations, and the state-indexed cover-

TABLE 2
DERIVED REQUIREMENTS

ID	Requirement	Term controlled	Case
DR-1	Every Class I mechanism shall be shown available ($a \geq a_{\min}$) in each hazard-active state, or the function shall be reallocated to a host (or passive device) available there — e.g., movement detection on an always-powered domain, or mechanical EPB engage-on-park default	$a(s)$, state-by-state	26V402
DR-2	Class D mechanisms shall carry an explicit deferral bound: expected time to next detection opportunity shall be shown less than the multiple-point-fault detection interval under worst-case usage (long-term parking)	$T_{\text{exp,eff}}$ via (3)	26V402
DR-3	Deployment of a designed safety mechanism shall be verified per configuration/population; a mechanism credited in the concept but absent from a build shall block release (allocation audit)	$A = 0$ boundary	26V030
DR-4	Diagnostic coverage shall be recorded state-indexed, and fault-injection campaigns [11, 12] shall include shutdown, sleep, and wake states in the campaign matrix	DC_{eff} via (2)	both
DR-5	DFA coupling-factor catalogs per ISO 26262-9 shall include operating-state coupling: any mechanism hosted on a sleep-managed ECU shares a coupling factor with every hazard active in slept states	dependent failure	both

age could be integrated into the Markov-based PMHF formulations [6, 7] rather than applied as a correction factor. Third, extension across domains: aircraft certification already prices dormancy into exposure time via certification maintenance requirements, and the industrial sector prices it via proof-test intervals [9]; a unified treatment of "who watches the watchman, and when is the watchman on shift" across ISO 26262, IEC 61508, and ARP4761A would let each domain borrow the others' strongest artifact. The parked vehicle is the automotive industry's low-demand mode; it deserves the same quantitative respect the process industries have given theirs for forty years.

REFERENCES

[1] Ford Motor Company, "Part 573 safety recall report 26V-402: Park system damage — 10R60/10R80MHT park-by-wire," Nat. Highway Traffic Safety Admin., Washington, DC, USA, Jun. 24, 2026. [Online]. Available: `https://static.nhtsa.gov/odi/rcl/2026/RCLRPT-26V402-1380.pdf`

[2] Volkswagen Group of America, "Part 573 safety recall report 26V-030: High-voltage battery may overheat," Nat. Highway Traffic Safety Admin., Washington, DC, USA, Jan. 21, 2026. [Online]. Available: `https://static.nhtsa.gov/odi/rcl/2026/RCLRPT-26V030-0889.pdf`

[3] *ISO 26262: Road Vehicles — Functional Safety*, Parts 1–12, 2nd ed., Int. Org. Standardization, Geneva, Switzerland, 2018.

[4] *Electronic Code of Federal Regulations, Title 49 §571.114 — Standard No. 114; Theft protection and roll-away prevention*. [Online]. Available: `https://www.ecfr.gov/current/title-49/section-571.114`

[5] *UN Global Technical Regulation No. 20: Electric Vehicle Safety (EVS)*, ECE/TRANS/180/Add.20, U.N. Econ. Commission Europe, Geneva, Switzerland, 2018.

[6] A. Sakurai, "Generic equations for a probabilistic metric for random hardware failures according to ISO 26262," in *Proc. Annu. Rel. Maintainability Symp. (RAMS)*, Palm Springs, CA, USA, 2020, pp. 1–6, doi: 10.1109/RAMS48030.2020.9153704.

[7] A. Sakurai, "Formulas of the probabilistic metric for random hardware failures to resolve a dilemma in ISO 26262," in *Proc. Annu. Rel. Maintainability Symp. (RAMS)*, Tucson, AZ, USA, 2022, doi: 10.1109/RAMS51457.2022.9893961.

[8] N. Das and W. Taylor, "Quantified fault tree techniques for calculating hardware fault metrics according to ISO 26262," in *Proc. IEEE Symp. Product Compliance Eng. (ISPCE)*, Anaheim, CA, USA, 2016, pp. 1–8, doi: 10.1109/ISPCE.2016.7492848.

[9] H. Jin, M. A. Lundteigen, and M. Rausand, "Reliability performance of safety instrumented systems: A common approach for both low- and high-demand mode of operation," *Rel. Eng. Syst. Saf.*, vol. 96, no. 3, pp. 365–373, Mar. 2011, doi: 10.1016/j.ress.2010.11.007.

[10] A. G. Phadke and J. S. Thorp, "Expose hidden failures to prevent cascading outages [in power systems]," *IEEE Comput. Appl. Power*, vol. 9, no. 3, pp. 20–23, Jul. 1996, doi: 10.1109/67.526849.

[11] L. Pintard, J.-C. Fabre, K. Kanoun, M. Leeman, and M. Roy, "Fault injection in the automotive standard ISO 26262: An initial approach," in *Dependable Computing (EWDC 2013)*, LNCS vol. 7869, Berlin, Germany: Springer, 2013, pp. 126–133, doi: 10.1007/978-3-642-38789-0_11.

[12] R. Almeida, V. Silva, and J. Cabral, "Virtualized fault injection framework for ISO 26262-compliant digital component hardware faults," *Electronics*, vol. 13, no. 14, art. 2787, Jul. 2024, doi: 10.3390/electronics13142787.

[13] Y. Sun, C. M. Poskitt, X. Zhang, and J. Sun, "REDriver: Runtime enforcement for autonomous vehicles," in *Proc. IEEE/ACM 46th Int. Conf. Softw. Eng. (ICSE)*, Lisbon, Portugal, 2024, doi: 10.1145/3597503.3639151.

[14] S. Gautham, A. Will, A. V. Jayakumar, and C. R. Elks, "STPA-driven multilevel runtime monitoring for in-time hazard detection," in *Computer Safety, Reliability, and Security (SAFECOMP 2022)*, LNCS vol. 13414, Cham, Switzerland: Springer, 2022, doi: 10.1007/978-3-031-14835-4_11.

[15] I. Aslam, N. Habib, A. Buragohain, M. Zhang, A. Rausch, V. Tiwari, and M. Benchat, "Connected dependability cage: Run-time function and anomaly monitoring for the development and operation of safe automated vehicles," arXiv:2604.27728, Apr. 2026.

[16] A. Eriksson and N. A. Stanton, "Takeover time in highly automated vehicles: Noncritical transitions to and from manual control," *Human Factors*, vol. 59, no. 4, pp. 689–705, Jun. 2017, doi: 10.1177/0018720816685832.

[17] X. Feng, M. Ouyang, X. Liu, L. Lu, Y. Xia, and X. He, "Thermal runaway mechanism of lithium ion battery for electric vehicles: A review," *Energy Storage Mater.*, vol. 10, pp. 246–267, Jan. 2018, doi: 10.1016/j.ensm.2017.05.013.

[18] M. Seo, T. Goh, M. Park, and S. W. Kim, "Detection method for soft internal short circuit in lithium-ion battery pack by extracting open circuit voltage of faulted cell," *Energies*, vol. 11, no. 7, art. 1669, Jun. 2018, doi: 10.3390/en11071669.

[19] S. Kang, H. Tu, and H. Fang, "BattBee: Equivalent circuit modeling and early detection of thermal runaway triggered by internal short circuits for lithium-ion batteries," arXiv:2506.13577, Jun. 2025.