JENRS

Received: 10 June 2026, Revised: 28 August 2026, Accepted: 30 August 2026, Online: 16 September 2026

DOI: https://doi.org/10.55708/js0509003

Fleet-Based Reliability Framework for Low-Voltage Electrical Subsystems in Autonomous Vehicles

Shrikant Pawar*

Independent Researcher, Mountain View, 94040, United States

*Corresponding author: Shrikant Pawar, California & Email: pawarshrikantsham@gmail.com

ABSTRACT: The low-voltage (LV) electrical subsystem of an autonomous vehicle (AV), spanning the auxiliary battery and management system, DC-DC conversion, power distribution, compute ECUs, in-vehicle communication networks, and harnesses, must operate continuously; its failures are a leading cause of in-service downtime even when propulsion is healthy. Yet no method connects fleet-derived reliability evidence to the assurance arguments AV safety cases require fleet-reliability methods are mature but not directed at AV LV electronics, while ISO 26262, ISO 21448, and UL 4600 specify what evidence a safety case needs, and ISO 26262 obliges field monitoring, none specifies how field evidence becomes an apportioned, sufficiency-justified assurance claim. **Objective.** This paper closes that gap with a direct bridge from field reliability estimates to standards-aligned assurance artefacts. **Methodology.** The bridge is an eight-layer, source-agnostic, continuous framework: component taxonomy; fleet data ingestion; duty-cycle-aware exposure modelling; reliability estimation reported in both per-hour (PMHF, MTBF, FIT) and per-mile (issues per million miles, IPMM) units; predictive health monitoring; safety-and-standards integration that apportions the item-level ISO 26262 target across the component register, classifies each fault, tests evidence sufficiency, and emits claim records rendered as Goal Structuring Notation fragments; containment; and governance. **Contributions.** (i) the first reliability framework scoped specifically to the AV LV electrical subsystem, including the communication architecture, connectors, and on-ECU power supply prior frameworks omit; (ii) a method translating field-derived failure intensities into safety-case evidence in both units, realized as a fourteen-field claim record, a PMHF apportionment, a fault-class transfer, an ES1-ES5 sufficiency criterion, and the conditions governing cross-operator comparison; and (iii) an AV-fleet containment ladder exploiting centralized dispatch and over-the-air channels. Because no public dataset exposes LV power-rail telemetry from a real AV fleet, validity is established by argument and illustrated with a worked example instantiated end to end for one safety goal; empirical validation is future work.

KEYWORDS: Autonomous vehicles, Fleet reliability, Low-voltage electrical subsystem, Functional safety, Assurance case

1. Introduction

1.1. Problem

Autonomous vehicles (AVs) depend on a continuously available low-voltage (LV) electrical subsystem that powers, networks, and monitors every safety-critical function on board. This subsystem comprises the 12 V auxiliary battery and its management system, the HV→LV DC-DC converter, the 48 V mild-hybrid rail where present, the LV power distribution network, the electronic control units (ECUs) that host the perception, planning, and control stack, the in-vehicle communication architecture (CAN, CAN-FD, Automotive Ethernet, LIN, FlexRay), and the wire harnesses, connectors, and grounds that physically interconnect them. Unlike a conventional vehicle, an AV cannot fall back on a human driver when any one of these components degrades. The LV subsystem must therefore operate continuously and degrade gracefully under nominal, off-nominal, and fault conditions.

Failures in this subsystem are a leading cause of in-service AV downtime. A 12 V auxiliary battery aged into low cold-cranking capacity, a DC-DC converter capacitor whose equivalent series resistance has drifted out of specification, an ECU whose decoupling network has degraded below its brown-out threshold, a CAN transceiver whose differential impedance has shifted under thermal cycling, a connector whose contact

resistance has risen from fretting corrosion: each can trigger an unplanned disengagement, mission abort, or service event even when the traction system is fully healthy. These failure modes are well documented in NHTSA Early Warning Reporting and recall data (https://www.nhtsa.gov/nhtsa-datasets-and-apis accessed May 2026) for conventional vehicles, but their incidence and severity under AV-specific duty cycles (continuous compute load, sustained sensor power draw, all-day operation in robotaxi service) are not characterized in the public literature.

Demonstrating that the LV subsystem of an AV meets its target failure rate is a statistical problem of unusual scale. Safety-critical AV functions are typically allocated an ASIL D probabilistic metric for random hardware failures (PMHF) on the order of 10^{-8} per hour [1], which corresponds to approximately 4×10^{-4} issues per million miles (IPMM) at the 25 mph reference speed used throughout this paper (an equivalence that holds only at that speed, for reasons developed in section 4.5) and translates to exposure on the order of 10^{8}–10^{9} vehicle-hours (equivalently 10^{9}–10^{10} vehicle-miles) before sufficient field evidence accumulates to confirm the target with reasonable confidence. No bench test, simulation campaign, or single-vehicle observation can produce that exposure. Only a fleet can. Reporting in both per-hour units (PMHF, MTBF, FIT) and per-mile units (IPMM) is required because the standards (ISO 26262) speak in per-hour terms while operators and regulators (Waymo, Cruise, California DMV) report in per-mile terms (California DMV, Autonomous Vehicle Disengagement Reports, https://www.dmv.ca.gov/portal/vehicle-industry-services/autonomous-vehicles/, accessed May 2026). A framework that produces both can address both audiences, but the units are not interchangeable: for a time-driven failure mechanism the per-mile figure inherits the fleet's duty cycle and is not comparable across operators without stated conditions. section 4.5 sets out those conditions and the conversion record the framework attaches to every claim. The contribution is that the conversion is explicit and auditable, not that one number serves every purpose.

1.2. The gap this paper addresses

The gap this paper addresses is precise: **no published framework ingests fleet-derived reliability evidence and emits it in the form an AV safety case can consume, for the LV electrical subsystem specifically**. Two mature bodies of work bracket this gap without closing it. On one side, fleet-reliability methods (mixed-Weibull regression, hierarchical Bayesian survival models, accelerated failure-time models [2], [3]) routinely estimate component-level failure intensities from warranty databases and operational telemetry, but they target conventional passenger vehicles and heavy-duty trucks and treat the LV side as one of several warranty cost centres. On the other, AV safety standards, namely ISO 26262 [1] (functional safety), ISO 21448 [4] (SOTIF), and UL 4600 [5] (goal-based assurance), define what evidence an assurance case requires, and ISO 26262 goes further, obliging a field monitoring process (Part 7) and admitting accumulated service experience as evidence (Part 8). What none of them specifies is how field-derived reliability evidence should be estimated, stratified, apportioned, and presented to satisfy the requirement, the step between the obligation and the artefact (section 2.2). Neither body targets the AV LV electrical subsystem: AV reliability research concentrates on perception, planning, and control [6], [7] rather than the electrical infrastructure that powers them, while the in-vehicle communication architecture (CAN, Automotive Ethernet, gateways) and the connectors, harnesses, and on-ECU power-supply circuitry that dominate NHTSA field data are treated as a first-class subsystem by neither community. The gap therefore falls exactly where fleet-reliability practice and AV assurance practice should meet, and that is what this framework closes.

1.3. Contribution

This paper presents an eight-layer fleet-based reliability framework specifically targeting the low-voltage electrical subsystem of autonomous vehicles. The framework is layered, source-agnostic, standards-aligned, reproducible, continuous, and AV-specific. Its layers are:

1. A scoped **component taxonomy** (Layer 1) covering energy storage and management, power distribution, ECUs, the in-vehicle communication architecture, harnesses, connectors, and sensor supply.
2. A source-agnostic **fleet data ingestion** model (Layer 2) that abstracts over internal OEM telemetry and public proxy data.
3. A duty-cycle-aware **data curation and exposure modelling** layer (Layer 3) for translating operational variability into reliability covariates.
4. A **reliability estimation** layer (Layer 4) combining statistical and physics-of-failure approaches to produce per-component failure intensities and FIT roll-ups.
5. A **predictive health monitoring** layer (Layer 5) for per-vehicle anomaly detection and remaining-useful-life estimation.
6. A **safety and standards integration** layer (Layer 6) that apportions the item-level ISO 26262 target across the Layer 1 register, classifies each fault against the architecture, applies the corresponding transfer function to the Layer 4 estimate, tests the result against an explicit evidence-sufficiency criterion, and emits a claim record renderable as an ISO 26262 / UL 4600-aligned GSN fragment, worked end to end for safety goal SG-LV-01 in section 3.7 and Figure 4. It reports in both the functional-safety per-hour unit and the regulatory per-mile unit, and attaches to every claim the conversion record and comparability conditions

that determine when a per-mile figure may be compared across operators (section 4.5).

7. An **AV-fleet containment and corrective-action** layer (Layer 7) that translates the framework's outputs into operator-controlled fleet actions (mission and ODD restriction, conservative-mode deployment, OTA mitigation, depot service, supplier 8D, regulator notification), exploiting the centralized dispatch and continuous OTA channels that distinguish AV fleets from passenger-vehicle deployments.
8. A **feedback, governance, and continuous-improvement** loop (Layer 8).

The paper is a methodological contribution. Its validity rests on three constructive arguments developed in section 4.1: completeness against the major reliability-evidence requirements of ISO 26262, ISO 21448, and UL 4600; compatibility with existing engineering artefacts such as component registers, FIT analyses, on-board diagnostics, and Goal Structuring Notation safety cases; and explicit closure of four gaps documented in section 2.6. Empirical validation on a real AV fleet dataset is deferred to future work, as no public dataset currently exposes LV power-rail telemetry from a real AV fleet, a constraint examined in section 4.3.

1.4. Paper structure

The remainder of the paper is organized as follows. section 2 surveys the five strands of literature on which the framework draws and identifies four gaps the framework targets. section 3 specifies the framework itself: core definition, motivation, subsystem scope, the eight-layer architecture with per-layer inputs and outputs, framework properties, and an illustrative scenario. section 4 establishes validity by argument, documents deployment requirements, discusses threats to validity, and compares the framework with existing practice. section 5 concludes with an adoption pathway for AV programmers suppliers, regulators, standards bodies, and researchers, and outlines future work.

2. Background & Related Work

The framework draws on five strands of literature. Each subsection summarizes the strand and identifies what the framework borrows or extends from it. The full reading list is preserved in the working notes at the end of this document.

2.1. Reliability engineering for vehicle fleets

Reliability of vehicle populations has been studied for decades through warranty-data analysis. Wu [2] reviews the field comprehensively, organizing methods around Weibull and mixed-Weibull distributions, mixture models that capture the bimodal infant-mortality and wear-out shapes typical of automotive electrical [8], and renewal-process treatments for repairable systems. Mixed-Weibull regression with covariates such as mileage-accrual rate, climate region, and use intensity was formalized by Attardi, Guida, and Pulcini [3] and remains the workhorse for fleet-wide failure-intensity estimation. Where data is aggregated rather than per-VIN (as is typical when only public reports are available), Wu [9] reviews coarse-data techniques that preserve the inferential power needed for population estimates. More recent work fuses telematics with warranty records: in [10], [11], the authors demonstrate end-to-end pipelines that combine Logged Vehicle Data with claim ground truth, using stacked ensembles for early failure prediction across heavy-duty truck fleets. The framework's Layer 4 builds directly on this lineage (mixed-Weibull regression for parametric reliability, hierarchical Bayesian survival models for partial pooling across model years, and accelerated-failure-time models for incorporating duty-cycle covariates), and Layer 3 extends them with the per-mile and per-hour paired exposure covariates that AV operations require.

2.1.1. Primer: Weibull analysis as the framework's parametric backbone

The Weibull distribution is the workhorse of automotive reliability analysis because its two parameters (a **shape parameter β** and a **scale parameter η**) together describe the three failure regimes that real automotive electrical components exhibit. Its probability density is given by Equation (1):

$$f(t)=\frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1}\exp\left[-\left(\frac{t}{\eta}\right)^{\beta}\right] \quad (1)$$

A shape parameter **β < 1** reproduces early-life infant-mortality failures driven by manufacturing defects, **β = 1** reduces to a constant-rate (exponential) regime characteristic of random failures, and **β > 1** captures wear-out failures whose hazard rate grows with age. The scale parameter η is the characteristic life, the time at which approximately 63.2 % of the population has failed. Real fleet failure data is rarely monotone in any one of these regimes; it is more often bimodal, with an infant-mortality population superposed on a wear-out population. A **mixed-Weibull** distribution (a weighted sum of two or more Weibull components) captures this shape and is the standard parametric model for automotive warranty data [8]. When per-vehicle covariates are available, mixed-**Weibull regression** substitutes a covariate-dependent scale parameter

$$\eta(x)=\eta_0\exp(-\beta x) \quad (2)$$

so that duty cycle, climate exposure, mileage accrual, and other operational variables modulate the characteristic life [3]. The framework's Layer 4 uses this form as its default parametric backbone, with hierarchical Bayesian and accelerated-failure-time extensions where data sparsity or covariate complexity warrants.

2.2. Automotive functional safety

Automotive functional safety is governed primarily by ISO 26262 [1], which defines a V-model lifecycle for safety-

critical electrical and electronic systems and a Hazard Analysis and Risk Assessment (HARA) procedure for assigning Automotive Safety Integrity Levels (ASIL-A through ASIL-D) to safety goals. For random hardware failures, the standard specifies quantitative targets: the Probabilistic Metric for random Hardware Failures (PMHF) for the safety goal, and the Single-Point Fault Metric and Latent Fault Metric (SPFM, LFM), which express the fraction of an item's total hardware failure rate that is either safe or covered by diagnostics. All three are **item-level** metrics, and they are item-level in two different ways. PMHF is evaluated for the item realizing a safety goal and aggregates that item's residual and multiple-point fault contributions; SPFM and LFM are *ratios* taken over the item's entire safety-related failure-rate population. A component therefore does not have a PMHF, an SPFM, or an LFM of its own: it has a *contribution* to the item's PMHF, it has no SPFM or LFM at all, and the per-fault diagnostic coverage that scales its contribution is a distinct quantity from the item's SPFM. This distinction is what makes an apportionment step necessary before field-derived per-component rates can be compared to anything the standard recognizes; section 3.7.2 supplies the apportionment and section 3.7.4 the aggregation back to item level. Compliance demonstration relies in part on field-derived failure rates, and the standard is correspondingly explicit about the field: it requires a planned process for monitoring the field during operation and service. What it does not specify is how field observations become the rates that compliance demonstration consumes. That distinction is developed at the end of this section, and it is where this framework integrates. In [12], the authors survey ISO 26262 application in AV contexts. In [13], the authors provides a landmark reliability analysis of a fail-operational brake-by-wire system with explicit LV power-architecture implications. ISO 21448 (Safety of the Intended Functionality, or SOTIF) [4] complements ISO 26262 by addressing functional insufficiencies and reasonably foreseeable misuse rather than random hardware failure [14], [15]. Its object in the operation phase is the functional insufficiency and the triggering condition rather than the random hardware failure, so while it too directs attention to the field, it does not bear on the estimation of component failure intensities that concerns this framework. The framework's Layer 1 produces the component register and ASIL annotations that feed ISO 26262 HARA; Layer 6 emits the field-evidence artefacts that close the random-hardware-failure portion of the safety case.

2.2.1. *Where ISO 26262 already contemplates field evidence, and where it stops*

ISO 26262 [1] does not treat in-service data as out of scope. It contemplates field evidence in two distinct places, and a framework proposing to supply field evidence to a safety case has to be precise about which of them it extends.

Part 7 requires a field monitoring process: the organization must collect and evaluate field information relevant to functional safety, define criteria by which a safety anomaly is recognized, and define the reaction when one is. This establishes a *duty to observe and respond*. It does not prescribe how the observation becomes a number. No exposure model is specified, no treatment of censoring, no covariate structure, no uncertainty quantification, and no rule for deciding when accumulated evidence is enough to support a claim rather than merely to trigger one. Part 7 tells an organization that it must notice a safety anomaly; it does not tell it how to estimate a failure intensity from a fleet.

Part 8 provides the proven-in-use route, under which the service history of a candidate element (accumulated operating hours and observed failures) may substitute for parts of the development-process evidence, subject to statistical requirements on the exposure achieved. This is genuinely quantitative use of field data within the standard, and it is the nearest existing relative of the present framework. But it answers a narrower question. Proven-in-use asks whether an element's service history is sufficient to stand in for development rigour; it does not ask what that element's contribution to the item's random-hardware-failure metric is, nor on what basis that contribution may be asserted as a claim in an assurance case. Its statistical requirements bound a substitution decision, not an apportioned per-element budget.

What neither part supplies (and what UL 4600 [5] requires but likewise does not specify) is the method between them: how a fleet-derived failure intensity is apportioned against an item-level target, classified by fault type against the architecture, tested against an explicit sufficiency criterion, and rendered as a claim with its rebuttals recorded. That method is this paper's Layer 6, and section 3.7 works it through end to end. The framework therefore does not compete with Part 7 or Part 8; it sits between the obligation they establish and the artefact UL 4600 demands, and its evidence-sufficiency criterion (section 3.7.5) can be read as a generalization of the proven-in-use exposure requirement from a binary substitution test to a per-claim sufficiency test.

2.3. *AV-specific safety frameworks*

Beyond ISO 26262 and SOTIF, autonomous products are increasingly assessed against UL 4600 [5], which prescribes a goal-based safety case typically expressed in Goal Structuring Notation (GSN). UL 4600 is explicit that the safety case must rest on evidence, including in-service field evidence, and that evidence sufficiency must be justified per claim. The standard does not, however, specify how that evidence is to be produced, structured, or maintained over time; this is where a fleet-based reliability framework adds value. SAE J3016 [16] scopes the automation levels (L0 through L5) to which a given safety case applies and is cited throughout the AV literature to bound system claims. In [6], the authors argues that AV

safety arguments must integrate ISO 26262, SOTIF, security, and behavioural-safety layers. In [7], the author survey risk-assessment methodologies aligned with these standards. Cross-domain GSN work in the rail-safety literature [17] provides transferable patterns for AV cases. In [18], the authors extend assurance arguments to machine-learning components specifically. The framework's Layer 6 produces GSN fragments and UL 4600 claim records that consume Layer 4's reliability estimates and document evidence sufficiency in both PMHF and IPMM terms.

2.4. Fleet data analytics and PHM in automotive

Predictive maintenance and Prognostics and Health Management (PHM) for automotive electronics have matured rapidly. In [19], [20] and [21], the authors review the field and extend it with thermomechanical sensor and machine-learning hybrids for ECU package prognosis. In [22], the researchers argue that fleet-derived prognostics can materially reduce reliability test duration, an argument the present framework operationalizes. For batteries specifically, In [23], a review on deep-learning approaches to state-of-charge, state-of-health, and remaining-useful-life estimation, while [24] demonstrate field-data-only end-of-life prediction methodology transferable to 12 V auxiliary battery fleets. For DC-DC and converter switching devices, IGBT precursor-parameter work [25], [26] and capacitor RUL methods [27] provide the mechanism-level signals Layer 5 monitors. CAN bus anomaly detection, although framed as security in much of the literature [28], [29], uses the same statistical and machine-learning techniques that PHM applies to fault diagnosis. In [30], the authors provide a recent overview emphasizing trustworthiness and explainability, qualities the framework's Layer 6 inherits when its outputs flow into the safety case. The framework's Layer 2 ingests the telemetry these methods require, Layer 3 prepares it, and Layer 5 applies anomaly detection and remaining-useful-life estimation at per-vehicle granularity.

2.5. LV electrical architecture in AVs

The architecture of vehicle low-voltage electrical subsystems is in active evolution. Zonal E/E architectures, in which a small number of zone controllers replace dozens of distributed ECUs, change how LV power and redundancy are physically distributed [31]. DC-DC converter topologies that tolerate single faults (multilevel, parallel-protection, and redundant designs) are surveyed by [32]. In [33] and [34], the researchers formalize mission-profile-driven reliability and the physics-of-failure approach for power electronics, which is directly applicable to LV DC-DC stages, 48 V mild-hybrid rails, and the on-ECU power-supply circuitry that the framework places in scope. In [35], the authors describe an open-access 48 V automotive architecture. In [36], the authors demonstrate hybrid battery-and-supercapacitor energy-buffer designs translatable to AV limp-home scenarios. Battery management is reviewed by [37], spanning ISO 26262 and IEC 61508 angles. Next-generation switching devices (GaN, SiC) and their failure modes are covered by [38]. FinFET aging in automotive memories (relevant to the SoC and MCU contribution to the LV-subsystem FIT budget) is treated by [39]. In [40], the authors provide a rare published compendium of failure rates for AV-relevant hardware (sensors, ECUs, communications), anchoring the bench-test FIT priors that Layer 4 cross-checks. Together, this body of work supplies the component-level reliability physics that Layer 4 uses as priors and the architectural awareness that Layer 1 uses to bound the subsystem scope for AV-specific deployments.

2.6. Gaps the framework targets

Synthesizing across section 2.1–section 2.5, four gaps are evident:

1. **The method joining fleet reliability estimates to AV assurance claims is unspecified**. Reliability methods produce $\lambda(t)$ and FIT; assurance methods consume claims. ISO 26262 establishes both the obligation to monitor the field (Part 7) and a route by which accumulated service experience may serve as evidence (Part 8, proven in use), and UL 4600 requires that evidence sufficiency be justified for every claim. None of them specifies the translation itself: how an estimated failure intensity is apportioned against an item target, classified by fault type, tested for sufficiency, and rendered as a claim. That method is what is missing, and section 3.7 supplies it.
2. **AV-LV electrical reliability is markedly under-studied** compared to traction-side power electronics, despite being a dominant cause of in-service AV downtime.
3. **In-vehicle communication architecture reliability** (CAN, Automotive Ethernet, gateways) is rarely treated as a first-class subsystem in fleet reliability literature.
4. **Connectors, harnesses, and ECU power-supply circuitry** are well-known failure modes in NHTSA data but absent from most published reliability frameworks.

3. The Framework

3.1. Core definition

A reliability framework is a structured, systematic approach that organizations use to assess, predict, improve, and manage the reliability of components, subsystems, or entire systems throughout their lifecycle. Reliability itself is the probability that a product performs its required functions under stated conditions for a specified period of time without failure. A reliability framework provides the overarching structure (processes, analytical tools, statistical models, testing methods, data management practices, and governance) needed to

achieve high reliability in a repeatable and measurable way.

A **fleet-based** reliability framework shifts emphasis from lab-only or simulation-only methods to **real-world operational data collected from a large population of deployed vehicles**. The framework continuously updates reliability estimates from in-service evidence rather than relying on bench-test extrapolation alone.

3.2. Why a fleet-based approach is essential for LV electrical subsystems in AVs

- **Statistical demonstration is impractical without fleets.** Demonstrating extremely low failure rates for safety-critical AV functions can require exposure on the order of 10^8–10^9 vehicle-hours; only a fleet provides that exposure.
- **LV failures cause real AV downtime even when propulsion is healthy.** 12 V battery sag, DC-DC converter dropouts, ECU brown-outs, CAN bus errors, harness intermittents, and ground-network corrosion are leading causes of disengagements, mission aborts, and unplanned service events.
- **AVs generate dense telemetry.** Bus voltages, currents, temperatures, DTCs, wake/sleep events, accessory loads, network error counters, and compute power draws are observable across an entire fleet at high cadence.
- **Real exposure captures variability lab testing under-represents**: continuous perception/compute load, sustained accessory draw, thermal cycling under all-day operation, software-version-dependent power profiles.
- **Reliability targets are continuously moving** as software, hardware, and ODD evolve; only fleet data enables continuous re-estimation.

3.2.1. Primer: Fleet exposure as the foundational quantity

Underlying every quantitative argument in this paper is the concept of **fleet exposure**: the total operating time, or mileage, accumulated across all vehicles in the fleet, summed and stratified by relevant covariates. If a fleet of N vehicles accumulates E vehicle-hours (or vehicle-miles) of operation over an observation window and k failures of a particular component are recorded, the empirical failure-rate estimate is simply

$$\hat{\lambda} = \frac{k}{E} \tag{3}$$

with uncertainty bounded by Poisson or survival statistics. Exposure is therefore the denominator that makes every reliability metric in this paper meaningful; without it, neither PMHF, MTBF, FIT, nor IPMM has any quantitative content.

A fleet enables three things that single-vehicle or bench testing cannot. First, it accumulates the **volume** of exposure required: confirming an ASIL-D PMHF target of 10^{-8}/hr to reasonable confidence requires on the order of 10^8–10^9 vehicle-hours, equivalently 10^9–10^{10} vehicle-miles. Second, it sources the **heterogeneity** of exposure (different duty cycles, climates, software revisions, aging trajectories) that lets Layer 4's mixed-Weibull regression attribute failure-rate variation to specific covariates rather than absorbing it as noise. Third, it generates exposure **continuously**, so that reliability estimates can be updated as software, hardware, and operating domain evolve. Layer 3 of the framework treats exposure as a first-class quantity, tracked in both vehicle-hours (for PMHF reporting) and vehicle-miles (for IPMM reporting), and stratified by the duty-cycle covariates that Layer 4 consumes.

3.3. Subsystem scope: what counts as "LV electrical subsystem" in an AV

Components **in scope**, grouped by functional role:

Energy storage, management and conversion

- 12 V auxiliary (SLI / LiFePO4) battery
- 12 V battery management system (BMS): state-of-charge / state-of-health estimation, cell balancing, fault detection, isolation monitoring
- HV→LV DC-DC converter (primary and any redundant unit)
- 48 V mild-hybrid rail where present, including its DC-DC stage
- Redundant LV energy buffer (supercapacitor / secondary battery) for fail-operational AV functions

Power distribution and protection

- Smart fuse boxes, smart power switches, eFuses
- LV bus bars, body harness, ground network
- Redundant LV supplies for safety-critical loads: brake-by-wire, electric power steering, redundant comms

Compute and control electronics (ECUs)

- AV compute ECUs: perception, planning, control hardware (Nvidia DRIVE, Mobileye EyeQ, OEM custom SoCs, AI accelerators)
- Domain and zone controllers in zonal E/E architectures
- Body control module (BCM), gateway ECU, wake/sleep controller, telematics control unit
- On-ECU power supply circuitry (LDOs, switching regulators, decoupling capacitors, brown-out detectors)
- Sensor ECUs (LiDAR, radar, camera, GNSS, IMU front-ends)

In-vehicle communication architecture

- CAN / CAN-FD: twisted-pair wiring, transceivers, terminations, common-mode chokes
- Automotive Ethernet (100BASE-T1 / 1000BASE-T1 / 10GBASE-T1): PHYs, switches, time-sensitive networking (TSN) hardware
- LIN, FlexRay where present
- Gateway and routing ECUs, bus arbitration hardware
- Bus terminations, ESD protection, EMC filtering

Wire harnesses, interconnects, and grounds

- Main LV harness and branch harnesses
- ECU-to-ECU links, sensor-to-compute links, power-over-data-line (PoDL) runs
- Connectors, splices, ring terminals, header pins
- Ground network and chassis bonding
- Conduits, grommets, strain reliefs

Sensor power and signalling

- Sensor supply rails (LiDAR, radar, camera, GNSS, IMU, ultrasonic)
- PoE / PoDL where used to power sensors over the data link

Components **out of scope** (kept explicit for boundary audit):

- HV traction pack, traction inverter, traction motor
- Mechanical actuators (only their LV supply path is in scope)
- Software functionality of ECUs (only the hardware reliability is in scope; software faults are an ISO 21448 / SOTIF concern, not ISO 26262 random-hardware-failure)

3.4. Framework architecture: eight layers

The framework comprises eight layers. Each layer has defined inputs and outputs, can be developed independently, and can be instantiated against either internal OEM telemetry or publicly available proxy data.

Figure 1 shows the layered architecture with the principal inputs and outputs of each layer. Figure 2 shows the same eight layers as an end-to-end data flow, from external evidence sources through estimation and assurance to containment, with the Layer 8 continuous-improvement loop that feeds field evidence back into design, supplier, and scoping decisions.

Eight-layer fleet-based reliability framework

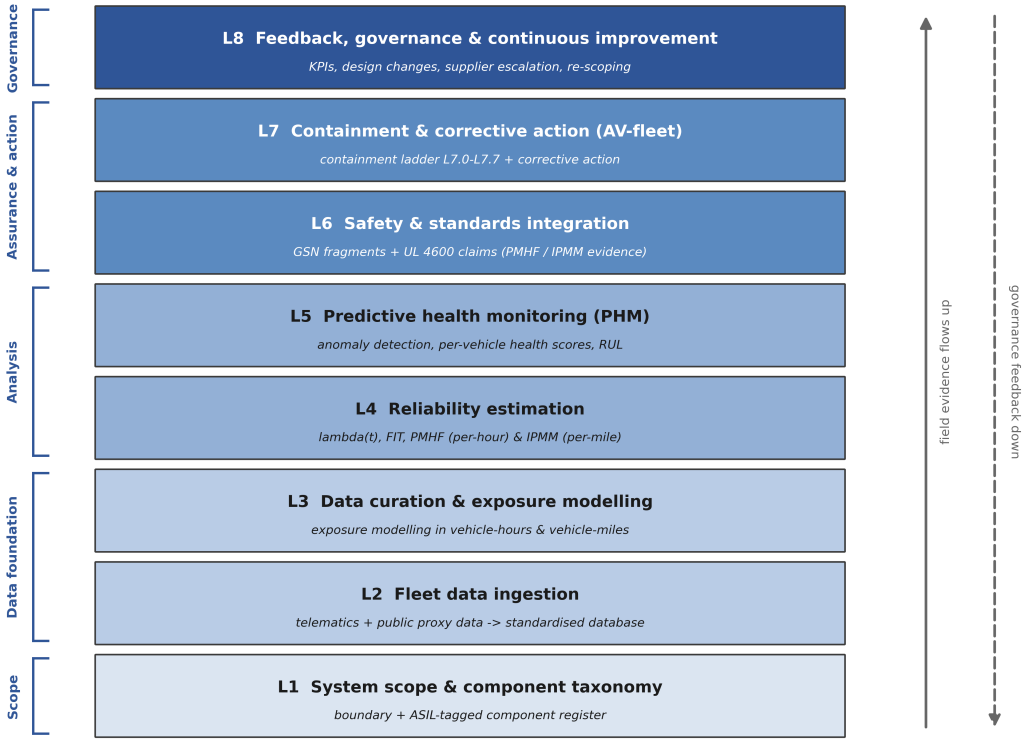

Figure 1: The eight-layer architecture. Layer 1 (component scope) is the foundation; field evidence flows upward through ingestion, curation, estimation, and predictive health monitoring into the safety-and-standards, containment, and governance layers, while governance feedback (Layer 8) flows back down. Each layer's principal inputs (left) and outputs (right) are shown; every layer can be developed independently and instantiated against either internal OEM telemetry or public proxy data.

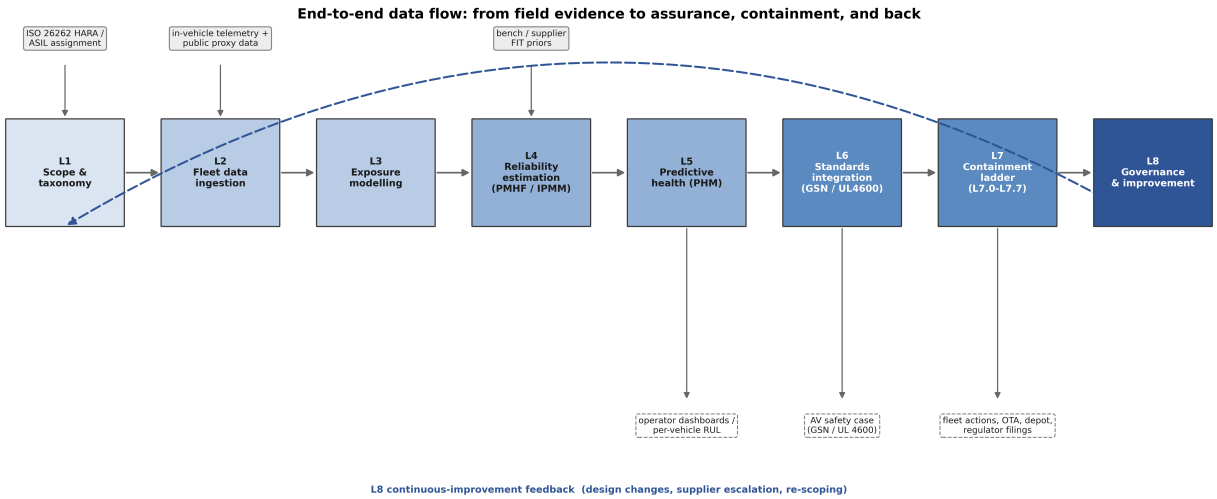

Figure 2: End-to-end data flow. External sources (ISO 26262 HARA/ASIL assignments, in-vehicle telemetry and public proxy data, bench/supplier FIT priors) enter the pipeline at the layers that consume them; the layers produce the operator-facing (dashboards, RUL), assurance-facing (GSN / UL 4600 safety case), and operational (fleet actions, OTA, depot, regulator filings) outputs; and Layer 8 closes the loop by returning continuous-improvement evidence (design changes, supplier escalation, and re-scoping) to Layer 1.

3.4.1. *Layer 1: System scope & component taxonomy*

- **Purpose:** establish the boundary of the LV electrical subsystem under analysis and define a component hierarchy.
- **Activities:** enumerate in-scope and out-of-scope components per section 3.3; build the system → subsystem → component → part hierarchy; assign each component a safety classification (ASIL-A through D) and a functional role (primary, redundant, monitor).
- **Inputs:** vehicle architecture documentation, ISO 26262 HARA outputs.
- **Outputs:** component register with safety attributes and interface specifications.

3.4.2. *Layer 2: Fleet data ingestion*

- **Purpose:** acquire and consolidate the evidence streams the framework operates on.
- **Activities:** define data contracts (schema, sampling rate, units, timestamp convention); ingest in-vehicle telematics (LV bus voltages, currents, temperatures, DTCs, wake/sleep events, CAN/Ethernet error counters); ingest external sources (NHTSA EWR / Recalls / VOQ for outcomes; FleetDNA for operational reference; vPIC for VIN decoding; state registration data for population denominators).
- **Inputs:** raw telematics, public dataset feeds, OEM service / warranty records.
- **Outputs:** standardized event + signal database, schema-validated.

3.4.3. *Layer 3: Data curation & operational exposure modelling*

- **Purpose:** turn raw data into a clean exposure-annotated dataset suitable for reliability inference.
- **Activities:** cleaning, deduplication, missing-data handling; population denominator estimation (vehicles-in-operation, accumulated vehicle-hours, and accumulated vehicle-miles by year/model); duty-cycle characterization (per-vehicle miles, idle hours, stop-start cycles, accessory load distribution, ambient temperature exposure, voltage time-at-low-SoC); construction of an **AV-equivalent duty-cycle envelope** when proxy data is used. Both per-hour and per-mile exposure are tracked as first-class quantities so that downstream reporting can be expressed in PMHF (per-hour) and IPMM (per-mile) units.
- **Inputs:** Layer 2 standardized database.
- **Outputs:** clean dataset with per-component exposure covariates including paired per-hour and per-mile exposure totals; and the fleet duty-cycle summary on which per-mile reporting depends: fleet-average speed over powered hours $\bar{v}$, idle fraction, and cycles per mile, each with its provenance and the operating/parked boundary used to compute it (section 4.5.3, CC2 and CC3).

3.4.4. *Layer 4: Reliability estimation*

- **Purpose:** convert exposure-annotated event data into quantitative reliability metrics expressed in both per-hour and per-mile units.
- **Activities:** statistical modelling (mixed-Weibull regression, hierarchical Bayesian survival models, accelerated-failure-time models with duty-cycle covariates); physics-of-failure models where mechanism is known (lead-acid sulfation, capacitor wear, solder-joint fatigue, connector corrosion and fretting, MOSFET aging); component-level FIT roll-up to subsystem level; cross-check against bench-test FIT (SN 29500 [41] / IEC 62380 [42] / Richter et

al. [40]); parallel estimation in per-mile units via the Layer 3 per-mile exposure totals.

- **Inputs:** Layer 3 clean dataset (with paired per-hour and per-mile exposure), component taxonomy from Layer 1.
- **Outputs:** per-component failure intensity $\lambda(t \mid \text{duty cycle})$ with credible intervals, expressed in both **per-hour units** ($\lambda(t)$, MTBF, FIT) and **per-mile units** (issues per million miles, IPMM); per-component FIT contributions, aggregated to item level for comparison against the apportioned ISO 26262 targets, with the apportionment itself being a Layer 6 activity (section 3.7.2), since it depends on the item definition and the safety goal rather than on the reliability estimate; and a subsystem-level IPMM rate for fleet-operator and regulator reporting. Each estimate additionally carries the **native denominator** of its dominant failure mechanism (time, distance, or cycles) as determined by the physics-of-failure model, since that determines which unit is the invariant and which is derived (section 4.5.2), and whether a per-mile figure is a comparable quantity at all (section 4.5.3, CC1).

3.4.5. *Layer 5: Predictive health monitoring (PHM)*

- **Purpose:** translate aggregate reliability estimates into per-vehicle, real-time health awareness.
- **Activities:** anomaly detection on LV signals (voltage drift, current spikes, ripple, ground-shift, CAN bus error-frame rate, Ethernet link quality / FEC corrections, ECU brown-out event rate); precursor feature extraction (rolling variance, spectral content, drift indicators); remaining useful life (RUL) estimation for monitored components (most directly the 12 V auxiliary battery, DC-DC converter capacitors and switching devices, network PHYs, and high-cycle connectors); integration with vehicle on-board diagnostics for early-warning flags.
- **Inputs:** Layer 4 reliability priors, real-time Layer 2 signal streams.
- **Outputs:** per-vehicle health scores, RUL projections, anomaly alerts.

3.4.6. *Layer 6: Safety & standards integration*

- **Purpose:** make the framework's reliability outputs directly consumable by the safety case for the AV and by fleet-operator and regulator reporting.
- **Activities:** apportion the item-level ISO 26262 PMHF target across the Layer 1 element register, recording the apportionment rule and holding an explicit reserve (section 3.7.2); classify each fault as single-point, residual, or multiple-point against the Layer 1 architecture, and apply the corresponding transfer function to convert the Layer 4 rate into a contribution to the item metric (section 3.7.3); compare each contribution to its apportioned share rather than to the item target; compute the item's SPFM and LFM from the aggregated fault populations, keeping these architectural metrics distinct from the per-fault diagnostic coverage that scales an individual component's residual; convert between PMHF (per-hour, the standards' native unit) and IPMM (per-mile, the operator and regulator unit) using fleet-average speed over powered hours from Layer 3, and emit a **conversion record** with every claim carrying the per-hour rate, the operational IPMM, the normalized IPMM at the declared reference speed, $\bar{v}$ and its provenance, the mechanism's native denominator from the Layer 4 physics-of-failure model, and the outcome of comparability conditions CC1–CC4 (section 4.5.3); generate Goal Structuring Notation (GSN) fragments and UL 4600-aligned claim structures backed by fleet evidence in both per-hour and per-mile units; document evidence sufficiency for each assurance claim (sample size, exposure equivalence in vehicle-hours and vehicle-miles, residual uncertainty); record assumptions and proxy substitutions.
- **Inputs:** Layer 4 reliability outputs (per-hour and per-mile), Layer 1 ASIL assignments, Layer 3 fleet-average-speed assumption.
- **Outputs:** standards-aligned assurance artefacts: GSN fragments, UL 4600 claim records, evidence-sufficiency declarations, and an explicit PMHF↔IPMM conversion record per claim, carrying the per-hour rate, the operational IPMM, the normalized IPMM at the declared reference speed, $\bar{v}$ and its provenance, the mechanism's native denominator, and the outcome of comparability conditions CC1–CC4 (section 4.5).

3.4.7. *Layer 7: Containment & corrective action (AV-fleet specific)*

- **Purpose:** translate the framework's reliability and safety outputs into immediate operator-controlled fleet actions that bound exposure and protect customers while a permanent fix is engineered. AV fleets differ structurally from passenger-vehicle deployments (operator-owned, centrally dispatched, fully observable, daily depot return, continuous OTA channels), and this layer exploits those properties to enable containment that is faster and finer-grained than traditional OEM recall workflows allow.
- **Activities:** bound the affected-vehicle population using Layer 5 per-vehicle risk scores stratified by Layer 4 covariates (software version, build date, supplier lot, mission profile); a Containment Review Board (cross-functional, drawing from engineering, quality, regulatory affairs, operations,

communications, and legal) selects the containment level appropriate to severity, confidence in root cause, mitigation availability, and operational impact. Containment levels available to an AV operator are:

- **L7.0 Soft monitor**: tag affected vehicles in the per-vehicle risk register; no operational change.
- **L7.1 Mission restriction**: withhold specific mission types (highway, airport runs, long shifts) from affected vehicles.
- **L7.2 ODD restriction**: tighten the operating design domain for affected vehicles (no rain, no night, specific geofences only).
- **L7.3 Conservative operating mode**: deploy a derated profile (lower speed, longer following distance, reduced accessory load) across the affected subset.
- **L7.4 OTA mitigation**: deploy a parameter or software change that addresses or works around the failure across the affected subset within hours.
- **L7.5 Pull from service**: route affected vehicles to the depot; no further missions assigned.
- **L7.6 Depot service action**: physical inspection, part swap, recalibration, or harness rework, typically overnight.
- **L7.7 Fleet ground**: pause the entire fleet or a major subset; suspend service in affected markets.

Beyond containment-level selection, this layer coordinates **supplier containment** (incoming-part inspection, quarantine of replacement stock, 8D activities at the supplier), **regulator notification** (NHTSA Standing General Order 2021-01 AV incident reporting, NHTSA Part 573 defect notification where applicable, CA DMV permit-related disclosures, EU type-approval authorities where applicable), **public and rider communication**, and **evidence preservation** (sealed-evidence procedures for failed parts, data preservation around failure events, chain of custody for safety-case and potential litigation).

- **Inputs:** Layer 4 fleet-level failure rate, Layer 5 per-vehicle risk scores, Layer 6 severity classification and assurance-case context, Layer 1 component scope; external context including regulatory thresholds, business posture, supplier capacity, and depot capacity.
- **Outputs:** containment-action record (audit trail of who decided what when, with rationale); updated per-vehicle risk register reflecting the action taken; communication artefacts (regulator filings, rider notices, supplier 8D records, internal customer-success notifications); failed-part collection log; updates to the Layer 6 assurance case reflecting the containment evidence; lessons-learned input to Layer 8 for governance updates.

3.4.8. *Layer 8: Feedback, governance & continuous improvement*

- **Purpose:** close the loop between fleet evidence and engineering, supplier, and operational decisions; ensure that the framework operates as a continuous lifecycle process rather than a one-shot exercise.
- **Activities:** KPI dashboard (per-component λ and IPMM trends vs. target, top-N failure modes ranked by IPMM, RUL distributions, containment-action frequency, time-to-permanent-fix); decision gates (design-change triggers, supplier escalations, OTA roll-out criteria, recall thresholds, defined in normalized IPMM at the governance reference speed where the audience is operators or regulators, and in PMHF where the audience is functional-safety reviewers; a gate set in operational IPMM silently moves whenever the fleet's mission mix shifts, and must be restated when it does); confidence-bound reporting on every metric; explicit data-quality scoring; documented review cadence; lessons-learned propagation across components and platforms.
- **Inputs:** outputs from Layers 4–7, engineering and operations team requirements, Layer 7 containment-action records.
- **Outputs:** reliability improvement workflow, governance policies, audit trail, design-change requests, supplier-development priorities.

3.5. *Framework properties*

- **Layered & modular**: each layer has defined inputs/outputs; teams can adopt incrementally rather than as a single waterfall deployment.
- **Source-agnostic**: works on internal OEM telemetry or on public proxy data; Layer 2 abstracts the source.
- **Standards-aligned**: Layer 6 produces artefacts directly consumable by ISO 26262 and UL 4600 safety cases.
- **Reproducible & auditable**: every transformation in Layers 2–4 is logged; assumptions in Layer 3's duty-cycle proxy are tracked.
- **Continuous**: designed for lifecycle use, not a one-shot reliability assessment.
- **AV-specific**: Layer 1 scope and Layer 3 duty-cycle envelopes are tailored to LV components and AV operational profiles; the framework is not a generic warranty-analysis pipeline.

3.6. *Illustrative scenario*

The following short sketch shows how the framework responds to a representative LV electrical condition. It is illustrative only and is not an empirical case study. To make the Layer 4 output concrete, the scenario carries

worked quantitative values; these are *illustrative figures anchored to the published priors cited inline* (in the manner of the illustrative FIT and λ budgets used by Sinha [13] and Mariani [6]), and are presented as a **population-growth curve over fleet exposure** (expected cumulative failures versus accumulated vehicle-hours and vehicle-miles) rather than as measurements of any real fleet. The values use a single illustrative fleet of N = 5 000 robotaxis operating 10 h/day at a Layer 3 fleet-average speed of 25 mph, the speed implied by the PMHF↔IPMM equivalence already stated in section 1.1 (ASIL-D 10^{-8}/h ≈ 4×10^{-4} IPMM), observed over a 24-month window.

3.6.1. *Scenario: Apparent 12 V auxiliary battery state-of-health drift caused by a leaky balancing MOSFET in the BMS*

A subset of the robotaxi fleet exhibits a steady downward drift in BMS-reported state-of-health and a slow rise in quiescent current draw. The dominant root cause turns out not to be battery aging itself but elevated drain-source leakage in a cell-balancing MOSFET on the BMS board, a known wear mechanism driven by gate-oxide degradation and time-dependent dielectric breakdown (TDDB) under repeated switching and elevated junction temperature. The leakage produces a parasitic discharge path that the BMS partially compensates for, while the cell-voltage spread it induces is reported as SoH degradation. The framework's value in this scenario is to attribute the failure to the BMS MOSFET rather than the battery, so that the corrective action and supplier accountability land in the right place.

- **L2 ingestion:** quiescent voltage at wake, quiescent current at sleep, cell-voltage spread reported by the BMS, balancing-FET ON-time and switching-cycle count, BMS board temperature where instrumented, cold-crank voltage at start, ambient temperature, and BMS-reported SoH per vehicle per day.
- **L3 curation:** clean and deduplicate; compute per-vehicle exposure covariates (daily km, idle hours, cumulative balancing-FET switching cycles, BMS-PCB time-at-high-temperature, cumulative quiescent-current integral that bounds parasitic discharge). Per-vehicle exposure is tracked in both vehicle-hours and vehicle-miles so that downstream estimates can be expressed in both PMHF and IPMM. At 24 months each vehicle has accrued 7 300 vehicle-hours / 182 500 vehicle-miles, for a cumulative fleet exposure of 36.5 M vehicle-hours / 912.5 M vehicle-miles.
- **L4 estimation:** fit Weibull regression for time-to-fault on "BMS balancing MOSFET leakage above threshold" with covariates of cumulative switching cycles, junction-temperature exposure, and gate-drive duty; anchor the prior on bench-test FIT for the specific MOSFET part and on published TDDB / gate-oxide-aging physics-of-failure models. The gate-oxide TDDB mechanism is wear-out dominated, so its time-to-failure follows a two-parameter Weibull with shape **$\beta > 1$**; published gate-oxide TDDB characterizations place the Weibull slope in the ≈ 1–4 range [43], with field acceleration on the thermochemical E-model [44]. This scenario uses an *illustrative* **$\beta = 2.0$** and characteristic life **η = 150 000 operating hours**. The unreliability $F(t)$ and hazard $\lambda(t)$ functions are given by Equations (4) and (5):

$$F(t) = 1 - \exp\left[-\left(\frac{t}{\eta}\right)^{\beta}\right] \tag{4}$$

$$\lambda(t) = \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1} \tag{5}$$

from which Layer 4 derives the illustrative quantities in Table 1, because the B_{10} life lies far beyond nominal service life, only the early tail of the wear-out population manifests in service. The Layer 4 headline output is the **population-growth curve**, expected cumulative fleet failures

$$E[\text{failures}] = N \cdot F(t) \tag{6}$$

evaluated at the per-vehicle exposure reached by each elapsed time, as shown in Table 2.

Table 1: Weibull-derived reliability quantities for the illustrative BMS balancing-MOSFET gate-oxide TDDB failure mode (β = 2.0, η = 150 000 h).

Quantity	Expression	Illustrative value
Mean time to failure	$\eta\cdot\Gamma(1 + 1/\beta)$	≈ 1.33×10^{5} h (≈ 36 yr at 10 h/day)
B_{10} life (10 % failed)	$\eta\cdot(-\ln 0.9)^{1/\beta}$	≈ 4.9×10^{4} h (≈ 13 yr at 10 h/day)
Hazard at 12 months (t = 3 650 h)	$(\beta/\eta)(t/\eta)^{\beta-1}$	≈ 3.2×10^{-7}/h (324 FIT)
Hazard at 24 months (t = 7 300 h)	$(\beta/\eta)(t/\eta)^{\beta-1}$	≈ 6.5×10^{-7}/h (649 FIT)

Table 2: Expected cumulative BMS-MOSFET failures and fleet-average failure rates over the 24-month observation window for the illustrative $N = 5\,000$-vehicle fleet.

Elapsed	Per-vehicle exposure	Cumulative fleet exposure	Expected cumulative MOSFET failures	Fleet-average rate
6 months	1 825 h / 45 625 mi	9.1 M veh-h / 228 M veh-mi	≈ 0.7	—
12 months	3 650 h / 91 250 mi	18.3 M veh-h / 456 M veh-mi	≈ 3.0	1.6×10^{-7}/h · 6.6×10^{-3} IPMM
18 months	5 475 h / 136 875 mi	27.4 M veh-h / 684 M veh-mi	≈ 6.7	—
24 months	7 300 h / 182 500 mi	36.5 M veh-h / 912 M veh-mi	≈ 11.8	3.2×10^{-7}/h · 1.3×10^{-2} IPMM

These per-hour and per-mile rates are reported alongside the downstream consequential $\lambda(t)$ for "12 V battery unable to crank" and "12 V brown-out under accessory load." The bare-component hazard (≈ 6.5×10^{-7}/h at 24 months) is deliberately *not* compared to an ASIL target here. PMHF is an item-level metric, so a component has a contribution to it and not a budget of its own, and the comparison only becomes meaningful once the item budget has been apportioned. section 3.7.2 performs that apportionment and gives this component a share of 5.0×10^{-10}/h; section 3.7.3 then shows that whether the share is met is decided by the fault's architectural classification rather than by the magnitude of the estimate. Figure 3 makes these three Layer 4 outputs concrete.

- **L5 PHM:** anomaly detector flags vehicles whose cell-voltage spread exceeds the fleet baseline by an illustrative ≥ 3σ, or whose quiescent-current integral or balancing-FET duty cycle trends upward beyond a 10×-baseline threshold, precursor signatures of MOSFET leakage, distinguishable from genuine battery wear, that lead the operational SoH threshold by an illustrative ~3 000–5 000 miles and so give Layer 7 a containment lead time. RUL projection issued per vehicle for the BMS board separately from RUL for the battery cell stack.
- **L6 standards:** the item is the LV supply to a fail-operational LV load (redundant brake-by-wire); its safety goal **SG-LV-01** carries **ASIL D** and an item PMHF target of 10^{-8}/h. Layer 6 apportions that target across the Layer 1 element register (section 3.7.2, Table 3), giving the balancing MOSFET a share of **5.0×10^{-10}/h**, and then classifies the fault. Read as a **residual** fault, the component contributes 6.5×10^{-9}/h (**13× its share**, and 65 % of the entire item budget on its own given in table 4), and the claim fails; recovering it would require a diagnostic coverage of **99.92 %** against the 99 % the Layer 5 detector provides, and aggregated back to item level (section 3.7.4, Table 5) the item's PMHF would reach **148 % of its ASIL-D budget**. Read correctly, the redundant LV energy buffer in scope per section 3.3 means the fault cannot alone violate SG-LV-01: it is a **dual-point** fault contributing **9.5×10^{-11}/h**, or

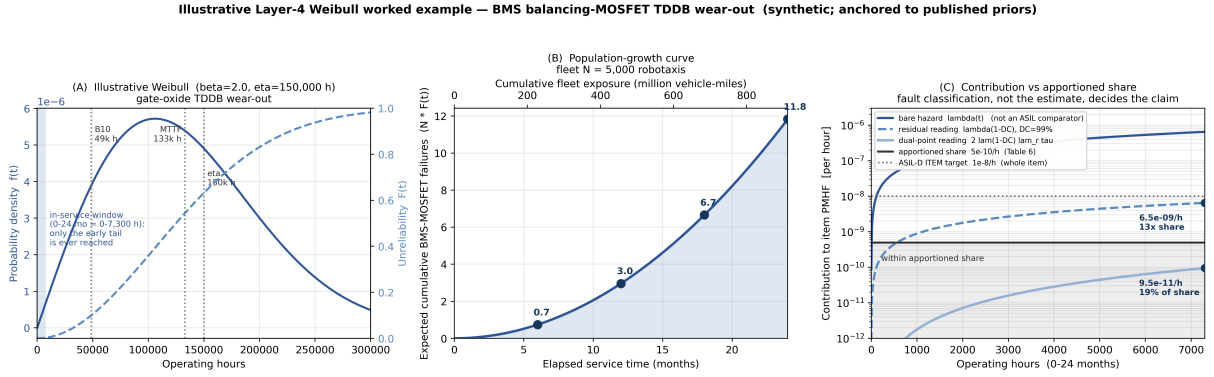

Figure 3: Illustrative Layer 4 Weibull worked example for the BMS balancing-MOSFET gate-oxide TDDB wear-out failure mode (synthetic; $\beta = 2.0$, $\eta = 150\,000$ h, anchored to published priors [8], [3], [43], [44]; not a measurement of any real fleet). **(A)** The Weibull PDF and CDF: the in-service exposure window (0–24 months ≈ 0–7 300 operating hours, shaded) reaches only the early tail of the distribution, with $B_{10} \approx 4.9\times10^4$ h and MTTF ≈ 1.33×10^5 h lying far beyond nominal service life. **(B)** The headline **population-growth curve**: expected cumulative failures across the $N = 5\,000$-vehicle fleet, $N{\cdot}F(t)$, rising convexly (the signature of wear-out, $\beta > 1$) to ≈ 11.8 vehicles by 24 months (912.5 M cumulative vehicle-miles). **(C)** The component's contribution to item PMHF against its **apportioned share** of 5.0×10^{-10}/h (section 3.7.2, Table 3), not against the item target, which a single component cannot be measured against. Three curves are shown: the bare hazard $\lambda(t)$, reaching ≈ 6.5×10^{-7}/h (649 FIT) at 24 months and not itself a comparator for any ASIL target; the residual-fault reading $\lambda(t)(1-K_{DC})$ at 99 % per-fault coverage, reaching 6.5×10^{-9}/h, or **13× the apportioned share**; and the dual-point reading $2\lambda(t)(1-K_{DC})\cdot\lambda_r\cdot\tau$, reaching 9.5×10^{-11}/h, or **19 % of the share**. The panel visualizes the section's substantive point: the same estimate supports or destroys the claim depending on how the fault is classified against the architecture, so it is the classification, not the reliability number, that carries the safety case. Reproduce: `python data/bms_mosfet_weibull.py`.

19 % of its share and the claim holds. The distinction is architectural, not statistical, and section 3.7 carries the residual reading as an explicit rebuttal rather than discarding it, because it states the condition (retention of the redundant buffer) on which the safety argument depends. The claim record documents the bench-test FIT prior for the MOSFET part, the field-derived λ, and the finding that the failure mode is BMS-induced rather than cell-induced, a distinction that changes both the safety-case argument and the corrective-action owner.

- **L7 containment:** Containment Review Board reviews the BMS-MOSFET fault. At 24 months the population-growth model and L5 risk scores identify **≈ 12 vehicles** past the fault threshold and an illustrative **~120 vehicles (2.4 % of the fleet)** with elevated-but-sub-threshold cell-voltage spread. For the ~120: **L7.0 Soft monitor** plus **L7.4 OTA mitigation** (revised balancing-FET duty profile cutting switching cycles ~40 % to slow leakage progression and extend η); for the ~12 past threshold: **L7.5 Pull from service** with routing to the depot for **L7.6 Depot service action** (BMS-board replacement overnight). Supplier 8D opened against the BMS vendor; failed BMS boards preserved for joint root-cause analysis; CA DMV and NHTSA notified per applicable thresholds if fault episodes occurred during autonomous operation.
- **L8 governance:** dashboard surfaces vehicles flagged for BMS-induced drift separately from vehicles with genuine battery aging; decision gate triggers supplier escalation against the BMS vendor (not the battery vendor) when the BMS-fault IPMM exceeds an illustrative governance threshold of 1×10^{-2} IPMM, crossed here at ~24 months (1.3×10^{-2} IPMM), noting that this gate is expressed in operational IPMM at the scenario's 25 mph duty cycle, and that the same component with an unchanged λ in a 45 mph highway fleet would report 7.1×10^{-3} and so would **not** cross the gate at all; section 4.5.4 sets out the normalized form a cross-fleet gate requires; lessons-learned propagated to other platforms using the same BMS architecture; the permanent fix is a BMS hardware revision (e.g., MOSFET part change or thermal-management improvement) tracked through the design-change workflow.

3.7. Layer 6 instantiated: a worked assurance fragment

section 3.6 produced a failure intensity. This section consumes it. The framework's central claim is that Layer 4 output can be carried into an assurance case without a manual re-derivation step, and that claim is only credible if the artefact at the far end is exhibited rather than described. What follows is one complete Layer 6 output for one LV safety goal: the claim record schema, the apportionment that gives the claim a target, the transfer that turns $\lambda(t)$ into a metric contribution, the criterion that decides whether the evidence suffices, and the assembled argument.

3.7.1. What Layer 6 emits, and why it is rendered in GSN

Neither standard the framework targets mandates a notation. ISO 26262 [1] requires a safety case as a work product but leaves its form open. UL 4600 [5] requires that the argument be structured, that evidence be identified per claim, and that the *sufficiency* of that evidence be justified, but it prescribes no diagram. Any framework that claimed to emit "the" standards-mandated artefact would therefore be overclaiming.

Layer 6 accordingly emits a **notation-neutral claim record**, of which Goal Structuring Notation is one rendering among several (Claims-Arguments-Evidence and plain tabular forms are equally admissible). The record carries fourteen fields:

`claim_id`	stable identifier, versioned
`statement`	the proposition asserted
`element`	Layer 1 component-register reference
`safety_goal`	goal identifier + ASIL from Layer 1 HARA
`scope`	element | item (which of the two the record asserts, section 3.7.4)
`metric`	if scope = element: PMHF contribution | operational IPMM | normalized IPMM if scope = item: PMHF | SPFM | LFM
`allocated_target`	apportioned share + pointer to the allocation rationale
`fault_class`	SPF | RF | MPF-latent | MPF-detected (ISO 26262-5)
`estimate`	Layer 4 point estimate + credible interval, per-hour AND per-mile
`exposure`	vehicle-hours, vehicle-miles, observation window
`evidence`	references to Layer 4 model, Layer 5 coverage evidence, bench prior
`sufficiency`	ES1-ES5 outcome with computed values (section 3.7.5)
`rebuttals`	conditions under which the claim fails, and their disposition
`status`	SUPPORTED | CONDITIONAL | NOT SUPPORTED, with owner and refresh date

The `scope` field is not bookkeeping. It is the schema-level guard against a category error: PMHF, SPFM, and LFM are item metrics, and a record that asserts one of them at element scope is malformed by construction rather than merely wrong. An element may carry a *contribution* to item PMHF, and it may carry its own per-mile rate; it has no SPFM and no LFM at all, for the reason section 3.7.4 gives. Table 8 shows the field populated for an element-scoped claim.

Against this, the two alignment claims the framework makes become checkable rather than rhetorical. A record is **ISO 26262-aligned** when its `metric`, `fault_class`, and `allocated_target` fields use the definitions of ISO 26262-5 (PMHF, the single-point/residual/multiple-point fault taxonomy, and an apportionment of the item-level target),

so that the record composes into a quantity the standard actually recognizes. A record is **UL 4600-aligned** when it carries the three things UL 4600 demands of an argument and reliability reporting alone does not supply: an explicit claim, an explicit sufficiency justification for the evidence offered, and explicit rebuttals. GSN is chosen for the rendering because it has a published community standard [45], a two-decade tool and practice base [46], and cross-domain precedent in exactly this role [17], not because either standard requires it.

3.7.2. *Apportioning the item budget*

PMHF is an item-level metric. It is evaluated for the item realizing a safety goal, and a single component has no PMHF of its own; it has a *contribution*. A framework that emits per-component field estimates must therefore also specify how the item target is divided, or its per-component claims have nothing to be measured against.

Take as the item the LV electrical supply to the fail-operational brake-by-wire channel, and as its safety goal **SG-LV-01**, *"the vehicle shall not lose braking capability due to loss of low-voltage electrical supply,"* assigned ASIL D by the Layer 1 HARA. ISO 26262-5 gives a recommended PMHF target of 10^{-8}/h for ASIL D.

Layer 6 apportions that target across the Layer 1 element groups in proportion to their bench-test FIT priors ([40], [42], [41]), normalized, with an explicit reserve held for elements not yet in the register. The apportionment must be complete and non-overlapping; that property is itself a claim in the argument (G2 in Figure 4), not an assumption.Within the energy-storage group, the BMS board takes half and the cell-balancing MOSFET half of that, giving the section 3.6 component an allocated share of **λ_alloc = 5.0×10⁻¹⁰/h**. Every downstream claim about that component is measured against this number and not against the item target.

3.7.3. *From λ(t) to a metric contribution*

This is the bridge, stated as a procedure. Given a Layer 4 output λ_e(t | duty cycle) for element *e*:

- **M1.** Evaluate at t = t_claim, the end of the claimed service interval. For a wear-out mechanism ($\beta > 1$) this is the conservative choice; the hazard is monotonically increasing, so the endpoint bounds the interval.
- **M2.** Classify the fault against the Layer 1 architecture using the ISO 26262-5 taxonomy: does the fault alone violate the safety goal (single-point or residual), or only in combination with another (multiple-point)? This step is architectural, not statistical, and it is where most of the assurance content lives.
- **M3.** Apply the transfer function for that class (Table 4).
- **M4.** Compare the result to λ_alloc,e from Table 3.
- **M5.** Convert to per-mile using the Layer 3 fleet-average speed, recording the speed and its provenance in the claim record (see section 4.5 on the conditions under which the per-mile figure is comparable across operators).
- **M6.** Run the sufficiency criterion (section 3.7.5) and emit the claim record with a status.

Table 3: Apportionment of the SG-LV-01 item PMHF budget (ASIL D, 10^{-8}/h) across Layer 1 element groups. Illustrative shares, proportional to normalized bench FIT priors.

Element group	Share	Allocated contribution
Energy storage & management (12 V battery, BMS)	20 %	2.0×10⁻⁹/h
DC-DC conversion (primary + redundant)	20 %	2.0×10⁻⁹/h
Power distribution & protection	15 %	1.5×10⁻⁹/h
On-ECU power supply circuitry	15 %	1.5×10⁻⁹/h
Harnesses, connectors, grounds	15 %	1.5×10⁻⁹/h
In-vehicle communication hardware	10 %	1.0×10⁻⁹/h
Unallocated reserve	5 %	5.0×10⁻¹⁰/h
Total	**100 %**	**1.0×10⁻⁸/h**

Table 4: Fault-class transfer functions and their evaluation for the section 3.6 BMS balancing MOSFET at t = 7 300 h, λ_e = 6.5×10⁻⁷/h, DC = 99 %, λ_alloc = 5.0×10⁻¹⁰/h.

Fault class	Contribution to item PMHF	Evaluated	vs. allocation
Single-point fault	λ_e	6.5×10⁻⁷/h	1 300×, not supported
Residual fault	λ_e (1 − DC)	6.5×10⁻⁹/h	13×, **not supported**
Multiple-point, latent	2 λ_e(1−DC) · λ_r(1−DC_r) · τ	9.5×10⁻¹¹/h	19 %, **supported**
Multiple-point, detected and repaired within τ_repair	≈ 0, documented	—	supported

The dual-point row uses the conservative symmetric form of the simplified dual-point estimator, with $\lambda\_r = 10^{-6}$/h for the redundant LV supply path (illustrative, no coverage credit taken) and $\tau = 7\,300$ h; the precise form depends on the detection and repair model assumed, which the claim record is required to state.

Two observations follow, and both matter more than the arithmetic.

First, **the classification decides the outcome, not the estimate.** The same $\lambda\_e = 6.5\times10^{-7}$/h supports the claim or destroys it depending on M2. This is the substantive reason a reliability estimate cannot be handed to a safety case unmediated, and therefore the substantive justification for Layer 6 existing at all.

Second, **the residual-fault reading fails.** If the redundant LV buffer is absent, or is de-scoped late in a programme, the MOSFET consumes 13× its share, and 65 % of the *entire* ASIL-D item budget on its own. The coverage that would be required to rescue the residual reading is 99.923 %, against the 99 % the Layer 5 detector provides. That is not a rounding difference; it is a design constraint, and Layer 6 is where it surfaces. It is carried in the fragment as rebuttal R1 rather than being suppressed.

3.7.4. *Aggregating to the item: PMHF, SPFM, and LFM*

Two things happen only at item level, and section 3.7.3 has set up both. The per-element PMHF contributions must be summed back up, or the apportionment of section 3.7.2 is a division whose result is never checked. And SPFM and LFM must be computed here or not at all.

PMHF is a rate, and rates compose additively over elements. That is what makes section 3.7.3 possible: an element has a contribution to item PMHF, the contributions sum, and a per-element claim is therefore well formed. SPFM and LFM do not behave this way, and the difference is not a technicality; it determines what Layer 6 is permitted to emit.

Both are *ratios* taken over the item's entire safety-related hardware failure-rate population:

$$\mathrm{SPFM} = 1 - \Sigma(\lambda\_\mathrm{SPF} + \lambda\_\mathrm{RF}) \,/\, \Sigma\,\lambda\_\mathrm{total}$$
$$\mathrm{LFM} = 1 - \Sigma\,\lambda\_\mathrm{MPF,latent} \,/\, (\Sigma\,\lambda\_\mathrm{total} - \Sigma(\lambda\_\mathrm{SPF} + \lambda\_\mathrm{RF}))$$

Every element appears in the numerator and the denominator at once. An element therefore has no SPFM of its own, and the quantity is not merely unmeasured but undefined: adding a highly reliable element to the item raises the item's SPFM without changing any existing element's properties, so no assignment of SPFM to an element could be stable. This is the same category error as treating PMHF as a per-component budget, one level up, and section 3.7.1's `scope` field exists to make it unrepresentable.

What Layer 4's output does supply is the input to the item-level computation, by way of the fault-class partition M2 already performs. Two further steps complete the mapping:

- **M7: Partition.** For each element in the Layer 1 register, split λ_total(t_claim) into λ_SPF, λ_RF, λ_MPF-latent, λ_MPF-detected, and λ_safe, using the Layer 1 FMEDA fault-mode distribution and the Layer 5 per-fault diagnostic coverage. Evaluate at t_claim per M1, so the partition is consistent with the PMHF contributions.
- **M8: Aggregate.** Sum each bucket across the register and evaluate all three item metrics: PMHF as $\Sigma(\lambda\_\mathrm{SPF} + \lambda\_\mathrm{RF}) + \Sigma\lambda\_\mathrm{DPF}$, and SPFM and LFM by the ratios above. Compare each to its ASIL target. The result is an **item-scoped** claim record: one per item, never one per element. Aggregating PMHF here is what closes the loop opened in section 3.7.2: the apportionment divides the budget top-down, and M8 checks bottom-up that the division was met.

Table 5: Item-level metric roll-up for SG-LV-01, in FIT (1 FIT = 10^{-9}/h), with each element group's λ_total evaluated at t_claim = 7 300 h. The → PMHF column is the group's contribution, λ_SPF+RF + λ_DPF, and *used* is that contribution as a fraction of the group's Table 3 allocation. Illustrative and consistent with the Table 3 register; not measurements of any real fleet.

Element group	λ_total	λ_SPF+RF	λ_MPF-latent	λ_DPF	→ PMHF	Table 3 alloc.	used
Energy storage & management (12 V battery, BMS)	900	1.5	78.0	0.35	1.85	2.0	92 %
DC-DC conversion (primary + redundant)	700	1.6	60.0	0.30	1.90	2.0	95 %
Power distribution & protection	400	1.1	34.0	0.22	1.32	1.5	88 %
On-ECU power supply circuitry	500	1.0	42.0	0.25	1.25	1.5	83 %
Harnesses, connectors, grounds	300	1.2	24.0	0.18	1.38	1.5	92 %
In-vehicle communication hardware	200	0.6	12.0	0.10	0.70	1.0	70 %
Item total	**3 000**	**7.0**	**250.0**	**1.40**	**8.40**	**9.5**	**88 %**
Unallocated reserve	—	—	—	—	—	0.5	—
Item budget						**10.0**	

PMHF = 7.0 + 1.40 = 8.40 FIT = 8.4×10^{-9}/h (84 % of the item budget)
SPFM = 1 − 7.0 / 3 000 = 99.77 %
LFM = 1 − 250.0 / 2 993 = 91.65 %

The PMHF figure is compared to the apportioned item budget of section 3.7.2. The two ratio metrics are reported as computed: their ASIL-dependent targets are set by ISO 26262-5 and are a matter for the programme's own safety plan, and nothing in the argument below turns on the particular thresholds adopted.

The first line is the aggregation the apportionment was for. section 3.7.2 divided the item budget top-down; M8 confirms bottom-up that the division was met, with every group inside its allocation and 16 % of the item budget unconsumed. Goal G1 in Figure 4 is discharged by this arithmetic, not by justification J1 alone; J1 states why a complete apportionment *would* imply the item claim, and Table 5 is where that implication is actually cashed.

The section 3.6 MOSFET sits inside the first row: 649 FIT of that group's 900 FIT, of which $\lambda_e(1-K_{DC})$ = 6.49 FIT is undetected and, under the dual-point classification of section 3.7.3, falls in the latent bucket and contributes 0.095 FIT to λ_{DPF}.

The rebuttal reaches the item level. If the redundant LV energy buffer is de-scoped (rebuttal R1), the MOSFET's fault becomes residual, those 6.49 FIT move from the latent bucket into λ_{SPF+RF}, and its 0.095 FIT dual-point contribution disappears. Nothing else about the item changes:

The energy-storage group alone then contributes 8.24 FIT against its 2.0 FIT allocation, and the MOSFET 6.49 FIT against its 0.5 FIT share. One reclassification of one component overruns the ASIL-D item budget by 48 %.

Two things follow, and the second is the more useful.

First, **this is the aggregation failure in its sharpest form.** A per-component reading that stopped at "6.5×10^{-9}/h is a small number" would not surface it; only summing the contributions back against the item budget does. It is also why R1 is carried in the fragment as an explicit rebuttal rather than as a remark in section 3.6: the condition it names is the difference between an item that meets its ASIL-D target and one that misses it by half again.

Second, **the ratio metrics do not notice.** SPFM moves 0.22 percentage points and LFM improves. No threshold applied to either would change its verdict on a movement of that size, so the reclassification that overruns the item's PMHF budget by half is effectively invisible to both: 6.49 FIT is negligible against a 3 000 FIT denominator. SPFM, LFM, and PMHF are independent criteria and none implies another, a point ISO 26262 makes structurally by requiring all three, and one this framework is well placed to illustrate because it produces the underlying rates. It is the concrete reason Layer 6 reports apportioned PMHF contributions alongside the architectural metrics rather than in place of them. Reproduce: `python figures/item_metrics_rollup.py`.

A final consequence deserves stating, because it is a property the framework introduces rather than inherits. Conventional FMEDA practice evaluates these metrics from constant base failure rates, so they are fixed for the life of the item. A fleet-derived $\lambda(t)$ with $\beta > 1$ makes all three **functions of service time**, and for the ratios the direction of movement is not always intuitive: a wear-out mechanism in an element whose faults are classified safe or detected inflates $\Sigma\lambda_{total}$ and therefore *raises* SPFM while the item is objectively degrading. A framework that produces time-resolved failure intensities is obliged to say which of its outputs survive the transition from constant rates to $\lambda(t)$. PMHF survives it and becomes more informative; the two ratio metrics survive it less cleanly, and should be read at a stated t_{claim} or not at all.

3.7.5. *When is the evidence sufficient?*

UL 4600 requires evidence sufficiency to be justified per claim but does not define a test. Layer 6 supplies one. A field-derived claim $\lambda_e \le \lambda_{target}$ is **sufficient** iff all five conditions hold:

- **ES1: Statistical adequacy.** The upper limit of the one-sided $(1-\alpha)$ credible interval satisfies $\bar{\lambda}_e(1-\alpha) \le \lambda_{target}$, with $\alpha \le 0.05$. For a homogeneous Poisson process with k failures in exposure E, $\bar{\lambda}(1-\alpha) = \chi^2_{\{1-\alpha\}}(2k+2) / 2E$ [47].
- **ES2: Exposure adequacy.** $E \ge E_{min} = -\ln \alpha / \lambda_{target}$, the exposure at which a zero-failure observation would carry the claim. Reported whether or not ES1 passes, because it states what the fleet would have to accumulate.
- **ES3: Representativeness.** The Layer 3 duty-cycle envelope of the contributing exposure covers the claimed ODD on every covariate Layer 4 found significant.

Table 6: Effect of rebuttal R1 on the three item-level metrics for SG-LV-01. Reclassifying the balancing MOSFET from a dual-point to a residual fault moves 6.49 FIT between buckets and nothing else about the item changes.

Metric	As classified	Under R1	Movement	Consequence
Item PMHF	8.40 FIT (84 % of budget)	**14.79 FIT (148 %)**	+64 pp of budget	**item budget exceeded**
Item SPFM	99.77 %	99.55 %	−0.22 pp	no material change
Item LFM	91.65 %	91.85 %	+0.20 pp	no material change

Table 7: Evidence sufficiency evaluated for the illustrative 5 000-vehicle, 24-month fleet (E = 3.65×10^{7} vehicle-hours / 9.12×10^{8} vehicle-miles, k ≈ 12).

Quantity claimed	λ_target	E_min (ES2)	vs. fleet exposure	ES1
λ_e itself	6.5×10^{-7}/h	4.6×10^{6} veh-h	0.13×	pass, 95 % UB 5.3×10^{-7}/h
ASIL-B item target	1.0×10^{-7}/h	3.0×10^{7} veh-h	0.82×	marginal
ASIL-D item target	1.0×10^{-8}/h	3.0×10^{8} veh-h	8.2×	fail
MOSFET apportioned share	5.0×10^{-10}/h	6.0×10^{9} veh-h	164×	fail

- **ES4: Observability.** Every failure mode in the element's FMEDA has a defined field-detectable signature in the Layer 2 data contract. Modes without one are excluded from the claim and carried as documented rebuttals.
- **ES5: Currency.** The estimate was refreshed within the Layer 8 governance cadence, and the claim carries the exposure window over which it was computed.

The last row is the important one, and it generalizes. **A component-level ASIL-D allocation is not confirmable by field evidence alone at any plausible fleet scale.** An operator would need 164× the exposure of a 5 000-vehicle fleet run for two years to retire a single 5.0×10^{-10}/h claim on observation. This is not a defect of the framework; it is the quantitative form of why safety cases for ASIL-D items are carried by architecture and diagnostic coverage rather than by demonstrated failure counts, and Layer 6 is where that becomes explicit rather than tacit. What field evidence *can* do (and does, comfortably, at 0.13× the available exposure) is size the input λ_e that the architectural argument then consumes. That division of labour is the honest statement of what fleet evidence contributes to an assurance case, and it also gives section 1.1's 10^{8}–10^{9} vehicle-hour figure a derivation rather than an assertion.

3.7.6. *The claim record, instantiated*

Section 3.7.1 specified the record; the preceding sections computed everything it needs. Table 8 is that record populated for the section 3.6 component. It is the artefact Layer 6 actually emits (Figure 4 is a rendering of it, not a substitute for it), and it is what a reviewer or auditor would receive.

Table 8: A populated Layer 6 claim record. Field names follow section 3.7.1 schema. Values are illustrative and inherit the section 3.6 synthetic parameters.

Field	Value
`claim_id`	`CLM-SGLV01-ESM-Q07-001`, v1.2
`statement`	Over the 0–7 300 h claimed service interval, the gate-oxide TDDB failure mode of the BMS cell-balancing MOSFET contributes ≤ 5.0×10^{-10}/h to violation of SG-LV-01.
`element`	`LV-ESM-BMS-Q07`: cell-balancing MOSFET, BMS board (Layer 1 register, section 3.3)
`safety_goal`	SG-LV-01, **ASIL D** (Layer 1 HARA)
`scope`	element
`metric`	PMHF contribution
`allocated_target`	5.0×10^{-10}/h, from Table 3: energy-storage group 20 % of the 10^{-8}/h item budget → BMS board 50 % → balancing MOSFET 50 %
`fault_class`	**MPF-latent**, per M2 against the Layer 1 architecture with the redundant LV energy buffer in scope (section 3.3)
`estimate`	Hazard at t_claim: λ_e(7 300 h) = 6.5×10^{-7}/h (649 FIT), Weibull β = 2.0, η = 1.5×10^{5} h. Window-average MLE 3.3×10^{-7}/h from k = 12 failures in E = 3.65×10^{7} veh-h; 95 % one-sided upper bound 5.3×10^{-7}/h. Contribution after the M3 transfer: **9.5×10^{-11}/h**, 19 % of the allocated share. Per-mile: 1.3×10^{-2} operational IPMM at $\bar{v}$ = 25 mph (normalized IPMM identical, $\bar{v}$ = $\bar{v}$_ref)
`exposure`	5 000 vehicles × 24 months = 3.65×10^{7} vehicle-hours / 9.12×10^{8} vehicle-miles; 0–7 300 operating hours per vehicle
`evidence`	Sn2: Layer 1 architecture + FMEDA (redundant buffer sustains brake-by-wire); Sn3: Layer 5 detector, cell-voltage spread + quiescent-current integral, K_DC = 99 %; Sn4: Layer 4 Weibull model (Figure 3); Sn5: sufficiency record; bench FIT priors [40], [42], [41]
`sufficiency`	ES1 **pass** (95 % UB 5.3×10^{-7}/h ≤ 6.5×10^{-7}/h). ES2 **pass** (E_min = 4.6×10^{6} veh-h, 0.13× available exposure). ES3 **pass** (Layer 3 envelope covers the claimed ODD on every significant covariate). ES4 **partial**: two FMEDA modes have no field-detectable signature in the Layer 2 contract; excluded from the claim and carried as R3. ES5 **pass**, refreshed within the Layer 8 cadence. *Sufficiency is asserted for λ_e, not for the 5.0×10^{-10}/h share, which is not field-confirmable at this fleet scale (Table 7, last row) and is carried architecturally.*

Field	Value
`rebuttals`	**R1**: if the redundant LV energy buffer is unavailable or de-scoped, the fault becomes residual at 6.5×10^{-9}/h (13× the share), recovery requires K_DC ≥ 99.92 %, and the item's PMHF reaches 148 % of its ASIL-D budget (section 3.7.4). *Disposition:* buffer confirmed present in the current architecture baseline; re-evaluate on any change to LV redundancy. **R2**: common cause between the BMS supply and the redundant buffer would void assumption A1 and force the residual reading. *Disposition:* open; referred to the Layer 1 dependent-failure analysis. **R3**: the two excluded FMEDA modes. *Disposition:* documented, not claimed.
`status`	**SUPPORTED**, conditional on R1 and R2 remaining discharged. Owner: LV subsystem safety owner. Refresh: next Layer 8 governance cycle, or on any change to the LV redundancy architecture, whichever is sooner.

Three things in this record are worth drawing out, because they are what a conventional reliability report does not carry and are the substance of the UL 4600 alignment claimed in section 3.7.1.

First, **ES4 is recorded as partial rather than rounded to a pass.** Two failure modes have no field-detectable signature, so the claim is narrower than the element, and the record says so. Second, **R2 is recorded as open.** The dual-point reading depends on an independence assumption the Layer 1 dependent-failure analysis has not yet discharged, and an argument that concealed this would be unsound in exactly the way an assurance case is meant to prevent. Third, **the status is conditional, and names its conditions.** The claim holds because of an architectural property, not because a number came out small, and the record makes the loss of that property a traceable event rather than a silent one.

G3 in Figure 4 is this record's `statement` field; Sn2–Sn5 are its `evidence` field; R1 is its first rebuttal. The fragment and the record are the same object in two notations, which is what section 3.7.1's notation-neutrality claim amounts to in practice.

3.7.7. *The assembled fragment*

Figure 4 renders the claim records above as a GSN structure. The top goal G1 is the item-level PMHF claim, in the context of the item definition (C1), the ASIL assignment (C2), and the target (C3). Strategy S1 decomposes it by apportionment, justified by J1 (the item metric is a sum over element contributions, so a complete, non-overlapping apportionment that every element meets implies the item claim), and bounded by assumption A1 on independence. G2 claims the apportionment itself is sound. G3 is the MOSFET's share; strategy S2 argues it by fault classification, and the three subgoals G3.1–G3.3 are precisely the three things M2 and M3 require: that the fault cannot act alone, that it is detected before a second fault accumulates, and that the resulting multiple-point contribution is within the share. Solutions Sn1–Sn5 tie each to a Layer 1, 4, or 5 artefact. G4–G9 are marked undeveloped: the fragment is one branch of a complete case, not the whole case.

R1 is the element most easily omitted and least safely omitted. It records that the whole argument turns on the redundant LV energy buffer remaining in the architecture, and states what the claim would cost if it did not: a residual contribution 13× the apportioned share, requiring 99.92 % diagnostic coverage to recover, and an item PMHF at 148 % of its budget. Carrying this as an explicit rebuttal rather than an unstated premise is what makes the fragment auditable, and it is the mechanism by which Layer 6 feeds design constraints back to Layer 1 through the Layer 8 loop.

4. Discussion

4.1. *Validity by argument*

Because the framework is presented without empirical validation on real AV fleet data, its validity rests on three constructive arguments.

4.1.1. *Completeness: does the framework cover what the standards require?*

The framework covers the standards' major reliability-evidence requirements end-to-end. Gaps deliberately not addressed (e.g., software functional safety, scenario coverage for ML perception) belong to SOTIF / UL 4600 workflows that consume (but do not duplicate) this framework.

4.1.2. *Compatibility: does each layer map to an existing engineering artefact?*

Each layer maps cleanly to an artefact engineering teams already produce; the framework's contribution is organizing them into a coherent end-to-end pipeline specific to LV AV subsystems.

4.1.3. *Gap closure: which gaps from section 2.6 does the framework close?*

- **Supplying the method that joins fleet reliability estimates to AV assurance claims:** L4 outputs flow into L6's apportionment, fault classification, and sufficiency test, and out as the claim records and GSN fragment worked through in section 3.7 and Figure 4. ISO 26262 Part 7 obliges the monitoring and Part 8 admits the evidence; neither constructs the claim, and that construction is the central novelty.
- **AV-LV electrical reliability under-study:** L1's scope explicitly enumerates the AV LV subsystems (BMS, DC-DC, ECUs, comms, harnesses, connectors) that prior literature largely omits.

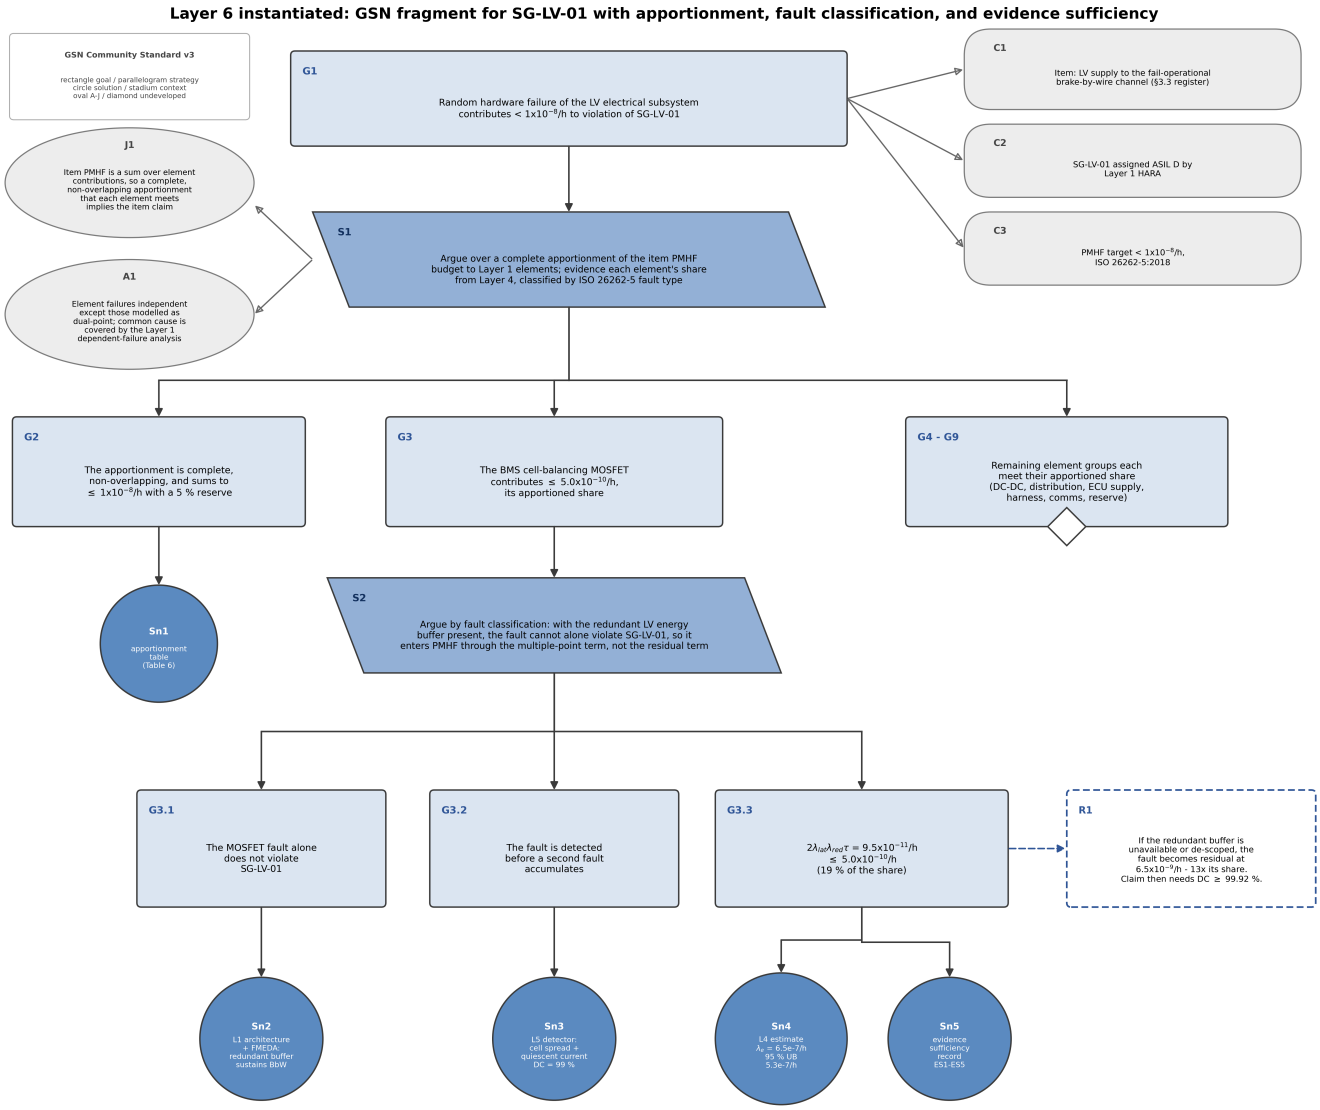

Figure 4: Layer 6 instantiated. A worked GSN fragment (GSN Community Standard v3 [45]) for safety goal SG-LV-01, showing the apportionment strategy, the fault classification that determines which transfer function applies, the evidence solutions drawn from Layers 1, 4, and 5, and the rebuttal R1 recording the architectural condition the claim depends on. Values are illustrative and inherit the section 3.6 synthetic parameters; they are not measurements of any real fleet. Reproduce: `python figures/gsn_fragment.py`.

Table 9: Coverage of the standards' principal reliability-evidence requirements by framework layer.

Standard requirement	Framework layer(s) that address it
ISO 26262 HARA inputs (component classification, ASIL assignment)	L1
ISO 26262 random-hardware-failure quantification (PMHF / SPFM / LFM)	L4 estimates contributions; L6 apportions the item target, classifies faults, and aggregates back to item level (section 3.7.2–section 3.7.4)
ISO 26262 field monitoring process (Part 7) and use of service experience as evidence (Part 8, proven in use)	L2, L3, L4, L8 supply the estimation method these provisions require but do not specify; L6 constructs the claim
ISO 21448 / SOTIF: exposure to operational variability	L3
UL 4600 goal-based safety case (GSN claims, evidence sufficiency)	L6
Field-incident response, corrective action, regulator notification (ISO 26262 Part 7; UL 4600; NHTSA SGO 2021-01; NHTSA Part 573)	L7
Continuous safety case maintenance	L8

Table 10: Mapping of each framework layer to an existing engineering artefact it produces or extends.

Layer	Existing engineering artefact it produces or extends
L1	Component register / parts list with safety attributes: extends ISO 26262 item definition
L2	Data pipeline / data lake: standard practice for connected-vehicle programmes
L3	Exposure / duty-cycle dataset: extends conventional warranty-analysis covariates
L4	Reliability prediction report: extends FIT analyses (SN 29500 [41], IEC 62380 [42]) with field-derived λ; emits IPMM as the operator- and regulator-facing figure, together with the mechanism's native denominator that determines how it may be used (section 4.5)
L5	OBD / PHM diagnostics: extends current on-board diagnostics with population-aware priors
L6	Safety case fragment: directly populates GSN / UL 4600 case
L7	Containment Review Board workflow + 8D corrective action: standard quality-engineering practice, with AV-fleet-specific levers (mission/ODD restriction, conservative mode, OTA, depot routing)
L8	Reliability KPI dashboard + governance / preventive action: standard quality-engineering practice

- **In-vehicle communication architecture as first-class subsystem:** L1 and L5 include CAN/Ethernet/LIN/FlexRay hardware and their characteristic failure modes (FEC corrections, link drops, bus error frames).
- **Connectors / harnesses / ECU power supply as first-class subsystem:** L1 and L4's physics-of-failure list explicitly cover them.

4.2. Deployment requirements

What an AV programme must provide to instantiate the framework:

- **Telemetry:** LV bus voltage and current at component-level granularity; BMS SoH/SoC; CAN/Ethernet error counters; ECU brown-out events; wake/sleep cycles; ambient temperature.
- **Population data:** vehicles-in-operation and operating-hours by configuration.
- **Failure events:** service records, fault codes, replacement records, joined to vehicle and component identifiers.
- **Bench / supplier data:** datasheet FIT and accelerated-life results to anchor Layer 4 priors.
- **Standards expertise:** to validate Layer 6 GSN fragments and PMHF mapping.
- **Organizational:** a Containment Review Board (Layer 7) with authority to deploy mission/ODD restrictions, OTA mitigations, and depot routing; a continuous-improvement function (Layer 8) with authority to act on dashboard outputs.

A regulator (NHTSA, CA DMV, EU type-approval authorities) instantiating the framework on aggregate public data would obtain a population-level reliability view but not per-vehicle PHM (Layer 5 degrades to forward-looking only).

4.3. Threats to validity / limitations

- **No empirical case study.** The framework is presented as a methodological proposal; per-layer effectiveness on a real AV fleet has not been measured. This is a deliberate scoping choice driven by the absence of public AV LV telemetry, not an oversight.
- **Layer 4 statistical methods rely on bench-test priors when fleet data is sparse.** Until production fleet exposure accumulates, FIT estimates inherit datasheet uncertainty.
- **The instantiated fragment is one branch of one safety goal.** section 3.7 works SG-LV-01 through to evidence solutions for a single element, and G4–G9 in Figure 4 are marked undeveloped. The Table 3 apportionment shares and the Table 5 FMEDA fault partition are illustrative rather than programme-derived, and a complete item-level case would require the same treatment for every element in the Layer 1 register. The ES1–ES5 evidence-sufficiency criterion is offered as a proposal: it has not yet been reviewed by independent assurance-case practitioners, nor exercised against a real certification submission, and the threshold choices within it ($\alpha \leq 0.05$, the Layer 8 currency cadence) are conventions rather than derived values.
- **The apportionment shares are illustrative.** Table 3 divides the item budget in proportion to normalized bench FIT priors, which is a defensible default but not the only one; programmes commonly apportion by architectural criticality, by supplier accountability boundary, or by negotiated margin. The framework requires that an apportionment exist, be complete, and be recorded, not that it take this particular form. Different shares move every per-component verdict in section 3.7.
- **The dual-point estimator is model-sensitive.** The contribution $2\lambda_e(1-K_{DC})\cdot\lambda_r\cdot\tau$ depends on the assumed multiple-point fault detection interval τ,

on the redundant path's rate λ_r (illustrative here at 10^{-6}/h, with no coverage credit taken), and on the assumption of independence recorded as A1 in Figure 4. A dependent-failure analysis that found common cause between the BMS supply and the redundant buffer would invalidate the dual-point reading and force the residual reading, the condition rebuttal R1 exists to surface.

- **AV-specific failure modes are partially unknown.** LV transients from compute inference power spikes, sensor-supply ripple under perception load, and redundant-supply changeover events are absent from prior literature; the framework anticipates them in scope but cannot yet quantify them.
- **Connector and harness failure modes are mechanical-electrical.** Layer 4's physics-of-failure models for fretting, vibration, and corrosion are mature in aerospace literature but under-applied to automotive LV; transferability requires confirmation.

4.4. Comparison with existing practice

The comparison of the framework with existing reliability and assurance practice is given in table 11.

4.5. The comparability of per-mile reporting

The framework reports in two units, and the reason is audience: ISO 26262 speaks in per-hour terms, while operators and regulators report per mile. Producing both is straightforward. Treating them as interchangeable is not, and this section states the conditions under which the per-mile figure supports the comparisons it is usually put to.

4.5.1. The per-mile figure inherits the duty cycle

Write λ_h for a component's per-hour failure intensity and $\bar{v}$ for the fleet-average speed over powered hours from Layer 3. Then $IPMM = \lambda_h \cdot 10^6 / \bar{v}$. If the underlying mechanism is time-driven, λ_h is the quantity fixed by the hardware and the environment, and every part of the variation in IPMM across fleets is contributed by $\bar{v}$:

$$IPMM_A / IPMM_B = \bar{v}_B / \bar{v}_A$$

Figure 5 evaluates this for the section 3.6 component across four AV operator archetypes. The reported IPMM spans 4.6×, from 5.41×10^{-2} for a dense-urban robotaxi averaging 12 mph to 1.18×10^{-2} for long-haul freight at 55 mph, with identical hardware, identical physics, and identical $\lambda_h = 6.5\times10^{-7}$/h in every case. The fleet that appears worst is not worse; it accumulates more powered hours per mile.

Table 11: Comparison of the framework with existing reliability and assurance practice.

Existing practice	Limitation	How this framework addresses it
Bench-test-only reliability (datasheet FIT, accelerated life)	Under-represents fleet operational variability	L3 + L4 incorporate field-derived covariates
Warranty data analysis (Weibull / mixed-Weibull)	Decoupled from safety case; produces a rate, not a claim, and no rate is comparable to an ASIL target without an allocation	L6 supplies the missing steps: apportionment of the item target (section 3.7.2), fault classification against the architecture (section 3.7.3), aggregation back to the item metric (section 3.7.4), and an evidence-sufficiency test (section 3.7.5)
OEM-internal PHM	Not standardized, rarely published, hard to audit	Source-agnostic layer contracts; auditable L2–L4 pipeline
UL 4600 safety cases	Often unsupported by quantitative field evidence, and the standard does not define a sufficiency test	L4 outputs feed L6, which emits claim records carrying an explicit ES1–ES5 sufficiency outcome and its rebuttals (section 3.7.5–section 3.7.6)
ISO 26262 Part 7 field monitoring	Establishes the duty to collect field data and react to safety anomalies, but prescribes no estimation method, no exposure model, and no route from an observed rate to an assurance claim	L2–L4 supply the estimation method and uncertainty treatment; L6 supplies apportionment, fault classification, and the sufficiency criterion (section 3.7)
ISO 26262 Part 8 proven in use	Uses accumulated service experience quantitatively, but to decide whether it substitutes for development evidence, not to establish an element's apportioned contribution to the item metric	L6 generalizes the exposure requirement into a per-claim sufficiency criterion (ES1–ES5)
AV disengagement / incident reports (CA DMV, NHTSA)	Aggregate, narrative, no LV root-cause stratification, and per-mile figures that are not cross-operator comparable, since mileage denominators and duty cycles differ between permit holders (section 4.5.3, CC3)	Provides a target consumer for the framework's outputs; suggests a minimal disclosure schema regulators could require

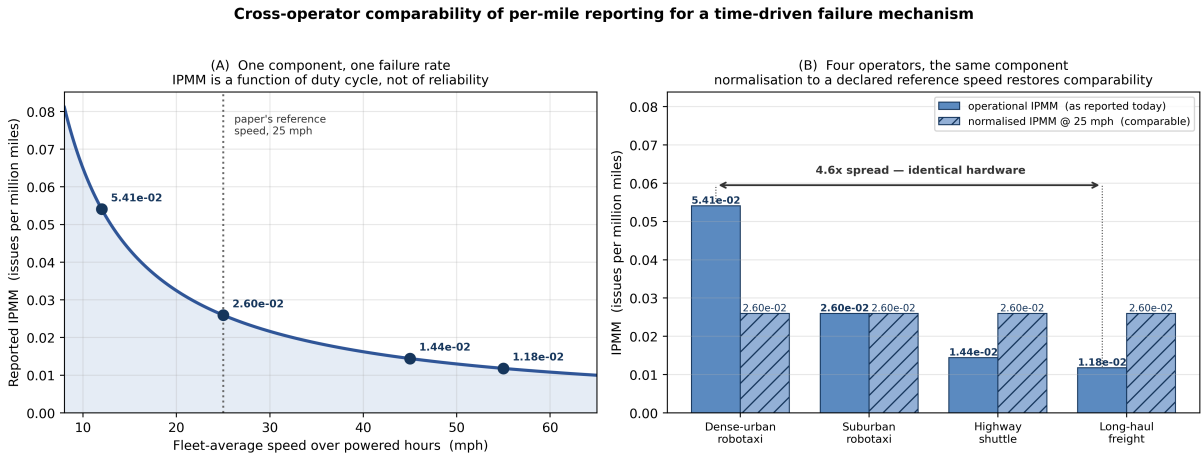

Figure 5: Cross-operator comparability of per-mile reporting for a time-driven mechanism, using the section 3.6 component at 24 months (λ_h = 6.5×10^{-7}/h; illustrative, not a measurement of any real fleet). **(A)** Reported IPMM as a function of fleet-average speed for a single fixed failure rate; the curve is a statement about duty cycle, not about reliability. **(B)** Four operator archetypes reporting the same component: operational IPMM spans 4.6× while normalized IPMM at a declared 25 mph reference speed is identical across all four. Reproduce: `python figures/ipmm_comparability.py`.

Note that $\bar{v}$ is not simply road speed. Time spent powered but stationary (depot dwell with compute energized, passenger loading, traffic) accumulates in the denominator's numerator, so to speak, and depresses $\bar{v}$ without any driving occurring. For an AV fleet running perception and compute continuously, idle dwell is a first-order contributor, which is why Layer 3 tracks idle hours and stop-start cycles as first-class exposure covariates rather than as diagnostics.

4.5.2. *Which unit is native depends on the mechanism*

The choice of denominator is not a reporting convention; it is a claim about the physics. A component's cumulative stress exposure is driven by some combination of powered time, distance, and cycle count, and the appropriate denominator is whichever term dominates. Layer 4's physics-of-failure model already determines this, so the framework has the information; it simply has to declare it.

For the LV electrical subsystem the answer is lopsided. Gate-oxide TDDB, capacitor ESR drift, electromigration, sulfation and corrosion are time-driven. Solder-joint fatigue, connector fretting, relay and contactor wear, and wake-sleep inrush stress are cycle-driven. Distance-driven mechanisms are the minority, and where distance does predict failure it is usually standing proxy for vibration cycles rather than acting as a cause. The per-mile unit is therefore, for this subsystem, an administrative unit rather than a physical one, a fact that does not disqualify it, but does determine how it may be used.

4.5.3. *Conditions for legitimate comparison*

Two IPMM figures (from different operators, or from one operator across a period in which the mission mix changed) are comparable only if all four conditions hold. Layer 6 records the outcome of each in the claim's conversion record.

- **CC1: Mechanism–unit match.** The dominant mechanism is distance-driven, per the Layer 4 physics-of-failure classification. Where it is time- or cycle-driven, the per-hour or per-cycle rate is the invariant and the per-mile figure is not the comparable quantity.
- **CC2: Duty-cycle equivalence.** The fleets' Layer 3 exposure profiles agree, within a stated tolerance, on every covariate Layer 4 found significant: at minimum $\bar{v}$, idle fraction, and cycles per mile.
- **CC3: Denominator congruence.** Both fleets compute mileage identically: the same boundary between operating and parked, the same treatment of deadhead and repositioning miles, and the same treatment of teleoperated and safety-driver periods. This condition is routinely violated in current public AV reporting, and is a substantial part of why California DMV disengagement figures resist cross-operator comparison.
- **CC4: Population congruence.** The comparison is between the same hardware revision and a comparable software-version distribution, or is stratified so that mixtures are not being compared.

Where CC1 fails but CC2–CC4 hold, the comparison remains legitimate in per-hour units. This is the practical reason Layer 3 tracks paired exposure: the framework can always retreat to the invariant unit rather than abandoning the comparison.

4.5.4. *Two quantities, not one*

Conflating two distinct things under one name is most of the difficulty. The framework separates them, and requires both to be labelled.

- **Operational IPMM**: $\lambda_h \cdot 10^6 / \bar{v}$ at the fleet's own observed speed. This is the operator's actual per-mile incidence, and it is the right input to spares provisioning, depot capacity, cost-per-mile, and rider-facing availability. It is **not** comparable across operators.

- **Normalized IPMM @ $\bar{v}$_ref**: λ_h · 10^6 / $\bar{v}$_ref at a declared reference speed. This is comparable across operators and across time for a time-driven mechanism, and it is the right input to regulatory benchmarking and sector-level oversight. It is **not** the operator's actual experience, and should never be presented as such.

The two coincide only when $\bar{v} = \bar{v}$_ref. Figure 5(B) shows the collapse: four archetypes whose operational IPMM spans 4.6× report an identical 2.60×10^{-2} once normalized to 25 mph. Choosing $\bar{v}$_ref is a governance decision rather than a technical one; what the framework requires is that it be declared, held stable, and recorded with every claim.

4.5.5. *What the framework contributes*

The claim that the framework supplies a common reporting unit is too strong, and this section withdraws it. A single number cannot serve a functional-safety reviewer, a fleet operations planner, and a regulator, because those three audiences are asking different questions of it.

What the framework does contribute is narrower and more defensible: **every reliability claim it emits carries an explicit, auditable conversion record** giving the per-hour rate, the per-mile rate, $\bar{v}$ and its provenance, the mechanism's native denominator from Layer 4, the reference speed if normalized, and the CC1–CC4 outcomes. That record is what makes the per-mile figure interpretable rather than merely available, and it is what allows a reader to determine whether a given comparison is admissible. Producing both units is the easy half; carrying the conditions under which each may be used is the contribution.

This also sharpens the recommendation to regulators in section 5.2. A disclosure schema that requires operational IPMM alone institutionalizes a metric an operator can improve by 44 % through raising average speed from 25 to 45 mph, with no change to any component. A schema that requires the per-hour rate, $\bar{v}$, and the mileage-denominator definition alongside it cannot be gamed in that way, and costs the operator three additional fields.

5. Conclusion & Call to Action

5.1. Summary

This paper presented an eight-layer fleet-based reliability framework targeting the low-voltage electrical subsystem of autonomous vehicles, covering energy storage and BMS, DC-DC conversion, ECUs and the in-vehicle communication architecture, wire harnesses and connectors, and sensor power. The framework is source-agnostic, standards-aligned (ISO 26262, ISO 21448, UL 4600), and emits quantitative reliability metrics, the assurance-case artefacts that consume them, and an AV-fleet-specific containment ladder that exploits centralized dispatch and continuous OTA to act faster and more finely than traditional passenger-vehicle recall workflows allow. Validity is established by completeness against the standards, compatibility with existing engineering artefacts, and explicit closure of gaps in the prior literature.

5.2. *Adoption pathway*

AV programmes: OEMs and robotaxi operators. Instantiate Layers 1–8 on internal fleet telemetry. Use Layer 6 outputs to populate UL 4600 cases and the random-hardware-failure portion of ISO 26262 hardware safety arguments. Report results in both PMHF (for functional-safety reviewers) and IPMM (for fleet operations and external communications). Adopt the Layer 7 containment workflow (Containment Review Board, AV-specific containment ladder L7.0 through L7.7) to act on field signals in hours-to-days rather than weeks-to-months. Adopt the Layer 8 governance discipline so that fleet evidence drives design changes, supplier escalations, and OTA roll-outs on an explicit cadence rather than ad hoc.

Tier-1 suppliers and component manufacturers. Publish FIT and accelerated-life data in formats directly consumable by Layer 4 priors. Participate in Layer 7 containment and 8D workflows so that field-incident response is fast, and in Layer 8 supplier-feedback flows so that field IPMM trends inform component-design iteration. Where datasheets carry only bench-test FIT, append the duty-cycle assumptions used so that Layer 3 can translate them into AV-equivalent exposure terms.

Regulators: NHTSA, California DMV, EU type-approval authorities. Require operators to publish a minimal disclosure schema (component, mileage-at-failure, **powered-hours-at-failure**, mission phase, fault code, severity classification), together with the three fields that make a per-mile figure interpretable: **fleet-average speed over powered hours, the mileage-denominator definition, and the mechanism's native denominator**. Without them, a per-mile disclosure rewards duty-cycle change over engineering: an operator improves reported IPMM by 44 % by raising average speed from 25 to 45 mph while every component behaves identically (section 4.5). With them, the schema supports Layer 2 ingestion, aggregate-level Layer 4 estimation, and legitimate cross-operator comparison in the normalized form of section 4.5.4. Existing California DMV disengagement reports could be extended with an LV-electrical root-cause field at minimal disclosure cost.

Standards bodies: UL, ISO, SAE. Consider an annex or technical report defining the Layer 6 PMHF↔IPMM mapping **including a sector reference speed and the CC1–CC4 comparability conditions**, GSN fragment templates for LV electrical subsystems, and evidence-sufficiency criteria for fleet-derived random-hardware-failure rates as recommended practice. This would lower the integration friction between fleet operators and

assurance practitioners and make safety cases auditable across organizations.

Researchers. Empirically validate the framework on the first available AV LV fleet datasets, which are likely to be OEM-internal in the near term. Refine Layer 4 covariate sets for AV-specific failure modes (compute-power-spike transients, sensor-supply ripple under perception load, redundant-supply changeover events) that have no analogue in conventional-vehicle warranty data. Develop physics-of-failure models for connectors, harnesses, and Automotive Ethernet PHYs under automotive duty cycles; these failure modes are well known but their physics is under-modelled in the open literature. Build open-source reference implementations of Layers 2–4 on public proxy data (NHTSA EWR, https://www.nhtsa.gov/nhtsa-datasets-and-apis; NREL FleetDNA, https://data.nrel.gov/submissions/68; ORNL ROAD [48]) to lower the adoption barrier for the wider community.

5.3. Future work

- Empirical validation on a real AV LV fleet dataset, when accessible.
- Quantification of AV-specific failure modes (compute power-spike transients, sensor-supply ripple, redundant-supply changeover).
- Open-source reference implementation of the Layer 2–4 pipeline on public proxy data (NHTSA EWR + NREL FleetDNA + ORNL ROAD) to lower the adoption barrier.
- Extension of the section 3.7 fragment from a single element branch to a complete SG-LV-01 case covering every element group in the Layer 1 register, and independent review of the ES1–ES5 sufficiency criterion by assurance-case practitioners against a live certification submission.
- Extension to high-voltage and mechatronic subsystems following the same layer structure.

5.4. Closing remarks

Reliable low-voltage electrical power is the quiet precondition for everything an autonomous vehicle does. Perception, planning, and control attract the research attention, but none of them runs for a moment without a healthy 12 V supply, a stable DC-DC stage, an uncorrupted communication bus, and sound connectors and grounds. As autonomous fleets scale from pilots to public infrastructure (robotaxis, autonomous freight, and shared-mobility services that millions may one day depend on daily), the dependability of this unglamorous electrical layer becomes a matter of public safety, operational viability, and societal trust rather than of engineering convenience alone. Unreliable LV electronics translate directly into stranded vehicles, blocked intersections, eroded rider confidence, and regulatory scrutiny that can stall an entire sector; demonstrable, fleet-evidenced reliability, by contrast, lowers operating cost, shortens the path to regulatory approval, and earns the public license to operate on which autonomous mobility ultimately depends. By turning the continuous stream of fleet evidence into auditable, standards-aligned assurance, the framework presented here aims to make that reliability not merely asserted but measured, defended, and sustained, and in doing so to help the autonomous-vehicle industry deliver on its promise of safer, more accessible, and more efficient transportation.

Conflict of Interest

The author declares no conflict of interest.

References

[1] *Road vehicles — Functional safety*, ISO Standard 26262, 2018.

[2] S. Wu, "Warranty data analysis: A review," *Quality and Reliability Engineering International*, vol. 28, no. 8, pp. 795–805, 2012, doi:10.1002/qre.1282.

[3] L. Attardi, M. Guida, G. Pulcini, "A mixed-Weibull regression model for the analysis of automotive warranty data," *Reliability Engineering & System Safety*, vol. 87, no. 2, pp. 265–273, 2005, doi:10.1016/j.ress.2004.05.003.

[4] *Road vehicles — Safety of the intended functionality*, ISO Standard 21448, 2022.

[5] *Standard for Safety for the Evaluation of Autonomous Products*, UL 4600, Underwriters Laboratories, 2020.

[6] R. Mariani, "An overview of autonomous vehicles safety," in *Proc. IEEE Int. Reliability Physics Symp. (IRPS)*, 2018, doi:10.1109/IRPS.2018.8353618.

[7] W. M. D. Chia, S. L. Keoh, C. Goh, C. Johnson, "Risk assessment methodologies for autonomous driving: A survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 10, pp. 16923–16939, 2022, doi:10.1109/TITS.2022.3163747.

[8] K. D. Majeske, "A mixture model for automobile warranty data," *Reliability Engineering & System Safety*, vol. 81, no. 1, pp. 71–77, 2003, doi:10.1016/S0951-8320(03)00073-5.

[9] S. Wu, "A review on coarse warranty data and analysis," *Reliability Engineering & System Safety*, vol. 114, pp. 1–11, 2013, doi:10.1016/j.ress.2012.12.021.

[10] R. Khoshkangini, P. Sheikholharam Mashhadi, P. Berck, S. Gholami Shahbandi, S. Nowaczyk, T. Niklasson, T. Jonsson, "Early prediction of quality issues in automotive modern industry," *Information*, vol. 11, no. 7, p. 354, 2020, doi:10.3390/info11070354.

[11] R. Khoshkangini, M. Tajgardan, J. Lundström, M. Rabbani, D. Tegnered, "A snapshot-stacked ensemble and optimization approach for vehicle breakdown prediction," *Sensors*, vol. 23, no. 12, p. 5621, 2023, doi:10.3390/s23125621.

[12] M. A. Gosavi, B. B. Rhoades, "Application of functional safety in autonomous vehicles using ISO 26262 standard: A survey," in *Proc. IEEE SoutheastCon 2018*, 2018, doi:10.1109/SECON.2018.8479057.

[13] P. Sinha, "Architectural design and reliability analysis of a fail-operational brake-by-wire system from ISO 26262 perspectives," *Reliability Engineering & System Safety*, vol. 96, no. 10, pp. 1349–1359, 2011, doi:10.1016/j.ress.2011.03.013.

[14] A. Schnellbach, G. Grießnig, "Development of the ISO 21448," in *Systems, Software and Services Process Improvement, EuroSPI 2019*, Springer CCIS, vol. 1060, 2019, doi:10.1007/978-3-030-28005-5_45.

[15] D. Kinalzyk, "SOTIF process and methods in combination with functional safety," in *Systems, Software and Services Process Improvement, EuroSPI 2021*, Springer CCIS, vol. 1488, 2021, doi:10.1007/978-3-030-85521-5_41.

[16] *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*, SAE J3016_202104, SAE International, 2021.

[17] M. Chelouati, A. Boussif, J. Beugin, E.-M. El Koursi, "Graphical safety assurance case using Goal Structuring Notation (GSN) — challenges, opportunities and a framework for autonomous trains," *Reliability Engineering & System Safety*, vol. 230, p. 108933, 2023, doi:10.1016/j.ress.2022.108933.

[18] Y. Dong, W. Huang, V. Bharti, V. Cox, A. Banks, S. Wang, X. Zhao, S. Schewe, X. Huang, "Reliability assessment and safety arguments for machine learning components in system assurance," *ACM Transactions on Embedded Computing Systems*, vol. 22, no. 3, pp. 1–48, 2023, doi:10.1145/3570918.

[19] M. Ahsan, S. Stoyanov, C. Bailey, "Prognostics of automotive electronics with data driven approach: A review," in *Proc. 39th Int. Spring Seminar on Electronics Technology (ISSE)*, 2016, doi:10.1109/ISSE.2016.7563205.

[20] A. Prisacaru, P. J. Gromala, M. B. Jeronimo, B. Han, G. Q. Zhang, "Prognostics and health monitoring of electronic system: A review," in *Proc. 18th Int. Conf. Thermal, Mechanical and Multi-Physics Simulation and Experiments in Microelectronics and Microsystems (EuroSimE)*, 2017, doi:10.1109/EuroSimE.2017.7926248.

[21] A. Prisacaru, P. J. Gromala, B. Han, G. Q. Zhang, "Degradation estimation and prediction of electronic packages using data-driven approach," *IEEE Transactions on Industrial Electronics*, vol. 69, no. 3, pp. 2996–3006, 2022, doi:10.1109/TIE.2021.3068681.

[22] A. Kleyner, A. S. S. Vasan, M. Pecht, "A new application for failure prognostics — reduction of automotive electronics reliability test duration," in *Annual Conference of the PHM Society*, vol. 9, no. 1, 2017, doi:10.36001/phmconf.2017.v9i1.2446.

[23] M. S. H. Lipu, S. Ansari, M. S. Miah, K. Hasan, S. T. Meraj, M. Faisal, T. Jamal, S. H. M. Ali, A. Hussain, M. M. Naser et al., "Deep learning enabled state of charge, state of health and remaining useful life estimation for smart battery management system: Methods, implementations, issues and prospects," *Journal of Energy Storage*, vol. 55, p. 105752, 2022, doi:10.1016/j.est.2022.105752.

[24] A. Aitio, D. A. Howey, "Predicting battery end of life from solar off-grid system field data using machine learning," *Joule*, vol. 5, no. 12, pp. 3204–3220, 2021, doi:10.1016/j.joule.2021.11.006.

[25] D. Astigarraga, F. M. Ibáñez, A. Galarza, J. M. Echeverria, I. Unanue, P. Baraldi, E. Zio, "Analysis of the results of accelerated aging tests in insulated gate bipolar transistors," *IEEE Transactions on Power Electronics*, vol. 31, no. 11, pp. 7953–7962, 2016, doi:10.1109/TPEL.2015.2512923.

[26] B. Ji, V. Pickert, W. Cao, B. Zahawi, "In situ diagnostics and prognostics of wire bonding faults in IGBT modules for electric vehicle drives," *IEEE Transactions on Power Electronics*, vol. 28, no. 12, pp. 5568–5577, 2013, doi:10.1109/TPEL.2013.2251358.

[27] A. Forouzandeh Shahraki, S. Al-Dahidi, A. R. Taleqani, O. P. Yadav, "Using LSTM neural network to predict remaining useful life of electrolytic capacitors in dynamic operating conditions," *Proceedings of the Institution of Mechanical Engineers, Part O: Journal of Risk and Reliability*, vol. 236, no. 6, pp. 1069–1080, 2022, doi:10.1177/1748006X221087503.

[28] S.-F. Lokman, A. T. Othman, M.-H. Abu-Bakar, "Intrusion detection system for automotive Controller Area Network (CAN) bus system: A review," *EURASIP Journal on Wireless Communications and Networking*, vol. 2019, no. 184, 2019, doi:10.1186/s13638-019-1484-3.

[29] S. Rajapaksha, H. Kalutarage, M. O. Al-Kadri, A. Petrovski, G. Madzudzo, M. Cheah, "AI-based intrusion detection systems for in-vehicle networks: A survey," *ACM Computing Surveys*, vol. 55, no. 11, pp. 1–40, 2023, doi:10.1145/3570954.

[30] A. Uçar, M. Karaköse, N. Kırımça, "Artificial intelligence for predictive maintenance applications: Key components, trustworthiness, and future trends," *Applied Sciences*, vol. 14, no. 2, p. 898, 2024, doi:10.3390/app14020898.

[31] J. Maier, H.-C. Reuss, "Design of zonal E/E architectures in vehicles using a coupled approach of k-means clustering and Dijkstra's algorithm," *Energies*, vol. 16, no. 19, p. 6884, 2023, doi:10.3390/en16196884.

[32] M. Gleißner, M.-M. Bakran, "Reliable & fault-tolerant DC/DC-converter structures," in *Proc. PCIM Europe 2014*, 2014.

[33] H. Wang, M. Liserre, F. Blaabjerg, "Toward reliable power electronics: Challenges, design tools, and opportunities," *IEEE Industrial Electronics Magazine*, vol. 7, no. 2, pp. 17–26, 2013, doi:10.1109/MIE.2013.2252958.

[34] H. Wang, M. Liserre, F. Blaabjerg, P. de Place Rimmen, J. B. Jacobsen, T. Kvisgaard, J. Landkildehus, "Transitioning to physics-of-failure as a reliability driver in power electronics," *IEEE Journal of Emerging and Selected Topics in Power Electronics*, vol. 2, no. 1, pp. 97–114, 2014, doi:10.1109/JESTPE.2013.2290282.

[35] A. Gupta, R. Ayyanar, S. Chakraborty, "Novel electric vehicle traction architecture with 48 V battery and multi-input, high conversion ratio converter," *IEEE Open Journal of Vehicular Technology*, vol. 2, pp. 448–463, 2021, doi:10.1109/OJVT.2021.3132281.

[36] A. Lahyani, P. Venet, A. Guermazi, A. Troudi, "Battery/supercapacitors combination in uninterruptible power supply (UPS)," *IEEE Transactions on Power Electronics*, vol. 28, no. 4, pp. 1509–1522, 2013, doi:10.1109/TPEL.2012.2210736.

[37] H. A. Gabbar, A. M. Othman, M. R. Abdussami, "Review of battery management systems (BMS) development and industrial standards," *Technologies*, vol. 9, no. 2, p. 28, 2021, doi:10.3390/technologies9020028.

[38] M. Buffolo, D. Favero, A. Marcuzzi, C. De Santi, G. Meneghesso, E. Zanoni, M. Meneghini, "Review and outlook on GaN and SiC power devices: Industrial state-of-the-art, applications, and perspectives," *IEEE Transactions on Electron Devices*, vol. 71, no. 3, pp. 1344–1355, 2024, doi:10.1109/TED.2023.3346369.

[39] G. Tshagharyan, G. Harutyunyan, S. Shoukourian, Y. Zorian, "Modeling and testing of aging faults in FinFET memories for automotive applications," in *Proc. IEEE International Test Conference (ITC)*, 2018, doi:10.1109/TEST.2018.8624890.

[40] A. Richter, T. Puig Walz, A. Marko, M. Rexer, F. Schnabel, "Components and their failure rates in autonomous driving," in *Proc. 33rd European Safety and Reliability Conference (ESREL 2023)*, 2023, doi:10.3850/978-981-18-8071-1_P398-cd.

[41] *Failure Rates of Components — Expected Values, General*, SN 29500, Siemens AG, 2018.

[42] *Reliability data handbook — Universal model for reliability prediction of electronics components, PCBs and equipment*, IEC Standard 62380, 2004.

[43] *Failure Mechanisms and Models for Semiconductor Devices*, JEDEC Publication JEP122H, JEDEC Solid State Technology Association, 2016.

[44] J. W. McPherson, H. C. Mogul, "Underlying physics of the thermochemical E model in describing low-field time-dependent dielectric breakdown in SiO_2 thin films," *Journal of Applied Physics*, vol. 84, no. 3, pp. 1513–1523, 1998, doi:10.1063/1.368217.

[45] Assurance Case Working Group, "Goal Structuring Notation Community Standard, Version 3," SCSC-141C, Safety-Critical Systems Club, 2021.

[46] T. Kelly, R. Weaver, "The Goal Structuring Notation — a safety argument notation," in *Proc. Dependable Systems and Networks Workshop on Assurance Cases*, 2004.

[47] W. Q. Meeker, L. A. Escobar, F. G. Pascual, *Statistical Methods for Reliability Data*, 2nd ed. Hoboken, NJ, USA: Wiley, 2021.

[48] M. E. Verma, R. A. Bridges, M. D. Iannacone, S. C. Hollifield, P. Moriano, S. C. Hespeler, B. Kay, F. L. Combs, "A comprehensive guide to CAN IDS data and introduction of the ROAD dataset," *PLOS ONE*, vol. 19, no. 1, p. e0296879, 2024, doi:10.1371/journal.pone.0296879.

SHRIKANT PAWAR has done his bachelor's degree from Shivaji University in 2022 in Mechanical Engineering. He has done his master's degree in Robotics Engineering from University of Michigan Dearborn in 2024. He published his research on Automotive Radar Sensor design and validation. He currently works as a reliability engineer specializing in low-voltage electrical systems for production autonomous vehicle fleets.