



TRUST. TRANSPARENCY. PERFORMANCE.

VEYRUM ROBOTICS STANDARD



VRS-GEN-005

Intrinsic Rating (IR) — Methodology

The central IR scoring standard — the seven categories, the point sum, tier ceilings, rating tiers and editions (VRS 2026).

VEYRUM RESEARCH INSTITUTE

DOCUMENT

VRS-GEN-005

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 Rating architecture	
6	5 IR categories	
7	6 Aggregation and tier ceilings	
8	7 Rating tiers	
9	8 Editions	
10	9 Versioning	
11	10 Worked example (fictional model)	
12	Bibliography	
13	Change history	

Foreword. This Standard is the master methodology for the Veyrum **Intrinsic Rating (IR)**. From the **VRS 2026** edition the IR is an unweighted sum of criterion points: every criterion in the seven IR Category Standards (VRS-GEN-101...107) awards an integer number of points against a ladder of verifiable states, each criterion states the maximum it can award at each of the three rating tiers, and the IR is the total. This edition replaces the weighted mean, the renormalisation of missing categories, the confidence model and the qualitative level names of earlier drafts; the reasons are recorded in the change history. Requirements use “shall”; recommendations use “should”.

The design follows the point-sum conventions of building-performance and product-rating schemes (LEED, BREEAM, Euro NCAP) rather than the weighted-index conventions of credit ratings, because an additive scale with frozen point meanings is the only form under which two ratings from different years can be compared without a conversion table.

1 Introduction

A rating that changes meaning when its weights are re-tuned cannot anchor an insurance price or a finance covenant. The IR is therefore built so that a point means the same thing in every edition: each criterion’s point range is sized once to its importance, the edition freezes those ranges, and later editions may only add. What the manufacturer can prove — publicly, by documents, or by demonstration — determines how many of those points are reachable; that is the whole of the tier mechanism, and it lives inside the criteria rather than as a rule applied to the total.

2 1 Scope

This Standard defines the **IR — Intrinsic Rating**: one integer per robot model, higher meaning better intrinsic quality, safety and reliability, independent of where or how the model is used. It specifies the seven IR categories and their point totals, the aggregation rule, the tiered ceiling mechanism, the three rating tiers, the display form, and how editions change the scale. Under the VRS 2026 edition the IR ranges from 0 to **400**; every IR shall be edition-stamped, e.g. **IR 292 (VRS 2026)**.

This Standard excludes: the per-criterion ladders, ceilings and evidence (the IR Category Standards VRS-GEN-101...107); sector suitability (VRS-GEN-006 and the sector protocols); the passport and registry record (VRS-GEN-009); evidence admissibility (VRS-GEN-012); and the information-pack and demonstration-protocol documents that derive from the criteria.

3 2 Normative references

- **VRS-GEN-001**, Vocabulary and Terminology (undated).
- **VRS-GEN-002**, Designation, Drafting and Style Rules — criterion layout §6.5 and criterion drafting rules §7.5 (undated).
- **VRS-GEN-004**, Change, Edition and Review Policy (undated).
- **VRS-GEN-006**, Application Rating (AR) Scheme (undated) — tier rule between IR and AR.

- **VRS-GEN-012**, Evidence and Data Requirements (undated) — admissibility by tier.
- **VRS-GEN-101 ... VRS-GEN-107**, IR Category Standards (undated) — the criteria.
- Anchors (dated at category level): ISO 14224:2016 (reliability taxonomy), IEC 62443-4-2:2019 (cyber), ISO 12100:2010 (risk assessment), ISO 10218-1:2025 / -2:2025 (industrial robot safety).

4 3 Terms and definitions

Terms per VRS-GEN-001 and VRS-GEN-002 §3.8–3.11 (criterion, rung, ceiling, evidence bar). Locally:

- **rating tier** — one of **Unverified, Verified, Certified** (Clause 7): the level of manufacturer involvement and evidence under which a rating was determined.
- **edition** — an annual release of the VRS library (e.g. *VRS 2026*) that fixes the criteria, their point ranges and the IR maximum for ratings issued under it (Clause 8).
- **edition stamp** — the edition named on every displayed rating value.
- **company-wide fallback** — the substitution, for a field-record criterion, of the manufacturer's data across its product line for data specific to the rated model, at the criterion's stated reduced ceiling.
- **cold-start parity** — the property that a model with no fleet history can reach the same IR as a mature one by evidence and demonstration (6.6).

5 4 Rating architecture

4.1 The IR shall be the **sum of the points awarded on every criterion** of the seven IR Category Standards (Clause 6). No weights, multipliers, normalisation factors or confidence adjustments shall be applied to the sum or to any criterion.

4.2 The IR is an integer on an **absolute scale**. The VRS 2026 edition defines a maximum of **400** points (Clause 5). An IR has no letter and no qualitative label; it shall be displayed with its edition stamp and rating tier, in the form **IR <n> (VRS 2026) – <Tier>** (VRS-GEN-002 §7.4).

4.3 Every rating shall record the tier under which it was determined (Clause 7), the edition, the hardware revision and software version range rated, and — per criterion — the points awarded, the rung reached, the ceiling that applied and whether the company-wide fallback was used (VRS-GEN-009).

4.4 The IR describes the model as built and supported; it is an opinion on intrinsic quality above the applicable legal floor, not a conformity attestation and not a statement about any deployment (VRS-GEN-006 §10, VRS-GEN-010).

6 5 IR categories

5.1 The seven categories, their Category Standards and their point totals under VRS 2026 are:

CATEGORY	CATEGORY STANDARD	CRITERIA	POINTS	UNVERIFIED CEILING	VERIFIED CEILING	CERTIFIED CEILING
Safety	VRS-GEN-101	10	100	66	79	100
Reliability and durability	VRS-GEN-102	8	66	37	56	66
Serviceability and support	VRS-GEN-103	8	55	41	49	55
Specification integrity	VRS-GEN-104	5	47	32	37	47
Cyber posture	VRS-GEN-105	8	64	44	51	64
Manufacturer support and quality system	VRS-GEN-106	7	39	33	39	39
Economics and residual value	VRS-GEN-107	5	29	27	29	29
Total (IR maximum)		51	400	280	340	400

5.2 The totals in 5.1 are the sums of the criterion ranges and ceilings in the Category Standards and shall be reproducible from them by addition. A Category Standard whose Clause 6 totals do not reproduce this table is malformed (VRS-GEN-002 §6.5.3).

5.3 The share of the IR carried by each category (Safety 25.0 %, Reliability 16.5 %, Serviceability 13.75 %, Specification integrity 11.75 %, Cyber 16.0 %, Manufacturer support 9.75 %, Economics 7.25 %) is a **consequence** of the criterion ranges, not a parameter. Changing a category's influence is done by adding or widening criteria in an edition (Clause 8), never by weighting.

5.4 References to the Category Standards are undated (VRS-GEN-002 §7.3.2). Adding a category is an amendment to this clause plus one new Category Standard.

7 6 Aggregation and tier ceilings

6.1 Point sum. $IR = \sum \text{points}(\text{criterion})$ over all 51 criteria. A criterion on which no admissible evidence meets the evidence bar scores **0**; it is not omitted, and no other criterion's points change because of it.

6.2 Tiered ceilings are the only cap. Each criterion states, in its own table, the maximum points it can award at Unverified, at Verified and at Certified (VRS-GEN-002 §6.5.2). These per-criterion ceilings are

the sole mechanism by which the tier limits the IR; there is no global cap per tier and no blanket percentage. The tier ceilings of 5.1 are therefore sums, not rules.

6.3 Ceilings equal rungs. Every ceiling is a rung value of its criterion (VRS-GEN-002 §7.5.1). The points awarded at a tier are the points of the highest rung the admissible evidence establishes, never exceeding that tier's ceiling.

6.4 Not every criterion needs three ceilings. A criterion fully establishable from public material has the same ceiling at every tier; a criterion establishable only by demonstration has an Unverified ceiling of 0. Both are stated in the criterion, not inferred.

6.5 Field-record criteria and the company-wide fallback. A criterion that scores a record over time (incidents, failures, availability, repairs, warranty conduct, recalls) accepts model-specific data at its stated ceiling and the manufacturer's company-wide data at its own, explicitly stated, lower ceiling. The reduced ceiling is a rung value and is stated as a number, never as a percentage.

6.6 Cold-start parity. The criteria shall allow a model with no fleet history to reach the same IR as a mature model: no criterion requires fleet age as such, and every field-record criterion offers the company-wide fallback of 6.5. A first-of-model assessment is an ordinary assessment.

6.7 Recognised third-party evidence. A certificate or test report from an accredited or notified body that covers a criterion establishes it at the criterion's **certificate rung** (VRS-GEN-002 §7.5.3). At Unverified, a **publicly verifiable** certificate may reach the Unverified ceiling — no reduction is applied merely because the evidence is a certificate. At Verified the report establishes the criterion at the Verified ceiling. At Certified a certificate counts only where the demonstration protocol names it as an accepted equivalent to a specific step. A statement of certification without a verifiable certificate or register entry is documentation, not a certificate (VRS-GEN-012 §6).

6.8 Null results. For a field-record criterion whose Unverified evidence is a search of public registers, a documented null result (registers searched, coverage stated, nothing found) scores the criterion's **null-result rung**; it is neither "no evidence" (0) nor "record clean" (top).

6.9 Regime-aware ceilings. Where a mandatory public reporting regime makes a field record genuinely public for the model's sector, the sector protocol may set the Unverified ceiling of the corresponding field-record criteria at their full value (VRS-GEN-002 §7.5.5). The 280 figure of 5.1 is the default-regime total.

6.10 Guidance is not disclosure. Manufacturer publications about what users *should* do, white papers and marketing establish nothing about the rated product; only statements about the product, its revision and its record do.

6.11 Parent-company evidence. For company-level criteria (quality system, vulnerability handling, going concern, warranty and support conduct), evidence held by a parent or group establishes the criterion only where the parent's system demonstrably governs the rated product line (VRS-GEN-002 §7.5.7); otherwise it is context.

6.12 Not applicable by class. IR criteria are written to apply to every robot form (claims are scored as a set; functions are named by role, not by sensor). Where a sector criterion declares itself inapplicable to an application class, it is removed from both numerator and denominator of the sector pool (VRS-GEN-006 §6); no IR criterion is ever scored 0 for inapplicability.

6.13 Priors. A numeric threshold inside a rung that has not yet been calibrated against assessment data is marked **(prior)** in the criterion. Priors change by edition (Clause 8), never by assessor judgement.

8 7 Rating tiers

7.1 Every rating shall be determined under exactly one tier:

	UNVERIFIED	VERIFIED	CERTIFIED
Manufacturer involvement	None	Information pack + officer attestation	Verified pack (same scope) + witnessed demonstration
Initiated by	Veyrum	Manufacturer	Manufacturer
Eligibility	Model publicly available ≥ 6 months	Manufacturer selling robots ≥ 6 months and $\geq 1,000$ units deployed (any models, any sectors)	As Verified
Earliest timing	6 months after launch	Launch day or pre-launch	Launch day or pre-launch
Admissible evidence	Public sources only (VRS-GEN-012 §11)	Public sources + the information pack	Public sources + pack + the Demonstration Session record
IR reachable (VRS 2026)	≤ 280	≤ 340	≤ 400
Best attainable AR letter	A	AA	AAA (VRS-GEN-006 §8)
Pack completeness	n/a	All-or-nothing: an incomplete pack yields no Verified rating	As Verified

7.2 Tiers are cumulative: Certified evidence contains Verified evidence contains public evidence. The Certified pack shall not exceed the Verified pack in scope.

7.3 An Application Rating's tier shall not exceed the IR's tier (VRS-GEN-006 §7). Sectors are otherwise independent.

7.4 Verified and Certified ratings are bound to a named hardware revision and software version range and are assessed on a production-representative unit. Publication timing for pre-launch ratings is agreed in writing (publish on issue; hold until launch; hold until a date) and recorded (VRS-GEN-009).

7.5 An Unverified rating shall follow the published Unverified Rating Procedure (VRS-GEN-204): collection period, manufacturer notice and comment window before publication.

9 8 Editions

8.1 The VRS library is released in **annual editions** (VRS-GEN-004). Each edition publishes the full criterion set, the IR maximum, each sector’s pool, and a change log of additions.

8.2 Additive only. An edition may add criteria or widen a criterion’s range at the top. The meaning of every existing point is frozen: a rung that awarded 6 points in VRS 2026 awards 6 points, for the same state, in every later edition. Retired criteria cease to be awarded in new editions; ratings already issued keep them.

8.3 Because the scale is absolute and additive, IR values from different editions **are** comparable raw; the edition stamp tells the reader what maximum applied. A rating cites the edition it was issued under and is not re-stated under a later one unless re-assessed.

8.4 Priors (6.13) and the thresholds inside rungs are changed only by edition, with the derivation published in the edition change log. Assessors shall not adjust them.

8.5 Upcoming additions are pre-announced in a published roadmap so that manufacturers can prepare evidence before the edition that introduces them.

10 9 Versioning

Editions and change instruments follow VRS-GEN-004 and VRS-GEN-002 Clause 8. A rating retains the edition and the document versions that produced it (reproducibility); the registry record carries them (VRS-GEN-009). A change to a criterion’s points, rungs or ceilings is an edition act, never an editorial correction.

11 10 Worked example (fictional model)

“*Meridian LMR-3*”, a fictional logistics mobile robot rated at **Verified** under VRS 2026. The example is illustrative only; it does not imply that any real robot holds a Veyrum rating. Category subtotals are the sums of the criterion points awarded under VRS-GEN-101...107; the per-criterion detail lives in those standards’ own worked examples.

CATEGORY	POINTS AWARDED	VERIFIED CEILING	NOTES
Safety (VRS-GEN-101)	61	79	Risk assessment, safety-function specification and validation reports supplied; field record company-wide (6.5), scored at the reduced ceiling
Reliability (VRS-GEN-102)	38	56	Environmental and endurance test reports; no model-specific fleet data — availability and failure rate from the company-wide record
Serviceability (VRS-GEN-103)	44	49	Full service documentation, parts catalogue with lead times, SLA terms
Specification integrity (VRS-GEN-104)	29	37	Four headline claims with manufacturer measurements within tolerance; one claim short by 8 %
Cyber posture (VRS-GEN-105)	40	51	SBOM, signed updates, vulnerability policy; no third-party security assessment
Manufacturer support (VRS-GEN-106)	33	39	ISO 9001 certificate whose scope names the plant; warranty record supplied
Economics and residual (VRS-GEN-107)	22	29	Pricing and licence terms disclosed; residual data thin (new model)
IR	267	340	

- **Aggregation (6.1).** $61 + 38 + 44 + 29 + 40 + 33 + 22 = 267$. Every criterion counts at face value.
- **Ceilings (6.2).** The model is 73 points below the Verified ceiling of 340; 49 of those points sit in Safety and Reliability criteria whose Verified rungs require validation evidence the pack did not contain. A Certified session could unlock a further 60 points at most ($400 - 340$), all in witnessable criteria.
- **New model (6.5–6.6).** The model has 4 months of fleet history. Its field-record criteria were scored on the manufacturer's company-wide record at each criterion's stated reduced ceiling; the rating is a full rating with no qualification attached.
- **Display (4.2).** IR 267 (VRS 2026) – Verified .
- **Comparison (8.3).** A competitor rated IR 251 (VRS 2026) – Unverified scored 16 points lower with no manufacturer involvement; the two numbers are directly comparable, and the tier tells the reader how much evidence stands behind each.

12 Bibliography

- ISO 14224:2016, *Petroleum, petrochemical and natural gas industries – Collection and exchange of reliability and maintenance data for equipment* – reliability taxonomy (VRS-GEN-102).
- IEC 62443-4-2:2019, *Security for industrial automation and control systems – Part 4-2* – cyber-posture anchor (VRS-GEN-105).
- ISO 12100:2010, *Safety of machinery – General principles for design – Risk assessment and risk reduction* – risk-assessment anchor (VRS-GEN-101).
- ISO 8373:2021, *Robotics – Vocabulary* – robot-term source (via VRS-GEN-001).
- U.S. Green Building Council, *LEED v4.1 Building Design and Construction* (2019) – additive point-sum rating with certification levels by threshold; precedent for Clause 4 and 6.
- Euro NCAP, *Assessment Protocol – Overall Rating* (v9.x, 2023) – point-based scoring with edition-stamped protocols; precedent for Clause 8.
- BRE, *BREEAM UK New Construction* (2018, updated 2023) – credits summed to a percentage with fixed band thresholds; precedent for the sector pools of VRS-GEN-006.

13 Change history

DATE	VERSION	STATUS	CHANGE	AUTHORITY
2026-09-15	1.0 (draft)	Draft	VRS 2026 rewrite (Batch A). IR becomes the unweighted sum of 51 criterion points (maximum 400; ceilings 280 / 340 / 400 by tier) per VRS-GEN-101...107; Clause 5 weights replaced by point totals with shares as consequences; Clause 6 rewritten as aggregation + tier-ceiling rules (ceilings equal rungs, company-wide fallback, cold-start parity, certificates, null results, regime-aware ceilings, parent evidence, not-applicable-by-class, priors); Clause 7 confidence model deleted and replaced by the three rating tiers; Clause 8 calibration-before-rating replaced by additive editions; qualitative levels (Exemplary...Poor), provisional flag, renormalisation and confidence band retired library-wide; worked example re-run at Verified.	CEO approval 2026-09-15 13:28; vrs-rewrite-instructions-v1 §5.1
2026-09-15	0.3ε2	Draft	DEEPEN: §10 worked example (sparse-weight renormalisation, tier, confidence band).	Curator Run; charter DEEPEN
2026-09-11	0.3ε1	Draft	COHERENCE: bare CALIBRATION-PENDING markers → stage labels.	Curator Run #23
2026-09-10	0.3	Draft	Clause 8 delegated calibration to VRS-GEN-202; k_safety anchor correction.	CEO direction 2026-09-10
2026-09-05	0.2	Draft	Clause 5 category derivations split into VRS-GEN-101...107.	CEO goal 2026-09-05
2026-09-03	0.1	Draft	Created by decomposition of VRS-STANDARD-v0.1.	Framework Design Brief v1.1 §3.3

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-005 · Version 1.0