



VEYRUM ROBOTICS STANDARD

**VRS-GEN-103**

IR Category: Serviceability and Support

Scoring basis, serviceability criteria SV-1...SV-8, evidence by tier and determination procedure for the serviceability-and-support intrinsic category.

VEYRUM RESEARCH INSTITUTE

DOCUMENT

VRS-GEN-103

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 SV-1 Service documentation quality (0–6)

7.2

6.2 SV-2 Field-replaceable design (0–10)

7.3

6.3 SV-3 Diagnostics and fault visibility (0–6)

7.4

6.4 SV-4 Restoration time (MTTR, critical repairs) (0–8)

7.5

6.5 SV-5 Parts availability and lead time (0–8)

7.6

6.6 SV-6 Support SLA and service-network coverage (0–6)

7.7

6.7 SV-7 Support responsiveness achieved (0–6)

7.8

6.8 SV-8 Support horizon and end-of-life policy (0–5)

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 **serviceability and support** (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 serviceability criteria (SV-1...SV-8, 55 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 restoration-time or parts-lead formula, no statistical blend toward a class median, no qualitative bands. Requirements use “shall”.

1 Introduction

Serviceability is the loss driver an insurer reads as claim *severity* — how long a failure keeps a robot out of service, and therefore how large the business-interruption exposure grows — and a lender reads as *recoverability* and *obsolescence*: can the asset be restored to earning condition, and for how long will the manufacturer keep it supportable. A robot that fails is a frequency event; a robot that cannot be repaired quickly, or at all, is a severity and residual-value event. This Standard turns “how serviceable is this model?” into a repeatable number by asking eight questions — is its service documentation complete, is it designed to be field-repaired, can its faults be seen, how fast is it restored, are its parts available, is its support commitment and network real, does it meet that commitment in practice, and how long will it stay supportable — and awarding integer points for each. The points sum to a category total of 0–55. 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 serviceability that is designed in *and* shown, not merely asserted, and it reuses the records operators already keep (ISO 14224:2016 maintenance events, IEC 62402:2019 obsolescence declarations) so that adoption is a mapping, not new testing.

2 1 Scope

1.1 This Standard defines, for the **serviceability-and-support** intrinsic category: what the category measures (Clause 4); the scoring basis (Clause 5); the eight criteria SV-1...SV-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), and the evidence-grade definitions (VRS-GEN-012). It measures how quickly and dependably the model returns to service when it fails and how long it remains supportable; the *frequency* of failure and the resulting availability are measured by reliability (VRS-GEN-102). Restoration time is defined here and consumed by VRS-GEN-102 as an availability input only; it is not re-scored there. This Standard 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 (maintainability, MTTR, mean down time, maintenance-support-performance definitions).
- **IEC 61703:2016**, Mathematical expressions for reliability, availability, maintainability and maintenance support terms — mean repair time, mean down time, logistic delay.
- **IEC 60706-2:2006**, Maintainability of equipment — Part 2: Maintainability requirements and studies during the design and development phase — active corrective maintenance time basis.
- **IEC 60300-3-14:2024**, Dependability management — Supportability and support — supportability attribute and maintenance-support framework.
- **ISO 14224:2016** (corrected 2016-10-01), Collection and exchange of reliability and maintenance data for equipment — repair-event taxonomy and the reused maintenance-record evidence lane.
- **IEC 62402:2019**, Obsolescence management — end-of-support / support-horizon declaration lane.

4 Terms and definitions

Terms per VRS-GEN-001, IEC 60050-192:2015 and IEC 60300-3-14:2024 (maintainability, supportability, maintenance support). Locally:

- **repair event** — a completed corrective maintenance action restoring a failed unit to its function (ISO 14224:2016), the counting unit for the restoration-time criterion (SV-4).
- **restoration time** **MTTR** — the mean restoration time per critical repair event, from failure detection to return-to-service, in hours, comprising active repair time (IEC 60706-2:2006) and logistic delay (IEC 61703:2016 mean down time).
- **parts lead time** — the logistic delay from a confirmed parts requirement to parts on-site, in days (IEC 61703:2016 logistic delay).
- **field-replaceable module** — a module a manufacturer declares replaceable in the field by an operator or technician with common tools (the FRU list of SV-2).
- **support horizon** — the time from the assessment date to the manufacturer's declared end-of-support date for the model (IEC 62402:2019 declaration; SV-8).
- **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 in lieu of a witnessed demonstration.

5 4 What this category measures

4.1 The category measures the model's **restoration behaviour and supportability under real duty** — service documentation, field-repairability, fault visibility, restoration speed, parts logistics, support commitment and coverage, support responsiveness achieved, and support horizon — bounded by the evidence that substantiates each claim.

4.2 It measures both **design and evidence and record**: whether the model is documented and built to be serviced and diagnosed (SV-1, SV-2, SV-3), what restoration time and parts availability actually are (SV-4, SV-5), what support is committed and delivered (SV-6, SV-7), and how long the model will remain supportable (SV-8). A declared serviceable design that is only asserted is measured at the Unverified ceiling; the observed repair record and the binding support declaration are measured at higher tiers. 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). Repair events and support arrangements shall not be pooled across models that differ in serviceability-relevant hardware or support arrangements unless the manufacturer declares them one model for serviceability 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 SV-1...SV-8 (Clause 6). The maximum is 55 points. There is no restoration-time or parts-lead 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 serviceability 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.

76 Criteria

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

§	CRITERION	POINTS	UNVERIFIED CEILING	VERIFIED CEILING	CERTIFIED CEILING
6.1	SV-1 Service documentation quality	0–6	6	6	6
6.2	SV-2 Field-replaceable design	0–10	6	6	10
6.3	SV-3 Diagnostics and fault visibility	0–6	4	4	6
6.4	SV-4 Restoration time (MTTR, critical repairs)	0–8	5	8	8
6.5	SV-5 Parts availability and lead time	0–8	6	8	8
6.6	SV-6 Support SLA and service-network coverage	0–6	5	6	6
6.7	SV-7 Support responsiveness achieved	0–6	4	6	6
6.8	SV-8 Support horizon and end-of-life policy	0–5	5	5	5
Total		55	41	49	55

7.1 6.1 SV-1 Service documentation quality (0–6)

ROW	
What is measured	Service manual, preventive-maintenance schedule, diagnostics guide, torque/consumable specs.
Ladder	6 = all, current for rated revision · 4 = manual + schedule · 2 = manual only · 0 = none
Unverified — evidence & ceiling	Public manual. Ceiling 6.
Verified — evidence & ceiling	Full service documentation. Ceiling 6.
Certified — evidence & ceiling	From Verified evidence. Ceiling 6.
Company-wide fallback	n/a
Evidence bar	None.
Applicability	All classes

7.2 6.2 SV-2 Field-replaceable design (0–10)

ROW	
What is measured	Declared field-replaceable modules can be swapped by an operator/technician with common tools within declared times.
Ladder	10 = witnessed within declared times · 6 = documented with times · 3 = documented, no times · 0 = factory-only repair
Unverified — evidence & ceiling	Manual. Ceiling 6.
Verified — evidence & ceiling	Module list with procedures and times. Ceiling 6.
Certified — evidence & ceiling	Witnessed battery + one module swap, timed. Ceiling 10.
Company-wide fallback	n/a
Evidence bar	FRU list declared.
Applicability	All classes

7.3 6.3 SV-3 Diagnostics and fault visibility (0–6)

ROW	
What is measured	Operators can reach fault codes, logs and health indicators without manufacturer intervention.
Ladder	6 = shown live · 4 = documented · 2 = partial · 0 = manufacturer-only
Unverified — evidence & ceiling	Manual. Ceiling 4.
Verified — evidence & ceiling	Diagnostics documentation. Ceiling 4.
Certified — evidence & ceiling	Shown on the robot’s own console. Ceiling 6.
Company-wide fallback	n/a
Evidence bar	None.
Applicability	All classes

7.4 6.4 SV-4 Restoration time (MTTR, critical repairs) (0–8)

ROW	
What is measured	Mean restoration time for critical repairs over window.
Ladder	8 = ≤ 4 h (prior) · 6 = ≤ 8 h (prior) · 5 = publicly evidenced only (no pack, no demonstration) · 4 = ≤ 24 h (prior) · 2 = ≤ 72 h (prior) · 0 = longer/not established. (Priors.)
Unverified — evidence & ceiling	Operator reports. Ceiling 5.
Verified — evidence & ceiling	Repair record. Ceiling 8.
Certified — evidence & ceiling	From Verified evidence. Ceiling 8.
Company-wide fallback	Company-wide → 6
Evidence bar	Repair record (model or company).
Applicability	All classes

7.5 6.5 SV-5 Parts availability and lead time (0–8)

ROW	
What is measured	Spare-parts catalogue, committed lead times, achieved lead times, independent parts channel.
Ladder	8 = catalogue + ≤ 5-day achieved + independent channel · 6 = catalogue + committed ≤ 10 days (prior) · 3 = catalogue only · 2 = parts covered under a paid support plan, no public catalogue (F5) · 0 = none
Unverified — evidence & ceiling	Public catalogue / parts store. Ceiling 6.
Verified — evidence & ceiling	Catalogue + lead-time record. Ceiling 8.
Certified — evidence & ceiling	From Verified evidence. Ceiling 8.
Company-wide fallback	Achieved lead time company-wide → 6
Evidence bar	Catalogue or support-plan parts coverage exists.
Applicability	All classes

7.6 6.6 SV-6 Support SLA and service-network coverage (0–6)

ROW	
What is measured	Response-time commitment and service coverage, read together (F15 split retained as two facts, one ladder).
Ladder	6 = numeric response commitment ≤ 24 h (prior) published and service presence in every region where the model is sold · 5 = ≤ 24 h (prior) commitment with partial regional coverage, or a longer numeric commitment with full coverage (publicly evidenced only reaches this rung) · 4 = numeric commitment (any) with partial coverage · 3 = numeric commitment with single-region / remote-only coverage, or channels only with full coverage · 2 = channels published without a numeric commitment, partial coverage · 1 = channels published, single region / remote only · 0 = no published support commitment
Unverified — evidence & ceiling	Published terms/plans. Ceiling 5.
Verified — evidence & ceiling	SLA + coverage map. Ceiling 6.
Certified — evidence & ceiling	From Verified evidence. Ceiling 6.
Company-wide fallback	n/a
Evidence bar	None.
Applicability	All classes

7.7 6.7 SV-7 Support responsiveness achieved (0–6)

ROW	
What is measured	Fraction of SLA targets met in window.
Ladder	6 = ≥ 95 % (prior) · 4 = ≥ 85 % (prior) · 2 = ≥ 70 % (prior) · 0 = lower/not established
Unverified — evidence & ceiling	—. Ceiling 4.
Verified — evidence & ceiling	Service-log adherence figure. Ceiling 6.
Certified — evidence & ceiling	From Verified evidence. Ceiling 6.
Company-wide fallback	Company-wide → 4
Evidence bar	SLA exists.
Applicability	All classes

7.8 6.8 SV-8 Support horizon and end-of-life policy (0–5)

ROW	
What is measured	Declared parts/software support horizon relative to the assessment date, and EOL notice policy.
Ladder	5 = ≥ 7 yr remaining + ≥ 12 -month notice · 3 = ≥ 5 yr remaining, or a stated horizon has passed but updates are demonstrably continuing (F20) · 1 = stated, shorter, or horizon passed with no renewal statement · 0 = none
Unverified — evidence & ceiling	Public declaration. Ceiling 5.
Verified — evidence & ceiling	Support/obsolescence declaration. Ceiling 5.
Certified — evidence & ceiling	From Verified evidence. Ceiling 5.
Company-wide fallback	n/a
Evidence bar	None.
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: the public service manual, preventive-maintenance schedule and diagnostics guide (SV-1, SV-3), the published field-replaceable-unit list (SV-2), operator reports of restoration times (SV-4), a public parts catalogue or parts store (SV-5), published support terms and service plans (SV-6), and the public end-of-support or obsolescence declaration (SV-8). Each criterion's Unverified ceiling is the most an honest reading of such evidence can earn; the category tier is **Unverified**. SV-7 has no public path and its ceiling rests on the Verified service-log figure.

7.2 Verified. The assessor additionally examines a manufacturer evidence pack: the full service documentation set (SV-1), the module list with its swap procedures and declared times (SV-2), the diagnostics documentation (SV-3), the repair record with restoration times (SV-4), the parts catalogue with a lead-time record (SV-5), the service-level agreement with a coverage map (SV-6), the service-log adherence figure (SV-7), and the support/obsolescence declaration (SV-8). 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: a battery and one further module swap performed and timed against the declared times (SV-2), and diagnostics shown live on the robot's own console (SV-3). The documentary and record criteria (SV-1, SV-4, SV-5, SV-6, SV-7, SV-8) carry their Certified points from the Verified evidence. The category tier is **Certified**.

7.4 Certificate rung. Where a criterion allows a publicly verifiable third-party certificate or mark in lieu of a witnessed demonstration, an artefact identified by registrar entry or certificate number earns the stated rung; an unverifiable mark shall not.

7.5 Company-wide fallback. Where a criterion offers a company-wide fallback (SV-4 → 6, SV-5 → 6, SV-7 → 4) and model-specific evidence is unavailable, manufacturer-wide records — a company repair record, a company achieved-lead-time record, or a company-wide service-log adherence figure — 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.

7.6 Support-horizon watch. A support horizon (SV-8) that will lapse within the assessment edition's validity shall be flagged as an obsolescence watch item (VRS-GEN-004 review); the flag records a finance recoverability and residual-value signal and does not itself change the awarded points.

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 serviceability 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 (SV-4, SV-5, SV-7), record the substitution and the fallback points (7.5).

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 55; otherwise it is 55 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 serviceability-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, the company-wide substitutions, and the support-horizon watch flag, so that an underwriter or lender can reproduce and audit the score from the passport alone. The per-criterion points carry the **business-interruption severity, recoverability, residual-value and obsolescence** loss drivers: restoration time and parts availability (SV-4, SV-5) size the outage-duration severity an insurer prices as a business-interruption sub-limit or waiting-period deductible and a lender reads into a debt-service-coverage covenant; support commitment and responsiveness (SV-6, SV-7) size recoverability; and the support horizon (SV-8) is an obsolescence and residual-value signal a lender reads into the residual-value curve, LTV and end-of-term covenant. Failure *frequency* and availability are the drivers of reliability (VRS-GEN-102) 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 = 49). The pack shows full service documentation, a field-replaceable-module list with swap procedures and times, diagnostics documentation, a repair record with a mean restoration time near 15 h, a parts catalogue with a committed lead time under 10 days, a service-level agreement with a ≤ 24 h numeric response commitment but service presence in only some of the regions where the model is sold, a service-log adherence figure of 0.92, and an end-of-support declaration roughly 33 months out.

§	CRITERION	RUNG EARNED	VERIFIED CEILING	POINTS
6.1	SV-1	6 (full service documentation in pack)	6	6
6.2	SV-2	6 (module list with procedures and times, documented)	6	6
6.3	SV-3	4 (diagnostics documentation)	4	4
6.4	SV-4	4 (mean restoration ≤ 24 h (prior); repair record)	8	4
6.5	SV-5	6 (catalogue + committed ≤ 10 days (prior))	8	6
6.6	SV-6	5 (≤ 24 h commitment with partial regional coverage)	6	5
6.7	SV-7	4 (≥ 85 % (prior) service-log adherence)	6	4
6.8	SV-8	1 (stated horizon, shorter than 5 yr)	5	1
Total			49	36

The serviceability-category total is **36 of a reachable 49** at the Verified tier. Arithmetic: $6 + 6 + 4 + 4 + 6 + 5 + 4 + 1 = 36$. Three effects are worth reading. Documentation, field-repairability and parts (SV-1, SV-2, SV-5) are strong, but two of the top rungs are unreachable at this tier by construction: SV-2 rung 10 needs a witnessed, timed module swap and SV-3 rung 6 needs diagnostics shown live on the console — both are Certified-tier steps, so the model would have to be assessed at the Certified tier to earn them. The binding weakness is SV-8: an end-of-support date only ~33 months out sits below the 5-year rung and earns 1 point, and it is flagged as an obsolescence watch item (7.6) — the finance recoverability and residual-value signal that the strong repair and parts picture alone would understate. No company-wide fallback was used (model-specific repair, parts and service-log 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 repair events / support arrangements not improperly pooled (4.3).
- [] Assessment tier fixed per criterion; not-applicable set recorded (Step 1); reachable maximum recorded.
- [] **SV-1** service documentation scored against the ladder.
- [] **SV-2** field-replaceable design scored; FRU list declared (evidence bar); witnessed swap only at Certified.
- [] **SV-3** diagnostics and fault visibility scored; live-console rung only at Certified.
- [] **SV-4** restoration time scored; repair record present (evidence bar); (prior) thresholds read as written.
- [] **SV-5** parts availability and lead time scored; catalogue or support-plan coverage exists (evidence bar).
- [] **SV-6** support SLA and coverage scored as one ladder from the two facts (commitment, coverage).
- [] **SV-7** support responsiveness scored; SLA exists (evidence bar); no public path (Unverified ceiling 4).
- [] **SV-8** support horizon and EOL policy scored; horizon-lapse flagged as an obsolescence watch item (7.6).
- [] Company-wide substitutions recorded (Step 3; 7.5).
- [] 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, and horizon watch flag recorded to the passport (Step 6).

12 Bibliography

- IEC 60050-192:2015, *International Electrotechnical Vocabulary — Part 192: Dependability* (maintainability, MTTR, mean down time, maintenance-support-performance terms).
- IEC 61703:2016, *Mathematical expressions for reliability, availability, maintainability and maintenance support terms* (mean repair time, mean down time, logistic delay).
- IEC 60706-2:2006, *Maintainability of equipment — Part 2: Maintainability requirements and studies during the design and development phase* (active corrective maintenance time).
- IEC 60300-3-14:2024, *Dependability management — Supportability and support* (supportability attribute and maintenance-support framework).
- ISO 14224:2016 (corrected 2016-10-01), *Petroleum, petrochemical and natural gas industries — Collection and exchange of reliability and maintenance data for equipment* (repair-event taxonomy).
- IEC 62402:2019, *Obsolescence management* (end-of-support / support-horizon declaration).

13 Change history

DATE	VERSION	STATUS	CHANGE	AUTHORITY
2026-09-15	1.0	Draft	VRS 2026 rewrite: replaced the rate-based scoring (input schema, restoration-time/support coding, exposure and class-median denominators, window/recency/small-sample rules, restoration-time and parts-lead formulae, end-of-support penalty and override cap, qualitative level descriptors) with eight point-scored criteria SV-1...SV-8 (55 points; Clause 6) with per-tier ceilings; added Clause 5 Scoring basis, Clause 7 Evidence by tier (certificate rung, company-wide fallback, parent-company evidence, support-horizon watch), Clause 8 point-sum determination and passport binding, regenerated worked example (Verified tier, 36/49) 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	0.5	Draft	Consolidation (REFACTOR, pair GEN-102/103): §9.1, the §9.2 passport-binding framing + GEN-009 caveat, and the §9.3 signals-not-prices/CAS preamble replaced by cross-references to their single home VRS-GEN-102 §9.1–§9.3; serviceability field map, loss-driver table and finance-driver line kept; no requirement removed (36 preserved); GEN-102 untouched	CEO direction 2026-09-11
2026-09-11	0.4	Draft	Insurance/finance fitness review (ACTUARIAL). Added §9.2 Passport binding and §9.3 Loss-driver and decision mapping; added VRS-GEN-009 normative reference and a §12 checklist line; anchors CAS Ratemaking Statement of Principles (reinstated 2021) and ASOP No. 25 (rev. 2014) verified. No requirement deleted; shall 32→34; net +330 words	Curator Run #36; charter ACTUARIAL
2026-09-11	0.3	Draft	Consolidation (REFACTOR, pair GEN-102/103): §5.1 preamble, §8.2 recency-weight bands, §8.3 credibility method and §9 Step 7 replaced by cross-references to their single home VRS-GEN-102; foreword clause list de-duplicated against §1. No requirement deleted, count unchanged (32); net –13 words	CEO direction 2026-09-11

DATE	VERSION	STATUS	CHANGE	AUTHORITY
2026-09-11	0.2	Draft	Expanded to publishable depth (DEEPEN): input schema, restoration-time decomposition + support rubric, class-median exposure and support-horizon denominators, window/recency/credibility rules, 8-step procedure with anchor-derived priors + end-of-support penalty and override cap, worked example, assessor checklist; dated anchors; CALIBRATION-PENDING replaced by stage labels	CEO direction 2026-09-11
2026-09-05	0.1	Draft	Split from VRS-GEN-005 §5.3; 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-103 · Version 1.0