



VEYRUM ROBOTICS STANDARD



VRS-GEN-102

IR Category: Reliability and Uptime

Scoring basis, reliability criteria RL-1...RL-8, evidence by tier and determination procedure for the reliability-and-uptime intrinsic category.

VEYRUM RESEARCH INSTITUTE

DOCUMENT

VRS-GEN-102

VERSION

1.0

STATUS

Draft

DATE

15 September 2026

Contents

1 Introduction

2 1 Scope

3 2 Normative references

4 3 Terms and definitions

5 4 What this category measures

6 5 Scoring basis

7 6 Criteria

7.1 6.1 RL-1 Rated life declared with basis (0–5)

7.2 6.2 RL-2 Environmental and durability test evidence (0–10)

7.3 6.3 RL-3 Field failure rate (0–12)

7.4 6.4 RL-4 Availability under duty (0–10)

7.5 6.5 RL-5 Fault tolerance and graceful degradation (0–8)

7.6 6.6 RL-6 In-session endurance sample (0–9)

7.7 6.7 RL-7 Software release quality (0–6)

7.8 6.8 RL-8 Warranty terms and claim rate (0–6)

8 7 Evidence by tier

9 8 Determination procedure

10 9 Worked example (fictional model)

11 10 Assessor checklist

12 Bibliography

13 Change history

Foreword. IR Category Standard for **reliability and uptime** (category weights live in the VRS-GEN-005 master and its data twin, never here). This edition scores the category as a **sum of integer points** across eight reliability criteria (RL-1...RL-8, 66 points; Clause 6), each earned against a plain-language ladder and bounded by an evidence tier — Unverified, Verified or Certified (Clause 7). It supersedes the rate-based model of the 0.x drafts: no availability/failure-rate formula, no statistical blend toward a class median, no qualitative bands. Requirements use “shall”.

1 Introduction

Reliability is the loss driver an insurer reads as claim *frequency* and a lender reads as *revenue continuity*. A robot that stops working does not merely fail its task: it triggers repair claims, business-interruption exposure, and — when a software release regresses a whole fleet at once — a correlated loss that behaves like a catastrophe rather than a series of independent events. This Standard turns “how dependable is this model?” into a repeatable number by asking eight questions — is its rated life declared with a basis, has it been durability-tested, how often does it fail in the field, is it available under real duty, does it fail gracefully, does it endure a working session, does its software change safely, and what do its warranty terms and claims say — and awarding integer points for each. The points sum to a category total of 0–66. What a robot can earn is bounded by how well its evidence is substantiated: public documents alone reach the **Unverified** ceiling, a manufacturer evidence pack reaches the **Verified** ceiling, and a witnessed demonstration reaches the **Certified** ceiling. The category rewards dependability that is designed in *and* shown, not merely asserted, and it reuses the data operators already keep (SEMI E10 equipment states, ISO 14224 maintenance records) so that adoption is a mapping, not new testing.

2 1 Scope

1.1 This Standard defines, for the **reliability-and-uptime** intrinsic category: what the category measures (Clause 4); the scoring basis (Clause 5); the eight criteria RL-1...RL-8 with their point ladders and tier ceilings (Clause 6); the evidence admissible at each tier (Clause 7); the determination procedure (Clause 8); a worked example (Clause 9); and the assessor checklist (Clause 10).

1.2 The category shall be assessed at one of three **tiers** — Unverified, Verified, Certified — and each criterion states the maximum points reachable at each tier. The category tier is the tier whose evidence bar the assessed evidence meets (Clause 8). An Application Rating tier shall not exceed the Intrinsic Rating tier.

1.3 This Standard excludes category weighting and aggregation into the Intrinsic Rating (VRS-GEN-005 §6), sector re-weighting (the sector **-201** protocols), the evidence-grade definitions (VRS-GEN-012), and repair-network / parts-logistics depth, which is measured by serviceability (VRS-GEN-103); mean-time-to-repair enters this category only through availability (RL-4). It does not itself assign the Intrinsic Rating.

3 Normative references

- **VRS-GEN-001**, Vocabulary and Terminology — robot and dependability terms.
- **VRS-GEN-005**, Intrinsic Rating — Methodology — point aggregation, tiers, cold-start parity and the company-wide fallback rule (§6).
- **VRS-GEN-009**, Robot Risk Passport and Registry Schema — the record that carries the category result (Clause 8).
- **VRS-GEN-012**, Evidence Grades and Data Requirements — evidence kinds and the tier evidence bars.
- **IEC 60050-192:2015**, International Electrotechnical Vocabulary — Part 192: Dependability (availability, failure rate, MTBF, MTTR definitions).
- **IEC 61703:2016**, Mathematical expressions for reliability, availability, maintainability and maintenance support terms.
- **ISO 14224:2016** (corrected 2016-10-01), Collection and exchange of reliability and maintenance data for equipment — failure-mode taxonomy and the reused maintenance-record evidence lane.
- **IEC 60529:1989+A2:2013**, Degrees of protection provided by enclosures (IP code) — publicly verifiable environmental-ingress ratings (RL-2).
- **SEMI E10 (2021 RAM revision)**, Definition and Measurement of Equipment Reliability, Availability and Maintainability — six-state equipment model behind the availability criterion (RL-4).

4 Terms and definitions

Terms per VRS-GEN-001 and IEC 60050-192:2015 (item, failure, fault, availability, failure rate, MTBF, MTTR). Locally:

- **op-hours** — operating hours: time during which the robot is powered and able to move or act under its own control, summed across the fleet in scope.
- **availability** — the fraction of scheduled operating time in which the model was able to perform its function, read from the SEMI E10 six-state equipment model (RL-4; IEC 61703:2016 operational availability).
- **critical failure** — a failure that stops the robot's primary function; a **degraded failure** reduces the function without stopping it; an **intermittent** fault is self-clearing and recurring; a **cosmetic** condition has no functional effect and is excluded from the failure count.
- **software-update regression** — a released software update that introduced, on the field fleet, a new failure mode or an availability loss not present on the prior version.
- **tier** — the evidence stringency at which the category is assessed: **Unverified** (public evidence only), **Verified** (manufacturer evidence pack examined), **Certified** (behaviour witnessed).
- **ceiling** — the maximum points a criterion can award at a given tier; points earned are capped at the ceiling of the assessed tier.
- **evidence bar** — the minimum evidence a criterion requires before any points are awarded; below it the criterion scores 0.
- **company-wide fallback** — where model-specific evidence is unavailable, the substitute points a criterion allows from manufacturer-wide evidence (VRS-GEN-005 §6); marked in the criterion row.

- **certificate rung** — a ladder step earned by a publicly verifiable third-party certificate or mark (registrar entry, certificate number, or IP class mark per IEC 60529) in lieu of a witnessed demonstration.

5 4 What this category measures

4.1 The category measures the model's **failure behaviour, availability and durability under real duty**, severity-differentiated (critical, degraded and intermittent failures; cosmetic conditions excluded) and bounded by the evidence that substantiates each claim.

4.2 It measures both **design and evidence and record**: whether rated life and durability tests are declared and shown (RL-1, RL-2), what the field failure rate and availability actually are (RL-3, RL-4), whether the machine fails gracefully and endures a working session (RL-5, RL-6), and how safely its software changes and what its warranty terms and claim rate say (RL-7, RL-8). A criterion is scored only to the tier its evidence supports.

4.3 The category shall be assessed at **model** granularity (manufacturer + model + declared major variant). Records shall not be pooled across models that differ in reliability-relevant hardware or in the reliability-relevant software baseline unless the manufacturer declares them one model for reliability purposes and evidence supports that declaration.

6 5 Scoring basis

5.1 The category score shall be the **sum of integer points** awarded across the eight criteria RL-1...RL-8 (Clause 6). The maximum is 66 points. There is no rate formula and no statistical blend; the point range of each criterion sizes its importance. Aggregation of this total into the Intrinsic Rating, tier rules and the cold-start parity and company-wide fallback rules are defined once in VRS-GEN-005 §6 and are not restated here.

5.2 Each criterion awards points against a **ladder** of rungs, highest first, decidable by inspection, measurement, test or documented evidence (VRS-GEN-002 §7.1). The assessor shall award the highest rung the evidence supports.

5.3 Each criterion states three **ceilings** — Unverified, Verified, Certified. Points earned shall be capped at the ceiling of the tier at which the criterion's evidence is assessed. A criterion may reach a higher tier than another; the category tier is recorded per Clause 8.

5.4 Each criterion states an **evidence bar**. Where the bar is not met the criterion scores **0**; no evidence is 0 points, never a null or a discard.

5.5 A criterion marked **not applicable by class** for the model's form (Clause 6 Applicability row) is excluded from the score and from the totals; the reachable maximum is reduced accordingly and recorded (Clause 8). For the reliability category all eight criteria apply to all classes.

5.6 Where a criterion offers a **company-wide fallback**, manufacturer-wide evidence may earn the stated fallback points when model-specific evidence is unavailable (VRS-GEN-005 §6); the substitution shall be recorded. Numeric thresholds inside a rung that carry (**prior**) are launch-edition values that change only through the calibration data twin, never in this document.

7 6 Criteria

6.0 Summary. Points and tier ceilings per criterion; totals reproduce by addition.

§	CRITERION	POINTS	UNVERIFIED CEILING	VERIFIED CEILING	CERTIFIED CEILING
6.1	RL-1 Rated life declared with basis	0–5	5	5	5
6.2	RL-2 Environmental and durability test evidence	0–10	6	8	10
6.3	RL-3 Field failure rate	0–12	6	12	12
6.4	RL-4 Availability under duty	0–10	5	10	10
6.5	RL-5 Fault tolerance and graceful degradation	0–8	5	5	8
6.6	RL-6 In-session endurance sample	0–9	0	4	9
6.7	RL-7 Software release quality	0–6	5	6	6
6.8	RL-8 Warranty terms and claim rate	0–6	5	6	6
Total		66	37	56	66

7.1 6.1 RL-1 Rated life declared with basis (0–5)

ROW	
What is measured	Design/rated life per major subsystem (drive, battery, actuators, sensors) with the test or analysis basis.
Ladder	5 = all subsystems with basis · 3 = headline life only · 0 = none
Unverified — evidence & ceiling	Datasheet. Ceiling 5.
Verified — evidence & ceiling	Rated-life table with basis. Ceiling 5.
Certified — evidence & ceiling	From Verified evidence. Ceiling 5.
Company-wide fallback	n/a
Evidence bar	None.
Applicability	All classes

7.2 6.2 RL-2 Environmental and durability test evidence (0–10)

ROW	
What is measured	Tests to stated standards for the declared operating envelope (IP, temperature, vibration/shock, duty-cycle endurance).
Ladder	10 = all four with reports · 8 = pack-evidenced (documents), not demonstrated · 7 = three · 6 = ratings certified by a publicly verifiable third-party mark/report (e.g. IP per IEC 60529, class mark) · 4 = ratings claimed, one report · 2 = ratings claimed, no report · 0 = none
Unverified — evidence & ceiling	IP/temperature ratings on datasheet; public certificate. Ceiling 6.
Verified — evidence & ceiling	Test reports. Ceiling 8.
Certified — evidence & ceiling	Operation in a declared environmental condition present at the site (sample) ; otherwise. Ceiling 10.
Company-wide fallback	n/a
Evidence bar	Operating envelope declared.
Applicability	All classes

7.3 6.3 RL-3 Field failure rate (0–12)

ROW	
What is measured	Class-weighted failures per 1,000 operating hours over the window.
Ladder	$12 = \leq 0.5 \cdot 9 = \leq 1 \cdot 6 = \leq 2 \cdot 3 = \leq 4 \cdot 0 = > 4$ or not established. (Thresholds: launch-edition priors.) Regime rule (F19): full Unverified ceiling where public malfunction reporting + disclosed exposure exist (e.g. MAUDE + procedure volumes)
Unverified — evidence & ceiling	Public reviews/operator forums; public malfunction databases (12 under a public regime). Ceiling 6.
Verified — evidence & ceiling	Model failure record + exposure. Ceiling 12.
Certified — evidence & ceiling	From Verified evidence. Ceiling 12.
Company-wide fallback	Company-wide → 9
Evidence bar	Exposure hours declared (model or company).
Applicability	All classes

7.4 6.4 RL-4 Availability under duty (0–10)

ROW	
What is measured	Availability per SEMI E10 states over the window versus declared.
Ladder	10 = ≥ 98 % (prior) and ≥ declared · 7 = ≥ 95 % (prior) · 5 = publicly evidenced only (no pack, no demonstration) · 4 = ≥ 90 % (prior) · 0 = < 90 % (prior) or not established
Unverified — evidence & ceiling	— (only if published by operators). Ceiling 5.
Verified — evidence & ceiling	Downtime/state record. Ceiling 10.
Certified — evidence & ceiling	From Verified evidence. Ceiling 10.
Company-wide fallback	Company-wide → 7
Evidence bar	Scheduled time and downtime declared.
Applicability	All classes

7.5 6.5 RL-5 Fault tolerance and graceful degradation (0–8)

ROW	
What is measured	On a component fault the robot enters the declared degraded or safe behaviour rather than an uncontrolled state; FMEA exists.
Ladder	8 = witnessed for declared faults · 5 = FMEA + test evidence, or fault-response behaviour independently tested/certified · 2 = documented only · 0 = none
Unverified — evidence & ceiling	Manual; public third-party tests. Ceiling 5.
Verified — evidence & ceiling	FMEA / fault-response test evidence. Ceiling 5.
Certified — evidence & ceiling	Induced faults (manufacturer's procedure) → declared behaviour. Ceiling 8.
Company-wide fallback	n/a
Evidence bar	Fault-response behaviour documented.
Applicability	All classes

7.6 6.6 RL-6 In-session endurance sample (0–9)

ROW	
What is measured	≥ 60 min or ≥ N declared cycles of representative work in session without a fault or intervention; counters on feed.
Ladder	9 = clean · 4 = one self-clearing fault · 0 = intervention needed
Unverified — evidence & ceiling	— (no public path; F4). Ceiling 0.
Verified — evidence & ceiling	Duty-cycle test report. Ceiling 4.
Certified — evidence & ceiling	Witnessed run (may span other steps). Ceiling 9.
Company-wide fallback	n/a
Evidence bar	Representative task declared.
Applicability	All classes

7.7 6.7 RL-7 Software release quality (0–6)

ROW	
What is measured	Regressions per release in window; rollback available.
Ladder	6 = no regressions, rollback shown · 5 = no known regressions in the public record, rollback documented · 4 = ≤ 1 regression, rollback documented · 3 = release notes + rollback documented, regressions unknown · 2 = regressions, no rollback · 0 = unknown
Unverified — evidence & ceiling	Public release notes; public advisories. Ceiling 5.
Verified — evidence & ceiling	Release history + regression record. Ceiling 6.
Certified — evidence & ceiling	Rollback performed on console. Ceiling 6.
Company-wide fallback	n/a
Evidence bar	Release history for window.
Applicability	All classes

7.8 6.8 RL-8 Warranty terms and claim rate (0–6)

ROW	
What is measured	Warranty duration/coverage relative to declared life; claims per unit-year.
Ladder	6 = ≥ 2 yr comprehensive, low claim rate · 5 = publicly evidenced only (no pack, no demonstration) · 4 = standard terms, claim rate disclosed · 2 = terms only · 0 = none
Unverified — evidence & ceiling	Published terms. Ceiling 5.
Verified — evidence & ceiling	Terms + claim-rate figures. Ceiling 6.
Certified — evidence & ceiling	From Verified evidence. Ceiling 6.
Company-wide fallback	Company-wide claim rate → 5
Evidence bar	Terms published.
Applicability	All classes

8 7 Evidence by tier

7.1 Unverified. The assessor works only from evidence anyone can obtain without the manufacturer's cooperation: datasheets and manuals (rated life, IP and temperature ratings), published declarations of conformity and their cited standards, public certificate registers and IP class marks, public malfunction, recall and incident databases and verified operator reviews or forums (field failure rate), operator-published availability figures where they exist, published release notes and advisories, and published warranty terms. Each criterion's Unverified ceiling is the most an honest reading of such evidence can earn; the category tier is **Unverified**. RL-6 has no public path and its Unverified ceiling is 0 (F4).

7.2 Verified. The assessor additionally examines a manufacturer evidence pack: the rated-life table with its basis, environmental and durability test reports, the model failure record with its exposure hours, the downtime and SEMI E10 state record, the FMEA and fault-response test evidence, the duty-cycle test report, the software release history with its regression record, and the warranty terms with claim-rate figures. Verified ceilings apply only to criteria whose pack items are present and examined; the category tier is **Verified**.

7.3 Certified. The assessor witnesses the declared behaviour: operation in a declared environmental condition present at the site (RL-2), induced faults resolving to the declared degraded or safe behaviour per the manufacturer's procedure (RL-5), an in-session endurance run against the declared representative task (RL-6), and a rollback performed on the console (RL-7). Documentary and record

criteria (RL-1, RL-3, RL-4, RL-8) carry their Certified points from the Verified evidence. The category tier is **Certified**.

7.4 Certificate rung. Where a criterion names a certificate or public-mark rung (RL-2, RL-5), a publicly verifiable third-party mark or report — an IP rating per IEC 60529, an accredited test report, or a fault-response certificate identified by registrar entry or certificate number — earns the stated rung without a witnessed demonstration; an unverifiable mark shall not.

7.5 Field-record regime rule. For RL-3 the reachable Unverified ceiling depends on the reporting regime the model operates under. Where a **mandatory public malfunction-reporting regime** covers the model with disclosed exposure — declared by the sector protocol (e.g. medical device reports via MAUDE with procedure volumes) — the assessor may treat the public record as sufficient and the Unverified ceiling is the full 12. Where no such regime applies, the public reviews-and-forums reading is bounded by the criterion's stated Unverified ceiling.

7.6 Company-wide fallback. Where a criterion offers a company-wide fallback (RL-3 → 9, RL-4 → 7, RL-8 → 5) and model-specific evidence is unavailable, manufacturer-wide records — a company failure record with exposure, a company downtime record, or a company-wide claim rate — earn the stated fallback points (VRS-GEN-005 §6); the substitution shall be recorded per Clause 8. Parent-company evidence is admissible on the same basis where the manufacturer is a declared subsidiary.

9 8 Determination procedure

Step 1 — Fix the tier and applicability. Record the tier at which each criterion is assessed (Clause 7) and, per 5.5, the set of any not-applicable criteria; for the reliability category all eight apply.

Step 2 — Award points. For each criterion, the assessor shall award the highest ladder rung the evidence supports (Clause 6) and shall **cap** the award at the ceiling of that criterion's assessed tier (5.3). Where the evidence bar is not met, the award shall be 0 (5.4). Numeric rung thresholds marked **(prior)** are read as written; they are not recomputed here.

Step 3 — Record substitutions. Where a company-wide fallback is used (RL-3, RL-4, RL-8), record the substitution and the fallback points (7.6).

Step 4 — Sum. The category score shall be the **sum of the capped per-criterion points**. With no not-applicable criteria the reachable maximum is 66; otherwise it is 66 minus the ranges of the excluded criteria, and this reachable maximum shall be recorded alongside the score.

Step 5 — Record the tier reached. The category tier is the lowest tier among the criteria that determined the score (a score resting on any Unverified criterion is an Unverified-tier result).

Step 6 — Output and passport binding. Record, for the model's Robot Risk Passport reliability-category record (VRS-GEN-009): the category score, the per-criterion points, the ceiling reached and tier for each criterion, the not-applicable set and reachable maximum, and the company-wide substitutions, so that an underwriter or lender can reproduce and audit the score from the passport alone. The per-criterion points carry the **frequency, business-interruption/availability and correlated-regression** loss drivers; residual value, recoverability and obsolescence are the finance drivers of **economics_residual** (VRS-GEN-107) and **serviceability** (VRS-GEN-103) and shall not be inferred from this score.

10 9 Worked example (fictional model)

“Northline AMR-7”, a warehouse AMR, assessed at the **Verified** tier: a manufacturer evidence pack is examined but no behaviour is witnessed. All eight criteria apply (reachable maximum at Verified = 56).

§	CRITERION	RUNG EARNED	VERIFIED CEILING	POINTS
6.1	RL-1	5 (all subsystems with basis)	5	5
6.2	RL-2	8 (four envelope areas, test reports in pack)	8	8
6.3	RL-3	9 (≤ 1 per 1,000 op-hours; model record + exposure)	12	9
6.4	RL-4	7 ($\geq 95\%$; downtime/state record)	10	7
6.5	RL-5	5 (FMEA + fault-response test evidence)	5	5
6.6	RL-6	9 (clean run in duty-cycle test report) → capped	4	4
6.7	RL-7	4 (one regression, rollback documented)	6	4
6.8	RL-8	6 (≥ 2 yr comprehensive, low claim rate)	6	6
Total			56	48

The reliability-category total is **48 of a reachable 56** at the Verified tier. Arithmetic: $5 + 8 + 9 + 7 + 5 + 4 + 4 + 6 = 48$. Two effects are worth reading. At RL-6 the duty-cycle test report shows a clean session, which would earn the top rung of 9, but the Verified ceiling of 4 caps the award — the model would have to be assessed at the Certified tier, with the endurance run witnessed, to earn the 9. At RL-3 the field failure rate sits in the “ ≤ 1 per 1,000 op-hours” band (rung 9) rather than the “ ≤ 0.5 ” band, and at RL-4 availability is at 95 % rather than 98 %, so neither reaches its top rung; at RL-7 one documented regression with a documented rollback earns rung 4 rather than the clean rung 6. No company-wide fallback was used (model-specific failure, downtime and warranty evidence were present), and the not-applicable set is empty. The score, per-criterion points, and tier and ceiling reached for each criterion are recorded to the passport (Step 6).

11 10 Assessor checklist

- ☐ Model granularity confirmed and reliability-software baseline scope declared (4.3).
- ☐ Assessment tier fixed per criterion; not-applicable set recorded (Step 1); reachable maximum recorded.
- ☐ **RL-1** rated life per major subsystem scored with its basis.
- ☐ **RL-2** environmental and durability tests scored; operating envelope declared (evidence bar); IP/public-mark rung applied only if publicly verifiable (7.4).

- [] **RL-3** field failure rate scored; exposure hours declared (evidence bar); regime rule applied per 7.5.
- [] **RL-4** availability scored from SEMI E10 states; scheduled time and downtime declared (evidence bar).
- [] **RL-5** fault tolerance and graceful degradation scored; fault-response behaviour documented (evidence bar); FMEA present.
- [] **RL-6** in-session endurance scored; representative task declared (evidence bar); Unverified ceiling is 0.
- [] **RL-7** software release quality scored; release history for window present (evidence bar).
- [] **RL-8** warranty terms and claim rate scored; terms published (evidence bar).
- [] Company-wide substitutions recorded (Step 3; 7.6).
- [] Each award capped at the assessed-tier ceiling (5.3); no-evidence criteria scored 0, never null (5.4).
- [] Score, per-criterion points, tier and ceiling reached recorded to the passport (Step 6).

12 Bibliography

- IEC 60050-192:2015, *International Electrotechnical Vocabulary — Part 192: Dependability*.
- IEC 61703:2016, *Mathematical expressions for reliability, availability, maintainability and maintenance support terms*.
- ISO 14224:2016 (corrected 2016-10-01), *Petroleum, petrochemical and natural gas industries — Collection and exchange of reliability and maintenance data for equipment*.
- IEC 60529:1989+A2:2013, *Degrees of protection provided by enclosures (IP Code)*.
- IEC 62402:2019, *Obsolescence management* (software support-horizon / end-of-support declarations).
- SEMI E10 (2021 RAM revision), *Specification for Definition and Measurement of Equipment Reliability, Availability, and Maintainability (RAM) — six-state equipment model*.
- High-availability “nines” / service-level availability convention (99.0 % $\approx$ 3.65 days/yr; 99.9 % $\approx$ 8.77 h/yr downtime) — used as a named anchor for the RL-4 rungs, not quoted.

13 Change history

DATE	VERSION	STATUS	CHANGE	AUTHORITY
2026-09-15	1.0	Draft	VRS 2026 rewrite: replaced the rate-based scoring (input schema, failure/equipment-state coding, exposure and IEC 61703 availability model, window/recency/credibility rules, availability and failure-rate formulae, regression-penalty and override cap, qualitative level descriptors) with eight point-scored criteria RL-1...RL-8 (66 points; Clause 6) with per-tier ceilings; added Clause 5 Scoring basis, Clause 7 Evidence by tier (certificate rung, field-record regime rule, company-wide fallback, parent-company evidence), Clause 8 point-sum determination and passport binding, regenerated worked example (Verified tier, 48/56) and assessor checklist; retired confidence bands, provisional flags, renormalisation, class/dataset medians and credibility blend, cold-start/GEN-201 dependency and qualitative IR levels; dropped GEN-201/GEN-202 references and the calibration/statistics bibliography that served the rate model; see GEN-005 §6	VRS 2026 rewrite Stage 4
2026-09-11	0.3	Draft	ACTUARIAL review: added §9.2 passport binding and §9.3 loss-driver/decision-mapping table; added VRS-GEN-009 to §2 and CAS Ratemaking Principles (2021) + ASOP No. 25 (rev. 2014) to Bibliography	ACTUARIAL run #28
2026-09-11	0.2	Draft	Expanded to publishable depth (DEEPEN): input schema, failure/equipment-state coding, exposure + IEC 61703 availability model, window/recency/credibility/regression-remediation rules, 8-step procedure with anchor-derived scale priors, worked example, assessor checklist; dated anchors; replaced the bare calibration marker with stage labels	CEO direction 2026-09-11
2026-09-05	0.1	Draft	Split from VRS-GEN-005 §5.2; level descriptors added per the category-standard wireframe	CEO goal 2026-09-05

About Veyrum Research Institute

Veyrum Research Institute is the research arm of Veyrum — the independent intelligence and trust layer for the robotics industry (research, ratings and registry, advisory, data, and specialist insurance & finance introductions).

Disclaimer

The opinions in this document are independent and current as of the date of publication, based on information believed reliable; they are not guarantees or warranties of safety, fitness, or compliance, and are not engineering, legal, or financial advice. This document may be revised. © 2026 Veyrum.

Contact: hello@veyrum.com · VRS-GEN-102 · Version 1.0