DRAFT

VRS
VEYRUM ROBOTICS STANDARD

TRUST. TRANSPARENCY. PERFORMANCE.

VEYRUM ROBOTICS STANDARD

VRS-GEN-105

IR Category: Cyber Posture

Scoring basis, cyber-posture criteria CY-1…CY-8, evidence by tier and determination procedure for the cyber-posture intrinsic category.

VEYRUM RESEARCH INSTITUTE

DOCUMENT	VERSION	STATUS	DATE
VRS-GEN-105	1.0	Draft	15 September 2026

Contents

13 Change history

Foreword. IR Category Standard for **cyber posture** (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 cyber criteria (CY-1…CY-8, 64 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 0.x drafts, which carried a capability/lifecycle formula, statistical exposure denominators and qualitative bands but no concrete scoring model: there is no rate formula and no letter cap. Cyber is the one loss driver whose failures *correlate across a fleet* — identical software makes one vulnerability an N-fold aggregation event — so the correlated/cyber-catastrophe reading is preserved as background (Clause 4) and as a passport signal (Clause 8), not as a score input. Requirements use “shall”.

1 Introduction

A robot model’s cyber posture is loss-relevant in a way no other category is: it is the one driver whose failures correlate across a fleet. A weak safety design harms one operator at a time; a shipped authentication flaw or an unpatched remote-access path is present in every identical unit at once, so a single adversary or worm can turn an idiosyncratic risk into a portfolio-wide aggregation event — precisely the correlated exposure an insurer prices with a catastrophe loading and a lender treats as a recoverability threat. This Standard turns “how secure is this model built and maintained to be?” into a repeatable number by asking eight questions — is its software inventory known and current, is access controlled, are updates trustworthy and reversible, is remote access contained, are vulnerabilities handled and patched, is data handling disclosed, are security events logged, and has a third party examined it — and awarding integer points for each. The points sum to a category total of 0–64. What a model can earn is bounded by how well its evidence is substantiated: public sources alone reach the **Unverified** ceiling, a manufacturer evidence pack reaches the **Verified** ceiling, and a witnessed demonstration reaches the **Certified** ceiling. The category rewards security that is designed in *and* shown, not merely asserted, and it reuses the conformity artefacts a manufacturer already produces (an IEC 62443-4-2:2019 capability evaluation, an IEC 62443-4-1:2018 lifecycle audit, an EU Cyber Resilience Act declaration, an ETSI EN 303 645 baseline) so that adoption is a mapping, not new testing.

2 1 Scope

1.1 This Standard defines, for the **cyber-posture** intrinsic category: what the category measures (Clause 4); the scoring basis (Clause 5); the eight criteria CY-1…CY-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 measures the model’s **inherent security capability and lifecycle posture** — what the model is built and maintained to withstand — not the configuration chosen at any one deployment site.

Whether a given operator isolated the network or rotated credentials is an Application-Rating matter and shall be scored only in the sector `-201` protocols; this Standard shall not score site-specific configuration. It excludes category weighting and aggregation into the Intrinsic Rating (VRS-GEN-005 §6), sector re-weighting (the `-201` protocols) and the evidence-grade definitions (VRS-GEN-012).

3 2 Normative references

- **VRS-GEN-001**, Vocabulary and Terminology — robot and security 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 62443-4-2:2019**, Security for industrial automation and control systems — Part 4-2: Technical security requirements for IACS components (component capability security levels, FR1–FR7).
- **IEC 62443-3-3:2013**, Part 3-3: System security requirements and security levels (SL 1–4).
- **IEC 62443-4-1:2018**, Part 4-1: Secure product development lifecycle requirements (maturity ML1–ML4).
- **ISO/IEC 30111:2019**, Information technology — Security techniques — Vulnerability handling processes.
- **ISO/IEC 29147:2018**, Information technology — Security techniques — Vulnerability disclosure.
- **ISO 10218-1:2025 / ISO 10218-2:2025**, Robotics — Safety requirements — cybersecurity clauses.
- **ETSI EN 303 645 V3.1.2 (2024-06)**, Cyber Security for Consumer Internet of Things — Baseline Requirements — the consumer-class mapping lane.
- **Regulation (EU) 2024/2847** (Cyber Resilience Act) — secure-by-default, vulnerability handling and support-period provisions in force at assessment.

4 3 Terms and definitions

Terms per VRS-GEN-001 and IEC 62443-4-2:2019. Locally:

- **software bill of materials (SBOM)** — the itemised list of software components and their versions in the rated build, against which known vulnerabilities are checked (CY-1).
- **capability security level** — the security level a component is capable of meeting for a foundational requirement when correctly configured (IEC 62443-4-2); background for what CY-1…CY-8 measure.
- **foundational requirement (FR)** — one of the seven IEC 62443 security requirement families (identification and authentication, use control, system integrity, data confidentiality, restricted data flow, timely response to events, resource availability); the capability picture Clause 4 summarises.
- **critical vulnerability** — a known defect scored CVSS ≥ 9.0 (or equivalent) exploitable in the model's supported configuration (CY-1, CY-5).

- **correlated exposure** — the fraction of the model's fleet running the same security-relevant build and network-reachable, over which a single vulnerability aggregates; a portfolio accumulation signal recorded to the passport, not a score input (Clause 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, audit or penetration-test report in lieu of a witnessed demonstration.

5 4 What this category measures

4.1 The category measures the model's **inherent security capability and lifecycle posture** — software transparency and currency, authentication and access control, update integrity and rollback, remote-access containment, vulnerability handling and patch cadence, data-handling disclosure, security logging, and independent security assessment — bounded by the evidence that substantiates each claim.

4.2 It measures both **design and evidence**: whether the model is built to resist and recover from attack (CY-2, CY-3, CY-4), whether its software surface is known and maintained (CY-1, CY-5), whether it is transparent and observable (CY-6, CY-7), and whether an independent party has examined it (CY-8). A declared capability that is only asserted is measured at the Unverified ceiling; the manufacturer evidence pack and the witnessed behaviour 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 security-software build range). Evidence shall not be pooled across builds that differ in a security-relevant capability unless the manufacturer declares them one security build and evidence supports that declaration.

4.4 The category surfaces the **correlated / cyber-catastrophe** loss driver: identical software across a fleet makes one vulnerability an N-fold aggregation event, which an insurer prices as a catastrophe/correlation loading and a lender treats as a recoverability threat. The fleet's correlated exposure and a lapsing security-support commitment are recorded to the passport as portfolio and obsolescence signals (Clause 8); neither changes the intrinsic point total.

6 5 Scoring basis

5.1 The category score shall be the **sum of integer points** awarded across the eight criteria CY-1…CY-8 (Clause 6). The maximum is 64 points. There is no capability or patch-latency 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 cyber-posture category all eight criteria apply to all classes, so the reachable maximum is always 64.

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	CY-1 Software bill of materials and component currency	0–8	6	8	8
6.2	CY-2 Authentication and access control	0–10	6	7	10
6.3	CY-3 Update integrity and rollback	0–10	6	7	10
6.4	CY-4 Remote access and network isolation	0–10	6	7	10
6.5	CY-5 Vulnerability handling and patch history	0–8	7	8	8
6.6	CY-6 Data handling and telemetry disclosure	0–5	4	4	5
6.7	CY-7 Security logging and audit	0–7	4	4	7
6.8	CY-8 Third-party security assessment	0–6	5	6	6
	Total	**64**	**44**	**51**	**64**

7.1 6.1 CY-1 Software bill of materials and component currency (0–8)

ROW	
What is measured	SBOM for the rated version; known-vulnerable components absent or mitigated.
Ladder	8 = SBOM + no unmitigated high/critical CVEs · **6 = publicly evidenced only (no pack, no demonstration)** · 5 = SBOM, mitigations partial · **2 = no SBOM, but all publicly known vulnerabilities for the product are mitigated per vendor advisories (F7)**, or component list only · 0 = none, or unmitigated known vulnerability
Unverified — evidence & ceiling	Public component disclosures vs CVE databases; vendor advisories. **Ceiling 6.**
Verified — evidence & ceiling	SBOM + CVE mitigation statement. **Ceiling 8.**
Certified — evidence & ceiling	From Verified evidence. **Ceiling 8.**
Company-wide fallback	n/a
Evidence bar	Rated version declared.
Applicability	All classes

7.2 6.2 CY-2 Authentication and access control (0–10)

ROW	
What is measured	No default credentials accepted; role separation; admin functions gated; **no unauthenticated control interfaces**.
Ladder	10 = shown live · 7 = documented in detail **and** covered by a public third-party security certificate/audit · 6 = access model documented · 3 = partial / a known unauthenticated interface exists · 0 = open/defaults
Unverified — evidence & ceiling	Manual / security documentation. **Ceiling 6.**
Verified — evidence & ceiling	Access-model description. **Ceiling 7.**
Certified — evidence & ceiling	Default credentials refused; roles shown on console. **Ceiling 10.**
Company-wide fallback	n/a
Evidence bar	None.
Applicability	All classes

7.3 6.3 CY-3 Update integrity and rollback (0–10)

ROW	
What is measured	Updates signed and verified; unsigned refused; rollback works; operator controls timing.
Ladder	10 = all shown · 7 = documented in detail (signing stated) **and** covered by a public third-party certificate/audit · 6 = mechanism documented · 3 = partial · 0 = unsigned/uncontrolled
Unverified — evidence & ceiling	Public update documentation. **Ceiling 6.**
Verified — evidence & ceiling	Update-mechanism description. **Ceiling 7.**
Certified — evidence & ceiling	Signed update applied, unsigned refused, rollback. **Ceiling 10.**
Company-wide fallback	n/a
Evidence bar	None.
Applicability	All classes

7.4 6.4 CY-4 Remote access and network isolation (0–10)

ROW	
What is measured	Remote access off by default, enablement audited; robot functions in segmented/offline mode as declared.
Ladder	10 = shown · 7 = documented in detail (offline/local mode, ports, defaults) **and** covered by a public third-party certificate/audit · 6 = documented · 3 = partial · 0 = always-on remote access
Unverified — evidence & ceiling	Network guidance / offline-mode documentation. **Ceiling 6.**
Verified — evidence & ceiling	Remote-access + isolation description. **Ceiling 7.**
Certified — evidence & ceiling	Remote access toggled; offline-mode operation shown. **Ceiling 10.**
Company-wide fallback	n/a
Evidence bar	None.
Applicability	All classes

7.5 6.5 CY-5 Vulnerability handling and patch history (0–8)

ROW	
What is measured	Disclosure channel, handling policy, time-to-patch for past issues.
Ladder	8 = policy + ≤ 30-day median patch · **7 = publicly evidenced only (no pack, no demonstration)** · 5 = policy + history · 2 = policy only · 0 = none
Unverified — evidence & ceiling	Public policy, public CVE response record. **Ceiling 7.**
Verified — evidence & ceiling	Policy + patch history. **Ceiling 8.**
Certified — evidence & ceiling	From Verified evidence. **Ceiling 8.**
Company-wide fallback	Company policy applies → **8**
Evidence bar	None.
Applicability	All classes

7.6 6.6 CY-6 Data handling and telemetry disclosure (0–5)

ROW	
What is measured	What is collected, where sent, retention, operator control.
Ladder	5 = shown on console and matches disclosure · 4 = documented · 2 = partial · 0 = undisclosed
Unverified — evidence & ceiling	Privacy/telemetry notice. **Ceiling 4.**
Verified — evidence & ceiling	Data-handling description. **Ceiling 4.**
Certified — evidence & ceiling	Console shows collection/destinations matching the description. **Ceiling 5.**
Company-wide fallback	n/a
Evidence bar	None.
Applicability	All classes

7.7 6.7 CY-7 Security logging and audit (0–7)

ROW	
What is measured	Security-relevant events (login, config change, update) logged and exportable.
Ladder	7 = shown + present in export · **4 = documented, or covered by a public third-party audit** · 0 = none
Unverified — evidence & ceiling	Security documentation. **Ceiling 4.**
Verified — evidence & ceiling	Logging description. **Ceiling 4.**
Certified — evidence & ceiling	Events visible on console and in the session export. **Ceiling 7.**
Company-wide fallback	n/a
Evidence bar	None.
Applicability	All classes

7.8 6.8 CY-8 Third-party security assessment (0–6)

ROW	
What is measured	IEC 62443 / ETSI EN 303 645 / penetration test coverage.
Ladder	6 = certified or full pen-test in last 24 months (prior) · **5 = publicly evidenced only (no pack, no demonstration)** · 4 = partial · 2 = self-assessment · 0 = none
Unverified — evidence & ceiling	Public certificate. **Ceiling 5.**
Verified — evidence & ceiling	Certificate + report summary. **Ceiling 6.**
Certified — evidence & ceiling	From Verified evidence. **Ceiling 6.**
Company-wide fallback	Company-level cert accepted → **6**
Evidence bar	None (0 valid).
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: public component disclosures checked against CVE databases and vendor advisories (CY-1), the security section of the user/administration manual (CY-2, CY-7), public update and network/offline-mode guidance (CY-3, CY-4), a published vulnerability-handling policy and public CVE-response record (CY-5), the privacy/telemetry notice (CY-6), and a public security certificate (CY-8). Each criterion's Unverified ceiling is the most an honest reading of such evidence can earn, and the category tier is **Unverified**. The "publicly evidenced only" rungs (CY-1 rung 6, CY-5 rung 7, CY-8 rung 5) fix what public evidence alone reaches when no pack and no demonstration are available.

7.2 Verified. The assessor additionally examines a manufacturer evidence pack: the SBOM with a CVE mitigation statement (CY-1), the access-model description (CY-2), the update-mechanism description (CY-3), the remote-access and network-isolation description (CY-4), the vulnerability-handling policy with a patch history (CY-5), the data-handling description (CY-6), the security-logging description (CY-7), and a third-party certificate with a report summary (CY-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 on the robot's own console: default credentials refused and roles shown (CY-2), a signed update applied and an unsigned update refused with a working rollback (CY-3), remote access toggled and offline-mode operation shown (CY-4), collection and destinations matching the disclosure (CY-6), and security events visible live and in the session export (CY-7). CY-1, CY-5 and CY-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, audit or penetration-test report in lieu of a witnessed demonstration (CY-2, CY-3, CY-4 rung 7; CY-7 rung 4; CY-8), an artefact identified by registrar entry, certificate number or report reference earns the stated rung; an unverifiable mark shall not. A mapped consumer-class baseline (ETSI EN 303 645) or a Cyber Resilience Act declaration is admissible on this basis where it covers the criterion.

7.5 Negative evidence. Absence of a defect is scored only where it is evidenced, not assumed: CY-1 rung 2 credits a product with no SBOM whose publicly known vulnerabilities are all mitigated per vendor advisories, and rung 0 records an unmitigated known vulnerability. Where no security evidence of any kind exists, the criterion scores 0 (5.4), never a null.

7.6 Company-wide fallback and parent-company evidence. Where a criterion offers a company-wide fallback (CY-5 → 8, CY-8 → 6) and model-specific evidence is unavailable, a manufacturer-wide vulnerability-handling policy or a company-level security certificate earns 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.

7.7 Regime-aware ceilings. A criterion's evidence may be satisfied through any recognised regime that covers it — an IEC 62443-4-2:2019 capability evaluation, an IEC 62443-4-1:2018 lifecycle audit, an ISO/IEC 30111:2019 / 29147:2018 vulnerability process, an ETSI EN 303 645 baseline, or a Regulation (EU) 2024/2847 declaration — read to the same ceiling regardless of jurisdiction; the regime is recorded, and the tier, not the regime, sets the ceiling.

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 cyber-posture category all eight apply and the reachable maximum is 64.

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 (CY-5, CY-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 64; otherwise it is 64 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 cyber-posture 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 portfolio signals of 8.1, so that an underwriter or lender can reproduce and audit the score from the passport alone. The per-criterion points carry the **cyber-loss frequency and severity** driver an insurer prices as a rate relativity and a security condition or warranty, and a lender reads into a recoverability assessment.

8.1 Portfolio and obsolescence signals. Cyber has no per-hour rate denominator; its exposure is **correlated**, and its denominator is an **accumulation base** — the count of identical, network-reachable units running the assessed build, of which the correlated exposure is the reachable fraction. This count, where obtainable, and a security-support commitment that will lapse within the assessment edition's validity shall be **recorded** to the passport so that portfolio underwriting can size a catastrophe/correlation loading and an accumulation limit (ASOP No. 38) and a lender can read an obsolescence and residual-value signal into a residual-value curve, LTV cap or re-securing covenant. Neither signal changes the intrinsic point total, which is context-independent; both are reported alongside it. No output of this category sets a price.

10 9 Worked example (fictional model)

"Northline AMR-7", a warehouse AMR on security build `4.2.x`, assessed at the **Verified** tier: a manufacturer evidence pack is examined but no behaviour is witnessed. All eight criteria apply (reachable maximum at Verified = 51). The pack shows an SBOM with a CVE mitigation statement and no unmitigated high/critical CVEs; a detailed access-model description; a signed-update mechanism with rollback, documented in detail; a remote-access and offline-mode description with default ports and defaults; a published vulnerability-handling policy with a patch history showing a median critical-patch time near 22 days; a data-handling description; a security-logging description; and a public IEC 62443-4-

2:2019 component-evaluation certificate with a report summary. The model's fleet correlated exposure is recorded as 0.9 for the portfolio signal.

§	CRITERION	RUNG EARNED	VERIFIED CEILING	POINTS
6.1	CY-1	8 (SBOM + no unmitigated high/critical CVEs)	8	8
6.2	CY-2	7 (documented in detail + public third-party certificate)	7	7
6.3	CY-3	7 (signing documented + public certificate)	7	7
6.4	CY-4	7 (offline mode/ports/defaults documented + certificate)	7	7
6.5	CY-5	8 (policy + ≤ 30-day (prior) median patch)	8	8
6.6	CY-6	4 (data-handling documented)	4	4
6.7	CY-7	4 (logging documented)	4	4
6.8	CY-8	6 (IEC 62443-4-2 certificate)	6	6
	Total		**51**	**51**

The cyber-posture total is **51 of a reachable 51** at the Verified tier. Arithmetic: 8 + 7 + 7 + 7 + 8 + 4 + 4 + 6 = 51. The reading is that this model earns everything reachable at the Verified tier but leaves real headroom that only a Certified assessment unlocks: CY-2, CY-3 and CY-4 each carry a top rung of 10 = "shown live" (default credentials refused on the console, an unsigned update refused with a working rollback, remote access toggled with offline operation), CY-6 carries a rung 5 for console-confirmed telemetry, and CY-7 a rung 7 for security events visible live and in the session export — together 13 further points (from 51 to the 64 maximum) that require witnessing the behaviour, not just reading the pack. No company-wide fallback was used (model-specific policy and certificate evidence were present), and the not-applicable set is empty. The correlated exposure of 0.9 is recorded as the portfolio accumulation signal (8.1) and does not alter the score. 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 + security-build granularity confirmed; evidence not pooled across differing security builds (4.3).
- [ ] Assessment tier fixed per criterion; not-applicable set recorded (Step 1); reachable maximum recorded.
- [ ] **CY-1** SBOM and component currency scored; rated version declared (evidence bar).
- [ ] **CY-2** authentication and access control scored; live-console rung only at Certified.

- [ ] **CY-3** update integrity and rollback scored; live demonstration rung only at Certified.
- [ ] **CY-4** remote access and network isolation scored; witnessed toggle/offline rung only at Certified.
- [ ] **CY-5** vulnerability handling and patch history scored; (prior) thresholds read as written; company policy fallback recorded if used.
- [ ] **CY-6** data handling and telemetry disclosure scored; console-match rung only at Certified.
- [ ] **CY-7** security logging and audit scored; live + export rung only at Certified.
- [ ] **CY-8** third-party security assessment scored; certificate/report verifiable; company-level cert fallback recorded if used.
- [ ] Certificate-rung artefacts verified by registrar entry, certificate number or report reference (7.4).
- [ ] Each award capped at the assessed-tier ceiling (5.3); no-evidence criteria scored 0, never null (5.4).
- [ ] Correlated exposure and any lapsing support commitment recorded as portfolio/obsolescence signals, not folded into the score (8.1).
- [ ] Score, per-criterion points, tier and ceiling reached recorded to the passport (Step 6).

12 Bibliography

- IEC 62443-4-2:2019, *Security for industrial automation and control systems — Part 4-2: Technical security requirements for IACS components*.
- IEC 62443-3-3:2013, *Part 3-3: System security requirements and security levels*.
- IEC 62443-4-1:2018, *Part 4-1: Secure product development lifecycle requirements*.
- ISO/IEC 30111:2019, *Information technology — Security techniques — Vulnerability handling processes* (reviewed and confirmed 2025).
- ISO/IEC 29147:2018, *Information technology — Security techniques — Vulnerability disclosure*.
- ISO 10218-1:2025 and ISO 10218-2:2025, *Robotics — Safety requirements* (cybersecurity clauses).
- ETSI EN 303 645 V3.1.2 (2024-06), *Cyber Security for Consumer Internet of Things — Baseline Requirements*.
- Regulation (EU) 2024/2847 of 23 October 2024 (Cyber Resilience Act), *horizontal cybersecurity requirements for products with digital elements*.
- NIST SP 800-82 Rev. 3 (2023), *Guide to Operational Technology (OT) Security*.
- FIRST, *Common Vulnerability Scoring System (CVSS) v3.1 Specification* (2019).
- ASOP No. 38, *Catastrophe Modeling (for All Practice Areas)* (Actuarial Standards Board; effective for work performed on or after 1 December 2021) — the correlated/accumulation basis for the cyber-catastrophe portfolio signal.

13 Change history

DATE	VERSION	STATUS	CHANGE	AUTHORITY
2026-09-15	1.0	Draft	VRS 2026 rewrite: replaced the capability/lifecycle formula model (input schema, FR-vector and lifecycle coding, correlated-exposure and accumulation-base denominators, assessment-window/recency/small-evidence rules, the `P_cap` / `P_life` formulae and override caps, and the qualitative L1–L5 level descriptors) with eight point-scored criteria CY-1…CY-8 (64 points; Clause 6) with per-tier ceilings — the category's first concrete scoring model; added Clause 5 Scoring basis, Clause 7 Evidence by tier (certificate rung, negative evidence, company-wide fallback, parent-company evidence, regime-aware ceilings), Clause 8 point-sum determination with passport binding and the correlated/obsolescence portfolio signals, regenerated worked example (Verified tier, 51/51) and assessor checklist; retired confidence framing, provisional flags, and qualitative IR levels; dropped the GEN-201/GEN-202 references and the calibration/statistics bibliography that served the formula model; see GEN-005 §6	VRS 2026 rewrite Stage 4
2026-09-13	0.4	Draft	RESEARCH (charter §2): added the named international-standard evidence-reuse lane for the `vuln_handling` input — ISO/IEC 30111:2019 and ISO/IEC 29147:2018 — the artefacts a maker already operates to evidence a CRA Annex I Part II posture. No constant, weight, threshold, band or enum value changed. Net +~70 words	RESEARCH run #98
2026-09-12	0.3	Draft	ACTUARIAL review: added passport binding + exposure denominator (Step 7 output recorded to the VRS-GEN-009 cyber record, auditable from the passport alone; exposure = correlated-exposure accumulation denominator, not op-hours) and the loss-driver/decision map; added ASOP No. 23 and ASOP No. 38 anchors; added a checklist line. No constant value changed	CEO direction 2026-09-11
2026-09-11	0.2	Draft	Expanded to publishable depth and reframed as intrinsic (capability + lifecycle, not as-deployed config): input schema incl. IEC 62443-4-2 SL-C vector, FR/lifecycle coding + evidence-reuse mapping, correlated-exposure/loss-driver framing, currency/window/small-evidence rules, 7-step procedure with priors + override caps, worked example, assessor checklist; dated cyber anchors	CEO direction 2026-09-11

DATE	VERSION	STATUS	CHANGE	AUTHORITY
2026-09-05	0.1	Draft	Split from VRS-GEN-005 §5.5; 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

Contact: hello@veyrum.com · VRS-GEN-105 · Version 1.0