

HU-RSC

Humanide Robotics Safety Class

A classification framework for the physical safety
of smart robotics in the German-speaking SME market

Discussion paper

Version: v0.2 rev9 (May 2026, Discussion Paper)

Date: 16 May 2026

Issuer: Humanide GmbH, Cologne (commercial register: HRB 115169)

Companion model: HU-RAC — Humanide Robotics Assurance Class (data and resilience), v0.3.1

Note: A German edition is available as HU-RSC Impulspapier v0.2.11 (May 2026). EU regulatory references are retained; a separate UK-localised edition (UKCA, PSTI Act, UK supply of machinery) is planned.

Status: Public discussion paper. Use and circulation under attribution to source (Humanide GmbH, HU-RSC v0.2 rev9, May 2026) requested.

Contents

- 0 Methodological disclaimer and reading notes.....4
 - 0.0 Purpose, audience, and scope.....4
 - 0.1 What HU-RSC is — and what it is not.....4
 - 0.2 How thresholds were derived.....5
 - 0.3 Relationship to other Humanide documents.....5
- 1 Position within the classification system (HU-RAC × HU-RSC).....6
 - 1.1 Two models, one system.....6
 - 1.2 What HU-RSC does not cover.....6
 - 1.3 Seven use-case domains in scope.....6
- 2 Regulatory framework.....8
 - 2.1 Machinery Regulation (EU) 2023/1230.....8
 - 2.2 AI Act (EU) 2024/1689.....8
 - 2.3 Product safety and product liability.....9
 - 2.4 Cyber Resilience Act (Regulation (EU) 2024/2847).....10
 - 2.5 Relevant safety standards.....10
- 3 Architecture of the model.....12
 - 3.1 Parallel structure with HU-RAC.....12
 - 3.2 Extensions specific to HU-RSC.....12
 - 3.3 Tier scale — a note on terminology.....12
- 4 The six axes — assessment content in detail.....14
 - 4.1 Axis 1 — Deployment environment.....14
 - 4.2 Axis 2 — Form factor, motion envelope, and kinetic injury potential.....14
 - 4.3 Axis 3 — Types and severity of harm.....15
 - 4.4 Axis 4 — Degree of autonomy.....15
 - 4.5 Axis 5 — Safety architecture (three pillars).....16
 - 4.5.1 Pillar 1 — Classical safety architecture.....16
 - 4.5.2 Pillar 2 — AI safety architecture.....17
 - 4.5.3 Pillar 3 — Lifecycle update infrastructure.....17
 - 4.6 Axis 6 — Exposed persons and exposure frequency.....18
- 5 Tier scale 1–5 with thresholds.....19
 - 5.1 Tier 1 — minimal residual risk.....19
 - 5.2 Tier 2 — limited residual risk.....19
 - 5.3 Tier 3 — elevated residual risk.....20
 - 5.4 Tier 4 — high residual risk.....20
 - 5.5 Tier 5 — critical residual risk.....20
- 6 Expectations matrices (E-Min, E-Plus, R).....22

6.1	Expectations matrix Axis 1 — Deployment environment.....	22
6.2	Expectations matrix Axis 2 — Form factor and kinetic injury potential.....	23
6.3	Expectations matrix Axis 3 — Types and severity of harm.....	24
6.4	Expectations matrix Axis 4 — Degree of autonomy.....	25
6.5	Expectations matrix Axis 5 — Safety architecture (three pillars).....	26
6.6	Expectations matrix Axis 6 — Exposed persons and exposure frequency.....	29
7	Secondary dimensions.....	31
7.1	Business criticality A/B/C.....	31
7.2	Supplier resilience R1/R2/R3.....	31
7.3	Relationship with Axis 5 Pillar 3.....	32
8	Aggregation and reclassification rules.....	33
8.1	Aggregation across the six axes.....	33
8.2	Aggregation within Axis 5 — the minimum-floor rule.....	33
8.3	Reclassification on substantial modification.....	34
9	Minimum coupling HU-RSC × HU-RAC.....	35
9.1	The asymmetric minimum coupling.....	35
9.2	Rationale for the coupling.....	35
9.3	Asymmetry.....	35
9.4	Operational consequences.....	36
10	Application in procurement practice.....	37
10.1	The three-step logic.....	37
10.2	Connection to HU-RAC and the Requirements Catalogue.....	37
10.3	Worked example — humanoid robot in an SME workshop.....	37
11	Final remarks.....	39
12	Annex — mapping tables for ISO classes.....	40
12.1	Notes on using the mapping table.....	40
13	Annex — Glossary.....	41
14	Annex — References.....	43

0 Methodological disclaimer and reading notes

0.0 Purpose, audience, and scope

HU-RSC, together with its companion model HU-RAC, is a communication framework. It is developed and used internally by Humanide GmbH as a shared vocabulary between three parties:

- Humanide GmbH itself, as the entity placing products on the market under its brands (Mowio, Wuffi, Robofixx, Robokitchen, Omtanke);
- service providers, distributors and manufacturers who supply or co-deploy components and platforms to Humanide GmbH;
- buyers, including the procurement and compliance functions of customer organisations, who evaluate Humanide products against their own internal requirements.

Its purpose is to enable a structured conversation about the data sovereignty profile (HU-RAC) and the physical safety profile (HU-RSC) of a robotics system before procurement, at a time when no comprehensive, publicly recognised standard yet covers humanoid robotics, AI-driven safety functions and the full lifecycle of cyber-physical robotics in an integrated manner.

HU-RSC is therefore neither a conformity assessment scheme, a certification, nor a competitor to ISO, IEC, EN or other industry standards. Statements of the form ‘this product reaches HU-RAC X / HU-RSC Y’ are internal classifications by Humanide GmbH; third parties who adopt these classifications do so voluntarily as a shared reference frame. The framework is explicitly designed to be superseded — in whole or in part — by harmonised standards as they become available (ISO 25785-1, ISO/IEC TS 22440, and others).

Disclaimer of warranty and liability. This paper is not a scientific publication, not a legal opinion, and not a binding conformity assessment. It is a working document of Humanide GmbH made publicly available for transparency and discussion purposes. The content has not been validated by external legal counsel prior to publication. Humanide GmbH provides the content without any warranty, express or implied, regarding accuracy, completeness, currency, fitness for any particular purpose, or non-infringement of third-party rights. Statements about the regulatory or normative landscape may become outdated without notice.

To the maximum extent permitted by applicable law, Humanide GmbH and the natural persons acting for it disclaim all liability for any direct, indirect, incidental, consequential, special or punitive damages arising out of or in connection with the use of, or reliance on, the content of this paper. Readers and adopters act on the content at their own risk. The legally binding judgement on conformity, safety and data protection remains with the responsible parties (manufacturers, notified bodies, qualified experts, supervisory authorities) and is not substituted by an HU-RSC tier. Should any provision of this disclaimer be or become invalid or unenforceable, the validity of the remaining provisions shall remain unaffected.

0.1 What HU-RSC is — and what it is not

HU-RSC is a procurement-stage assessment framework, not a safety-engineering tool.

An HU-RSC tier classification **does not substitute for the ISO 12100 risk assessment**, and it is not a conformity assessment within the meaning of the Machinery Regulation (EU) 2023/1230. Responsibility for conformity, risk assessment, the duties of those who place products on the market, and the duties of operators under the German Industrial Safety Ordinance (BetrSichV) and equivalent national rules

remains entirely with the parties on whom these duties fall.

HU-RSC supports buyers and operators in making a structured judgement, *before purchase*, about the level of safety a robotics system needs to achieve, and whether a particular product plausibly meets that level. The legally binding assessment is performed by the responsible parties (manufacturers, notified bodies, qualified experts) and lies outside the scope of this model.

0.2 How thresholds were derived

The numerical tier thresholds in this model — kinetic energy in joules, minimum PL/SIL per tier, severity classes — are derived from the relevant standards: ISO 12100 Table B.3 for the severity scale, ISO 13849-1 for performance levels, ISO/TS 15066 for human-robot collaboration thresholds, and IEC 61508 for safety integrity levels. They are not derived from the example products of the suppliers we assess.

This sequence matters. Deriving thresholds from products would mean calibrating the model against the suppliers who happen to be on the market today — a textbook case of confirmation bias. Supplier validation tests the thresholds we have set in advance; it does not define them.

0.3 Relationship to other Humanide documents

This discussion paper defines the HU-RSC model. Concrete supplier-specific assessments are not held here but in the internal Requirements Catalogue v2.x (Document 01) and the external Requirements Profile v2.x (Document 03). This separation keeps the discussion paper stable over time — supplier assessments can be updated without republishing the model itself.

The companion model HU-RAC covers data sovereignty and supplier resilience and is positioned in detail in Chapter 1.

1 Position within the classification system (HU-RAC × HU-RSC)

1.1 Two models, one system

The Humanide classification system consists of two nominally independent, structurally parallel models with one asymmetric coupling. HU-RAC assesses the digital sovereignty and supplier resilience of a robotics system; HU-RSC assesses its physical safety in relation to people. The two models share their architecture: a tier scale of 1 to 5, six axes, an aggregation principle, and the structure of the expectations matrices.

This parallel structure is not cosmetic. It exists so that buyers can apply a single mental model to both. Anyone familiar with HU-RAC can pick up HU-RSC without a fresh learning curve, and the other way round. External communication takes a single form: 'This product reaches HU-RAC X / HU-RSC Y'.

1.2 What HU-RSC does not cover

Three assessment dimensions sit deliberately outside HU-RSC and are handled elsewhere:

- **Digital sovereignty** — Data protection and supplier resilience in the sense of data sovereignty — the subject of HU-RAC.
- **Cybersecurity** — Cybersecurity is held exclusively in HU-RAC Matrix 5.3. HU-RSC does not address cybersecurity directly. It addresses it indirectly, through the minimum coupling described in Chapter 9, because safety-relevant telemetry has to be integrity-protected.
- **Functional capability** — Payload, MTBF, reach, and degree of autonomy as a performance feature are handled in Part 5 of the Requirements Catalogue as a context-specific assessment per use case, not as a tier that is universally comparable.

Damage to property is not handled in HU-RSC. It appears as a secondary expectation in Part 5 of the Requirements Catalogue. HU-RSC focuses on personal safety in the sense of the ISO 12100 logic.

1.3 Seven use-case domains in scope

HU-RSC is calibrated to the robotics application domains that matter commercially in German-speaking SMEs over the next few years:

- **Consumer-facing robotics** — Lawn-mowing, vacuuming, and cleaning robots in the consumer and light-commercial market (B2C/SoHo).
- **Service robotics in semi-public settings** — Cleaning, inspection, and security robots in hotels, hospitals, care homes, and schools.
- **Industrial robotics in SMEs** — Production, logistics, and inspection robots in the SME industrial space, outside the traditional caged-cell setting of large industry.
- **Outdoor robotics** — Outdoor robots in agricultural, municipal, and commercial outdoor settings.

- **Humanoid robotics (pilot phase)** — Early pilot deployments of humanoid robots in SMEs (e.g. pick-and-place, workpiece feeding).
- **Assistive robotics** — Personal assistance, care support, and edutainment.
- **Mobile logistics robotics** — AGVs and AMRs in the SME logistics context.

The tier scale is calibrated so that all seven domains are covered. Specialist domains such as medical devices (MDR), the perimeter of nuclear installations, and military robotics fall outside this scope. Note on terminology: a humanoid robot is a robotic system with anthropomorphic morphology (upper body and two arms, plus two legs or other suitable means of locomotion), equipped with environment-sensing capability and intended to perform a variety of human-designed tasks; it may operate autonomously or be teleoperated. The robot may include grippers; otherwise grippers are part of the surrounding humanoid robot system, which comprises at least one humanoid robot, end-effectors, workpieces, peripheral components, and charging infrastructure. Adapted from Fraunhofer IPA KMUmanoid (2025).

Real-world placement reference (consumer-facing robotics). A typical Mowio-class robotic lawn-mower with a rotating cutting blade in a private garden, operating without local RTK infrastructure but with cloud-based positioning, places on Axis 2 around Tier 3 (kinetic energy in the 10–80 J range plus rotating tool motion) and on Axis 1 around Tier 3 (semi-controlled environment with possible third-party encounter). Axis 6 typically reaches Tier 3 with pet encounters as a residual risk. The overall placement at Tier 3 is plausible for most current market offerings.

2 Regulatory framework

HU-RSC sits within a bundle of regulatory instruments that will tighten markedly between 2026 and 2027. This chapter sketches the rules and standards that matter for buyers and operators. It does not aim at legal completeness — for concrete compliance questions, professional legal or expert advice is required.

2.1 Machinery Regulation (EU) 2023/1230

The Machinery Regulation 2023/1230 replaces the Machinery Directive 2006/42/EC and applies from 20 January 2027 under Article 54. It applies directly as a Regulation; no national transposition is required. For robotics it brings several material tightenings:

The applicable date follows the corrigendum in OJ L 169 of 4 July 2023; the original publication in OJ L 165 of 29 June 2023 contained a calculation error ('14 January 2027') which has since been officially corrected.

- **Annex I categories — Annex I Part A lists categories of machinery requiring mandatory third-party conformity assessment. Point 5 explicitly covers 'safety components with fully or partially self-evolving behaviour using machine learning approaches ensuring safety functions'.**
- **Substantial modification —** Substantial modifications to machinery already placed on the market — for instance through software updates that change safety-relevant or functional behaviour — may trigger a fresh conformity assessment. The detailed interpretation is still being worked out.
- **Safety component definition — Article 3(3) defines 'safety component' as 'a physical or digital component, including software, of a product within the scope of this Regulation, which is designed or intended to fulfil a safety function and which is independently placed on the market, the failure or malfunction of which endanger the safety of persons, but which is not necessary in order for that product to function'. Software-based safety components are explicitly in scope.**
- **Annex III EHSRs — Annex III contains the essential health and safety requirements, including section 1.1.6 on self-evolving systems, section 1.1.9 ('Protection against corruption'), and section 1.2.1 ('Safety and reliability of control systems'). Where this discussion paper later references 'MR Annex III' as a tier-level requirement, the reference is to the EHSRs as a whole, with particular emphasis on the AI and cybersecurity sections.**

Commission guidance status. The European Commission is preparing a guidance document and an implementing decision listing harmonised standards. Drafts on 'AI-supported safety functions' and 'software as a safety component' are expected mid-2026; the implementing decision on harmonised standards is expected by the end of 2026. As of March 2026, industry has withdrawn its request for postponement of the AI provisions of the Machinery Regulation but continues to request alignment of the cybersecurity requirements with the Cyber Resilience Act application date of 11 December 2027 — a 10-month gap between MR application (20 January 2027) and CRA application (11 December 2027) that procurement decisions must factor in.

2.2 AI Act (EU) 2024/1689

The AI Act entered into force in August 2024; its obligations apply in stages under Article 113. For HU-RSC two paths must be distinguished. Under Article 6(1) — AI as a safety component of a product covered by Annex I harmonisation legislation, including the Machinery Regulation — the high-risk obligations apply from 2 August 2027. Under Article 6(2) — AI systems referred to in Annex III — the high-risk obligations apply from 2 August 2026. Most consumer-facing and industrial robotics with safety-relevant AI functions fall in the Article 6(1) path. Under its Digital Omnibus proposal of 19 November 2025, the European Commission put forward postponements: Annex III high-risk obligations to 2 December 2027, and Annex I-embedded AI systems to 2 August 2028. Council and Parliament adopted negotiating positions in spring 2026 broadly supporting these postponed dates; formal adoption was still outstanding at the time of writing. Until the formal text appears in the Official Journal, the Article 113 dates above remain the binding default. Article 6(5) of the AI Act obliges the Commission to provide guidelines on the practical implementation of Article 6 by 2 February 2026 together with a comprehensive list of practical use-case examples; the Commission missed this deadline, and a final draft for consultation was subsequently absorbed into the Digital Omnibus negotiations. Procurement decisions in 2026 are therefore made without binding Commission guidance on the high-risk classification of robotics AI. Three points are central for HU-RSC:

- **High-risk AI in the machinery context** — Robotics systems that act as a safety component for machinery within the meaning of the Machinery Regulation typically fall into the high-risk class of the AI Act. The duties cover risk management, data quality, transparency documentation, and human oversight mechanisms.
- **Explainability** — The AI Act requires a degree of explainability of decision logic that has not yet been achieved across the board for many robotics AI systems — particularly for vision-language-action models (VLAs).
- **Interface between MR and AI Act** — How the AI Act and the Machinery Regulation (Annex III) fit together is still an open interpretive question. ISO/IEC TR 5469 'Artificial Intelligence — Functional Safety and AI Systems' offers a first technical framework.

2.3 Product safety and product liability

Three further instruments frame the procurement decision:

- **GPSR** — Regulation (EU) 2023/988 on general product safety has applied since December 2024. It supplements the sector-specific rules, especially for consumer products such as consumer-facing robotics.
- **Revised PLD** — **Directive (EU) 2024/2853 must be transposed into German law by 9 December 2026 under Article 22 of the Directive. The German Federal Ministry of Justice published a draft law (Referentenentwurf) on 11 September 2025. The revised regime extends 'product' explicitly to include software (Article 4(1) of the Directive). The previous 500-EUR deductible and the 85-million-EUR liability ceiling are abolished. A claimant-favourable disclosure**

obligation for evidence is introduced. For robotics manufacturers and integrators, this is a structural increase in product-liability exposure.

- **BetrSichV and duties to safeguard third parties** — In Germany, the Industrial Safety Ordinance (Betriebssicherheitsverordnung, BetrSichV) governs the duties of the operator (employer) when work equipment, including robotics, is used. Duties to safeguard third parties (Verkehrssicherungspflichten) apply in any area to which third parties have access.
- **Radio Equipment Directive (RED)** — Directive 2014/53/EU applies to any robotics system with radio-frequency components (4G/5G, Wi-Fi, Bluetooth, GPS receivers). In Germany, the Bundesnetzagentur is the market surveillance authority. RED conformity is assessed separately from Machinery Regulation conformity. For HU-RSC scope, RED applies in particular to consumer-facing robotics, mobile logistics robotics, and humanoid robots with wireless connectivity.

2.4 Cyber Resilience Act (Regulation (EU) 2024/2847)

The Cyber Resilience Act addresses products with digital elements. Under Article 71, the Regulation applies from 11 December 2027. Two earlier dates apply by Article 71: Chapter IV (notification of conformity assessment bodies) applies from 11 June 2026, and Article 14 (reporting obligations) applies from 11 September 2026. Procurement decisions in 2026 must already factor in the Article 14 reporting obligations for products placed on the market thereafter. The points relevant for HU-RSC are:

- **Risk assessment** — Robotics systems with digital elements must undergo a cybersecurity risk assessment.
- **Update obligation** — Manufacturers must provide security updates over the expected lifetime — Article 13 CRA establishes a support period of at least five years; where the product is expected to be in use for less than five years, the support period corresponds to the expected use time. Article 13 also empowers the Commission to adopt delegated acts specifying the minimum support period for specific product categories where market surveillance data suggests inadequate support periods.
- **Reporting obligations** — Article 14 CRA requires manufacturers to notify actively exploited vulnerabilities via the single reporting platform (Article 16) simultaneously to the designated CSIRT coordinator and to ENISA: an early warning notification within 24 hours of becoming aware, and a full vulnerability notification within 72 hours. From a procurement perspective, the supplier must demonstrably have the operational capacity to meet these reporting timelines.

How the CRA and the Machinery Regulation interact around the term ‘substantial modification by software update’ is a central open interpretive question. HU-RSC reflects this interaction through Axis 5 Pillar 3 (lifecycle update infrastructure).

2.5 Relevant safety standards

The following standards inform the tier thresholds and the expectations matrices:

- **ISO 12100** — ISO 12100 — fundamental risk assessment; Table B.3 provides the severity scale used in HU-RSC Axis 3.
- **EN ISO 13849-1:2023** — Safety of machinery — Safety-related parts of control systems — Part 1: General principles for design. Fourth edition published 2023, listed in the OJEU in May 2024. The previous edition EN ISO 13849-1:2015 remains relevant during the OJEU transition period; its presumption of conformity expires on 15 May 2027, while CEN lists a Date of Withdrawal of 31 May 2026. Major 2023 changes include subsystem terminology (replacing SRP/CS), an explicit Safety Requirements Specification chapter, an EMC immunity Annex L, and an alternative PL determination method. Informs HU-RSC Axis 5 Pillar 1.
- **ISO 10218** — ISO 10218-1 and ISO 10218-2 (revised editions, February 2025) — safety for industrial robots and robot applications; relevant to the SME industrial context.
- **ISO/TS 15066** — ISO/TS 15066 — safety of collaborative robots; the collision thresholds inform HU-RSC Axis 2 (kinetic injury potential).
- **ISO 25785-1 — Robotics — Safety requirements for dynamically stable industrial mobile robots. Currently Working Draft, developed by ISO TC 299 WG 12 with industry participants from Agility Robotics, Boston Dynamics, A3 Association and Fraunhofer IPA. Final publication expected in 2027 or 2028. When published, this will become the primary Type C reference for humanoid pilot deployments under HU-RSC; the mapping table in Annex 12 will be updated accordingly.**
- **ISO/IEC TR 5469** — ISO/IEC TR 5469 — functional safety of AI systems; relevant to HU-RSC Axis 5 Pillar 2.
- **IEC 61508** — IEC 61508 — Safety Integrity Levels (SIL 1–4); relevant for safety-related electrical and electronic systems.
- **IEC TS 62998-1** — Safety of machinery — Safety-related sensors used for the protection of persons; relevant for HU-RSC Axis 5 Pillar 1 volumetric safety monitoring (cameras, LIDAR, radar) and required reference for sensor-based safe human detection from Tier 3 upwards.
- **ISO/IEC TS 22440** — Artificial Intelligence — Functional Safety and AI Systems (Parts 1, 2, 3); Committee Draft as of February 2026, comment period closed April 2026; will supersede ISO/IEC TR 5469 as the binding Technical Specification for AI functional safety, with publication expected in 2027 or 2028. Together with ISO 25785-1, the most relevant upcoming standards for HU-RSC Pillar 2.
- **ISO 13482** — Safety requirements for personal care robots; relevant to HU-RSC assistive-robotics and semi-public service-robotics domains; covers humanoids only for non-industrial use cases.
- **ISO 3691-4** — Industrial trucks — Safety requirements for driverless industrial trucks; relevant to HU-RSC mobile logistics-robotics domain (AGVs, AMRs). Humanoid robots are not in scope of ISO 3691-4.

3 Architecture of the model

Methodological note on this discussion paper (v0.2 rev9). The paper expresses three model-formation decisions that are deliberately left provisional in this discussion draft: (1) the selection of six assessment axes (Chapter 4); (2) the choice of an HU-RSC-native tier scale 1–5 with ISO mapping tables, rather than direct aggregation of existing ISO classes (Chapter 11); (3) the asymmetric minimum coupling $HU-RSC \geq 4$ requires $HU-RAC \geq 3$ (Chapter 9). All three may be revised in v0.3 in response to practitioner feedback. Adopters of HU-RSC should treat them as tendencies of the current model, not as finalised definitions.

3.1 Parallel structure with HU-RAC

HU-RSC carries the same model architecture as HU-RAC. The following architectural features are identical for both models, and we adopt them here without restating the rationale:

- **Axes** — six assessment axes with a tier scale of 1 to 5
- **Aggregation** — the weakest-link principle for aggregation across axes — the worst-performing axis sets the tier of the system as a whole
- **Matrices** — expectations matrices E-Min, E-Plus, and R per axis and tier
- **Secondary dimensions** — two secondary dimensions: business criticality (A/B/C) and supplier resilience (R1/R2/R3)
- **Dynamics** — dynamic reclassification on substantial modification within the meaning of the Machinery Regulation 2023/1230

3.2 Extensions specific to HU-RSC

Three extensions distinguish HU-RSC from HU-RAC architecturally:

- **Three-pillar Axis 5** — One of the six axes — Axis 5 (safety architecture) — is built on three pillars: classical safety functions to SIL/PL, AI safety architecture, and lifecycle update infrastructure. Aggregation within this axis follows a modified rule (see Chapter 8).
- **Asymmetric minimum coupling** — There is a one-way minimum coupling between HU-RSC and HU-RAC: $HU-RSC \geq 4$ requires $HU-RAC \geq 3$. The rationale appears in Chapter 9.
- **Distinct R-column function** — In HU-RSC, the R column of the expectations matrices carries an explicit functional clarification as a duty-of-care warning. It is neither a disclaimer nor a model gap. The reasoning is set out in Chapter 6.

3.3 Tier scale — a note on terminology

The HU-RSC tier scale of 1 to 5 measures residual risk after protective measures, not the inherent risk of the technology before any measures are applied. A humanoid robot with an elaborate safety architecture can reach a lower tier than a lawn-mower without any meaningful safety architecture, even though the humanoid is inherently more dangerous.

This reading follows the ISO 12100 risk-assessment logic: protective measures reduce risk, and what remains after measures is what must be assessed and communicated. A

terminology note: HU-RSC tiers 1 to 5 denote residual risk levels of a deployed system. Vendor materials (for instance Synapticon's Positron Tier 1 to 3) use 'tier' for architectural layers of a safety stack — these are orthogonal to HU-RSC tiers and should not be confused.

4 The six axes — assessment content in detail

This chapter sets out each of the six axes in detail. Every axis carries one dimension of physical safety and is orthogonal to the other five — in the sense of the model's logic, not in any strict mathematical sense. Overlaps with HU-RAC are excluded by design: in particular, cybersecurity is held only in HU-RAC Matrix 5.3, never in HU-RSC.

The axes are presented in the order in which a buyer would naturally read them: first the deployment context (Axis 1), then the system itself (Axes 2 to 5), then the people exposed to it (Axis 6).

4.1 Axis 1 — Deployment environment

Axis 1 distinguishes the environments in which the system operates and how those environments shape safety-relevant factors. Three classes guide the assessment:

- **Controlled environment** — fully controlled (safety cage, fenced industrial cell, reliable lighting, defined paths, no human access during operation).
- **Semi-controlled environment** — human access is possible without continuous physical separation, but with organisational protective measures (private home, hospital corridor outside patient rooms, hotel corridor, workshop with employees present).
- **Public, unprotected environment** — human access is not protected, third parties without prior briefing are present, and children, animals, and road users may appear (public outdoor areas, parks, pavements, shopping centres).

Secondary factors that can shift the class: lighting (day, night, mixed), floor conditions (level, uneven, slippery), exposure to weather (indoor, outdoor), noise levels (whether acoustic warnings can be heard), and visibility conditions (mirrors, blind corners).

4.2 Axis 2 — Form factor, motion envelope, and kinetic injury potential

Axis 2 assesses the physical injury potential of the system in the event of contact with a person or a collision with an object. Four factors are central:

- **Kinetic energy** — impact energy as $E_{\text{kin}} = \frac{1}{2} m v^2$. Reference values come from ISO/TS 15066 (collision thresholds for human-robot collaboration).
- **Motion envelope** — the size of the free motion envelope, the reach of the gripper, the maximum lift height, the swing radius.
- **Tools and impact surfaces** — rotating blades (lawn-mowers), brushes and suction openings (vacuum robots), tools fitted to grippers, welding electrodes, sharp edges on the chassis.
- **Stability under emergency stop** — can the system stand still on its own? Can it fall in a guided way (Safe Guided Falling)? Can stopping itself destabilise the platform — the inverted-pendulum problem familiar from humanoids?

Humanoid robots carry a special case rooted in the inverted-pendulum problem: stopping is not the same as being safe, unlike a classical industrial robot. The tier assessment reflects

this through stricter requirements on Axis 5 Pillar 1 (Safe Guided Falling, Safe Stability Control).

Transparency note on ISO/TS 15066. The biomechanical limit values originally specified in ISO/TS 15066:2016 are — as disclosed by the literature survey in the standard's own annex — based on data that has not been experimentally validated. Behrens et al. (Frontiers in Robotics and AI, 2022) provided the first experimentally validated set of limits in a study covering 112 subjects across 28 body regions, with an explicit gender effect. Further work (Svarný et al., 2020; Kirschner et al., 2021; Fischer et al. 2023) has shown that the ISO/TS 15066 force estimation formulas may both underestimate and overestimate actual impact forces. These findings support the methodological stance of HU-RSC: ISO/TS 15066 serves as an evidentiary basis, not as a conformity proof.

4.3 Axis 3 — Types and severity of harm

Axis 3 assesses which types of harm are plausible and how severe the harm would be if a safety event occurred. The structure follows ISO 12100 Table B.3:

- **Mechanical** — crushing, impact and collision injuries, lacerations from tools or sharp edges, falls of third parties triggered by the system.
- **Thermal** — burns from hot surfaces, frostbite from coolant, fire risk.
- **Chemical** — escaping cleaning agents, hydraulic fluid, battery electrolyte; rarely relevant in the use cases addressed here.
- **Electrical** — electric shock on a fault, sparking in explosive atmospheres; usually low for battery-powered systems.
- **Biological** — contamination from grippers in hospitals or care settings that have not been sterilised properly; use-case specific.
- **Psychological** — startle responses (with consequent falls), loss of trust in robotics generally, in rare cases trauma; relevant where vulnerable people are present. For prolonged co-presence with humanoid robots and AGVs in workplace settings, structural psychological strain must additionally be assessed in line with DIN EN ISO 10075-1 (Ergonomic principles related to mental workload). This is operator-side, not supplier-side, but it interacts with HU-RSC tier in the R column from Tier 3 upwards.

Empirical anchor for human-robot interaction. Morandini et al. (2025) synthesise 23 studies in a preregistered scoping review on cobots that adapt their behaviour to operators' psychological states (workload, attention, situational awareness). This marks a bridge between the AI behaviour axis and the HRI / mental workload axis: dynamic behavioural adaptation based on physiological markers (heart rate variability, electrodermal activity, gaze fixation) is methodologically not merely conceivable but empirically grounded.

A secondary dimension of Axis 3 is reversibility: fully reversible, reversible over time, irreversible. The ISO 12100 scale distinguishes four severity classes (slight, moderate, serious, very serious or fatal), which map onto HU-RSC tiers 1 to 5.

4.4 Axis 4 — Degree of autonomy

Axis 4 assesses the extent to which the system makes safety-relevant decisions of its own. The classification follows a five-step autonomy scale, drawn from ISO 25785-1 (Working Draft) and from established robotics autonomy taxonomies:

- **A0** — teleoperated or pre-programmed; no autonomous action; the operator carries every safety-relevant decision.
- **A1** — assisted; the system supports a human operator with sensor and path-planning data; the operator decides.
- **A2** — partly autonomous; the system handles routine decisions on its own (obstacle avoidance, route planning), but escalates safety-relevant decisions.
- **A3** — highly autonomous with oversight; the system also takes safety-relevant decisions on its own; a person can intervene at any time and retains ultimate control.
- **A4** — fully autonomous, without human ultimate control at the moment of operation; the system decides on its own, including in safety situations.

Axis 4 measures the degree of autonomy as a demand on the safety architecture, not the architecture's response to that demand. How well the safety architecture meets the demand is captured in Axis 5; the two together produce the picture.

A secondary factor is how often and how long the system makes safety-relevant autonomous decisions. An A3 system that takes three such decisions per day calls for a different judgement than one that takes three thousand.

4.5 Axis 5 — Safety architecture (three pillars)

Axis 5 is the only three-pillar axis in the model. It assesses how well the safety architecture meets the demands set by Axis 4 (degree of autonomy) and by the relevant ISO standards. Three pillars are distinguished:

4.5.1 Pillar 1 — Classical safety architecture

Pillar 1 covers classical functional safety to ISO 13849 and IEC 61508. The relevant factors are:

- **Certified safety functions** — the presence and class of certified safety functions (emergency stop, Safe Torque Off, Safe Stop, Safe Limited Speed, Safe Brake Control). Performance Level to ISO 13849-1.
- **Redundancy and diversity** — the architecture of the safety functions (Categories B, 1, 2, 3, 4 to ISO 13849-1); tolerance to hardware and software faults.
- **Safe-state concept** — does a universally safe state (stop) exist? If not — as for humanoids and AGVs in tight turns — are Active Damping (e.g. through Active Short Circuit at the joint level), Safe Guided Falling, and Safe Stability Control in place?

Empirical evidence for dynamic safety thresholds. Rustler et al. (IEEE Humanoids 2024) show in an empirical study on a UR10e covered with electronic safety skin (AIRSKIN) that thresholds set dynamically — based on per-link velocity and effective mass — yield a measurable productivity gain over static uniform settings without sacrificing safety. For HU-RSC this is methodologically relevant: static PFL conformity and productive human-robot collaboration are not mutually exclusive.

4.5.2 Pillar 2 — AI safety architecture

Pillar 2 covers the safety architecture of the AI components specifically. The relevant factors are:

- **AI Act conformity** — classification as high-risk under the AI Act 2024/1689 — where the component acts as a safety component within the meaning of the Machinery Regulation, the high-risk obligations apply (risk management, data quality, transparency, human oversight).
- **Explainability** — explainability of the decision logic (determinism, reproducibility, confidence values, post-hoc explanations).
- **Behavioural Safety** — a risk assessment of the AI components in line with ISO/IEC TR 5469; Behavioural Safety, Safe Situational Motion, Safe Behaviour Supervision (proprietary feature definitions in current industry practice, not yet covered by harmonised standards).
- **MR Annex III** — conformity with Annex III of the Machinery Regulation — safety components with self-evolving behaviour.

Note on terminology: Behavioural Safety refers to safety measures that work neither by classical-mechanical means (emergency stop, fenced enclosure) nor by classical-electronic means (certified safety functions to ISO 13849), but through AI-driven behaviour adjustment by the robot — for example slowing down when a person is nearby, anticipatory motion adjustment, or a behaviour-supervision layer that checks AI decisions against expected patterns. Terms such as Safe Situational Motion and Safe Behaviour Supervision are proprietary feature definitions used by individual manufacturers (for instance Synapticon and Neura) and are not yet covered by harmonised standards. ISO/IEC TR 5469 provides a first framework; further standards are in preparation. Synapticon (2026) defines AI Behavioral Safety as 'utilises multimodal foundation models, pre-trained on internet data and post-trained for robot-specific contexts; it semantically understands danger beyond geometry'. This is a vendor-specific working definition and not yet a harmonised standard term.

Empirical anchor for AI behavioural safety. SafeVLA (Zhang et al., 2025) demonstrates an empirically validated method for safety alignment of Vision-Language-Action models via constrained reinforcement learning: an 83.58 % reduction in cumulative safety violations alongside a 3.85 % gain in task success rate over the previous state of the art. For HU-RSC assessment this matters because a concrete aligner mechanism now exists — the question of verifiably safety-constrained VLA control is no longer a purely theoretical one.

4.5.3 Pillar 3 — Lifecycle update infrastructure

Pillar 3 assesses the technical update capability of the system as a property of the system itself — not the empirical track record of the supplier in delivering updates (that is the secondary dimension R1/R2/R3). The relevant factors are:

- **Patchability** — is the software architecture of the system update-capable? Modular components? Clear API boundaries? Documented update procedures?
- **Secure OTA infrastructure** — cryptographic signing, rollback mechanism, atomic updates, secure boot chain, protected update backend.
- **Substantial-modification detection** — does the system itself detect, on receiving an update, whether the update amounts to a substantial modification within the

meaning of the Machinery Regulation 2023/1230? Is the impact of updates on safety functions logged?

- **CRA conformity** — conformity with the Cyber Resilience Act 2024/2847 (security-update obligations, reporting duties, risk assessment).

Important distinction:

Pillar 3 measures technical update capability as a property of the system ('is the system patchable?'). The empirical likelihood that the supplier will actually deliver updates ('will the supplier ship the patches?') is a property of the supplier and is assessed in the secondary dimension R1/R2/R3 (Chapter 7). A patchable system is of no use without a supplier that ships updates — and the other way round. Both dimensions are necessary.

4.6 Axis 6 — Exposed persons and exposure frequency

Axis 6 assesses who is exposed to the system and how often. The structure follows the ISO 12100 logic (number of persons exposed, frequency of exposure):

- **Adult employees** — trained and briefed staff, no vulnerable groups present.
- **Adult householders** — private adults, untrained, with rights of householder — typical in the consumer setting.
- **Mixed audiences** — trained clinical staff together with untrained patients and visitors — typical in healthcare.
- **Vulnerable groups** — children, older people with reduced mobility or attention, patients with cognitive impairments, persons with physical impairments.
- **General public** — third parties without prior briefing, including children and older people — typical in public outdoor space.

A secondary dimension is the frequency and duration of exposure. A short, one-off exposure (a robot passes through a corridor and someone walks past) is weighted differently from a prolonged exposure (a care robot in a patient's room over several hours).

Procurement note on anthropometric heterogeneity. Fischer et al. (2024) demonstrate in a 31-subject study that operator anthropometry significantly influences which body region is impacted under an identical robot configuration — chest, abdomen, or pelvis may be affected by the same trajectory depending on stature. For procurement this implies that biomechanical verification against a single standard anthropometric model is insufficient when the workforce is anthropometrically heterogeneous. This finding feeds into HU-RSC tier assessment via the affected-persons axis.

5 Tier scale 1–5 with thresholds

The HU-RSC tier scale of 1 to 5 measures residual risk after protective measures. The tier descriptions below are stated in semi-quantitative form: they cite reference values from the relevant standards without claiming engineering precision. The legally binding judgement remains with the ISO 12100 risk assessment.

A note on how to read the axis values in the tier descriptions that follow. The values cited per tier are typical configurations in which a robotics system would fall on that tier. They are not necessary conditions per axis. Under the weakest-link principle, a single axis can set the overall tier even when the others sit lower. A humanoid robot with Tier-5 stand and fall potential on Axis 2 reaches Tier 5 overall, even if Axis 4 (degree of autonomy) only reaches A2.

Note on Axis 5 aggregation:

Aggregation within the three-pillar Axis 5 follows a modified minimum-floor rule. Pillar 3 (lifecycle update infrastructure) only enters axis aggregation from Tier 3 upwards, because the CRA reach and the weight of the update obligation are marginal at Tier 1 and Tier 2. Pillars 1 and 2 are aggregated across all tiers under the weakest-link principle. The detailed rationale is set out in Chapter 8.

5.1 Tier 1 — minimal residual risk

Character: stationary or very light mobile systems in a fully controlled environment; kinetically near-harmless; no safety-critical AI functions.

- **Axis 2** — $E_{kin} < 1$ J (oriented at the ISO/TS 15066 collision thresholds for soft tissue) or fully stationary.
- **Axis 3** — severity slight under ISO 12100 Table B.3, fully reversible.
- **Axis 4** — degree of autonomy A0–A1 (teleoperated or assisted).
- **Axis 5** — Pillars 1–2: ISO 13849 PL a or basic protective measures sufficient.

Examples: a desk-top robot arm used for edutainment in a trained setting; simple vacuum robots without rotating brush rollers or sharp tools; stationary inspection arms with limited payload in a controlled industrial cell. Note: Pillar 3 (lifecycle update infrastructure) is not aggregation-relevant at Tier 1 under the minimum-floor rule (Chapter 8.2). However, Cyber Resilience Act conformity remains binding regardless of HU-RSC tier — see the Pillar 3 R column in the Axis 5 expectations matrix.

5.2 Tier 2 — limited residual risk

Character: light mobile systems in a semi-controlled environment; reversible harm is plausible, irreversible harm is unlikely.

- **Axis 2** — E_{kin} between 1 and 10 J; low speed (typically below 0.5 m/s); no cutting or rotating tools.
- **Axis 3** — severity slight to moderate, reversible over time.
- **Axis 4** — degree of autonomy A1–A2.

- **Axis 5** — Pillar 1: PL b, simple emergency stop; Pillar 2: simple AI plausibility check; Pillar 3: not yet aggregation-relevant.

Examples: a vacuum robot in a private home; simple AGVs in an industrial setting with safety distances; inspection robots in commercial premises. Note: Pillar 3 (lifecycle update infrastructure) is not aggregation-relevant at Tier 2 under the minimum-floor rule (Chapter 8.2). CRA conformity remains binding regardless of HU-RSC tier.

5.3 Tier 3 — elevated residual risk

Character: mobile systems with material kinetic energy or with cutting or rotating tools; human contact is possible in a semi-controlled environment.

- **Axis 2** — E_{kin} between 10 and 80 J, OR tool motion that could cause injury (rotating blades, brushes).
- **Axis 3** — severity moderate to serious; partially irreversible (laceration, crush injury with sequelae).
- **Axis 4** — degree of autonomy A2–A3.
- **Axis 5** — typical expectations: Pillar 1 PL c or higher; Pillar 2 ISO/IEC TR 5469-aligned AI risk assessment; Pillar 3 enters aggregation here, with patchability and secure OTA infrastructure.

Examples: lawn-mowers with rotating blades in a private garden; cleaning robots in hotel corridors outside guest hours; AGVs in mixed human-robot environments.

5.4 Tier 4 — high residual risk

Character: systems with material injury potential in a semi-controlled or public environment; vulnerable groups are plausibly exposed; AI-driven decisions carry safety relevance.

- **Axis 2** — typically E_{kin} between 80 and 200 J as a guide; alternatively, humanoid form factor with fall potential; reach and lift height matter.
- **Axis 3** — severity typically serious; irreversible harm is plausible.
- **Axis 4** — typically degree of autonomy A3 with a high frequency of safety-relevant autonomous decisions.
- **Axis 5** — typical expectations: Pillar 1 PL d, redundant architecture; Pillar 2 documented and risk-assessed AI Behavioural Safety functions (third-party assessment where available), AI Act high-risk obligations met; Pillar 3 CRA-aligned update infrastructure with substantial-modification detection.
- **Axis 6** — vulnerable groups typically exposed plausibly or regularly.

Examples: humanoid pilot deployments in SME industry; cleaning robots in hospitals outside sensitive zones; outdoor lawn-mowers in publicly accessible parks.

5.5 Tier 5 — critical residual risk

Character: Tier 5 is not a quantitative step beyond Tier 4 but a qualitative one. Three configurations define Tier 5 — any one of them is enough: cumulative harm types in combination, severity reaching fatality, or direct and prolonged exposure of vulnerable groups

without human ultimate control. The Tier 4 / Tier 5 boundary for humanoid systems is not crisp: Tier 4 applies to humanoid systems with fall potential operating in structured environments — defined work area, geofencing, regular surface checks, supervised co-presence with trained adults. Tier 5 applies where the environment is unstructured (variable surfaces, free third-party access, no geofencing) or where multiple Tier-5 conditions accumulate (compound hazards, vulnerable groups exposed continuously). The buyer documents the rationale for placement in the procurement file.

- **Axis 2** — typically very high E_{kin} (above the Tier-4 reference threshold) or compound harm (kinetic with thermal, kinetic with chemical); humanoid form factor in unstructured environments. The joule threshold is heuristic — engineering judgement is made through the body-region-specific force and pressure thresholds in ISO/TS 15066:2016 Annex A (Tables 2 and 3), now also reproduced in ISO 10218-2:2025, not through E_{kin} alone.
- **Axis 3** — severity typically very serious to fatal; only irreversible.
- **Axis 4** — typically degree of autonomy A4 (fully autonomous without human ultimate control); alternatively A3 with so high a frequency of safety-relevant decisions that human oversight is in practice unsustainable.
- **Axis 5** — typical expectations: Pillar 1 PL e, Cat-4 redundant architecture; Pillar 2 fully aligned with the AI Act high-risk regime, MR Annex III met, Behavioural Safety documented, risk-assessed, and independently reviewed where available; Pillar 3 CRA conformity with emergency-update routes.
- **Axis 6** — vulnerable groups typically directly exposed (care, child supervision).

Examples: care assistance robots with bodily contact in care homes; fully autonomous healthcare robotics making safety decisions on patients.

Methodological outlook: the Behrens S0 framework. Behrens et al. (IEEE ISR 2021) propose a refinement of the PFL framework that introduces a new severity 'S0' with three sub-levels of tolerable stress at different risk levels — an approach methodologically close to the HU-RSC tier logic that may serve as a reference for tier refinement in a later version.

6 Expectations matrices (E-Min, E-Plus, R)

This chapter is the heart of the model. For each axis and each tier, three expectation sets are kept, parallel to the HU-RAC structure:

- **E-Min** — the minimum measures the product or supplier must demonstrably meet for the tier to be assigned.
- **E-Plus** — additional measures that secure the tier robustly (advanced state of the art, preparation for substantial modifications).
- **R** — residual risks that lie outside the tier and must be addressed in addition by the buyer or operator.

Function and use of the R column

The R column is neither a disclaimer nor a gap in the model. It is a duty-of-care warning: it makes transparent to the buyer which residual risks the tier structurally does not cover, and which additional organisational, technical, or contractual measures are needed.

This clarification carries liability implications. A buyer who reads the R column as *'risks that are acceptable as they stand'* rather than as *'risks that require additional measures'* will make poor procurement decisions.

Cross-references rather than repetition

Several safety measures appear logically across multiple axes — emergency stop, for instance, sits in Axis 5 Pillar 1 and also in Axis 3 (mechanical harm prevention). To avoid redundancy and inconsistency, the model holds each measure primarily on one axis and cross-refers from the others. Where this happens, the text uses 'see Axis X'.

6.1 Expectations matrix Axis 1 — Deployment environment

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
Tier 1	<ul style="list-style-type: none"> • fully controlled environment • reliable lighting, defined paths • no human access during operation 	<ul style="list-style-type: none"> • sensor-based environmental monitoring • logged access control • redundant lighting 	<ul style="list-style-type: none"> • any change of environment class requires reclassification • third-party access outside specification triggers reclassification
Tier 2	<ul style="list-style-type: none"> • semi-controlled environment with organisational protective measures • signage and staff briefing • defined exclusion times or zones kept clear of people 	<ul style="list-style-type: none"> • sensor-based person detection at low sensitivity • clear separation between robot operation and human access 	<ul style="list-style-type: none"> • unannounced third-party access calls for tighter organisational measures • pet encounters not covered
Tier 3	<ul style="list-style-type: none"> • semi-controlled environment; clear conduct rules for staff • audible or visual warning on 	<ul style="list-style-type: none"> • documented risk assessment of the deployment environment • barriers during sensitive 	<ul style="list-style-type: none"> • encounters with children or animals in semi-controlled settings are residual risks to be addressed organisationally (exclusion

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
	approach <ul style="list-style-type: none"> adequate lighting at all operating hours 	operations <ul style="list-style-type: none"> lighting-condition sensors 	times, supervision) <ul style="list-style-type: none"> varying floor conditions may become stability-critical
Tier 4	<ul style="list-style-type: none"> documented risk assessment of the deployment environment geofencing or defined operating zone clear escalation routes when the environment changes 	<ul style="list-style-type: none"> AI-supported environmental analysis with confidence values pre-mapped environment with live updates weather and lighting sensors with adaptive behaviour 	<ul style="list-style-type: none"> unexpected environmental changes caused by third parties (obstacles, other robots) require human intervention to be available public areas typically lie outside the expected operating envelope
Tier 5	<ul style="list-style-type: none"> full risk assessment for the most demanding environmental conditions redundant environment perception with independent sensor paths emergency hand-over protocol to human oversight when conditions are unforeseen 	<ul style="list-style-type: none"> multi-channel sensory environment modelling geofencing with hard-stop zones adaptation to weather, light, and acoustics with documented validation 	<ul style="list-style-type: none"> untrained environmental conditions are a residual risk to be addressed through operational supervision and emergency plans outdoor weather events (snow, ice, storm) remain a structural residual risk

6.2 Expectations matrix Axis 2 — Form factor and kinetic injury potential

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
Tier 1	<ul style="list-style-type: none"> $E_{kin} < 1$ J or fully stationary no cutting or rotating tools no fall potential 	<ul style="list-style-type: none"> redundant mechanical energy limiting soft materials at contact surfaces 	<ul style="list-style-type: none"> any hardware modification (attachments, tool change) should trigger a reclassification check
Tier 2	<ul style="list-style-type: none"> E_{kin} between 1 and 10 J speed typically below 0.5 m/s rounded edges, soft bumpers 	<ul style="list-style-type: none"> redundant mechanical speed limiting deformable outer shell defined motion-envelope limits 	<ul style="list-style-type: none"> fall consequences for third parties (trip incidents) remain a residual risk damage to property at contact surfaces is outside HU-RSC scope (see Requirements Catalogue Part 5)
Tier 3	<ul style="list-style-type: none"> E_{kin} between 10 and 80 J, OR cutting or rotating tools ISO/TS 15066-aligned collision avoidance tool emergency stop on human contact 	<ul style="list-style-type: none"> sensor-based tool detection with stop-on-touch dual mechanical safeguard (e.g. blade-lift plus stop) speed adapted to human presence 	<ul style="list-style-type: none"> with children or persons of reduced mobility, minimum detection time is critical — see Axis 6 damage to property remains a residual risk (see Requirements Catalogue)

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
			Part 5)
Tier 4	<ul style="list-style-type: none"> • E_{kin} between 80 and 200 J as a guide; alternatively humanoid form factor • comprehensive environmental sensing • Safe Limited Speed and Safe Limited Position as standard 	<ul style="list-style-type: none"> • multi-modal sensing (vision, LiDAR, tactile) • Active Damping on imminent fall (e.g. Synapticon ASC or equivalent) • model-based collision prediction • mechanically backdrivable joint design as passive impact protection 	<ul style="list-style-type: none"> • for humanoids: the inverted-pendulum problem means stopping is not the same as being safe — see Axis 5 Pillar 1 • fall consequences for third parties or vulnerable persons remain the primary residual risk
Tier 5	<ul style="list-style-type: none"> • E_{kin} above the Tier-4 guide (above 200 J) or compound hazards (kinetic with thermal, etc.) • redundant environmental sensing with diversity • Safe Guided Falling and Safe Stability Control for humanoid form factors 	<ul style="list-style-type: none"> • multi-channel motion and environment modelling • anticipatory fall avoidance • complete mechanical energy limiting in every operating state 	<ul style="list-style-type: none"> • under compound hazards, emergency plans and first-aid provision are an operator duty • complete collision protection against unexpected third parties remains structurally uncertain • the joule threshold is heuristic — engineering judgement is made through the body-region-specific thresholds in ISO/TS 15066:2016 Annex A (Tables 2 and 3), now also reproduced in ISO 10218-2:2025

6.3 Expectations matrix Axis 3 — Types and severity of harm

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
Tier 1	<ul style="list-style-type: none"> • harm type exclusively mechanical and slight • fully reversible severity under ISO 12100 Table B.3 	<ul style="list-style-type: none"> • documented harm-type analysis as an annex to the risk assessment 	<ul style="list-style-type: none"> • psychological consequences (startle responses in third parties) sit outside model scope at this tier
Tier 2	<ul style="list-style-type: none"> • harm types mechanical, occasionally electrical (electric shock on a fault) • severity slight to moderate, reversible over time 	<ul style="list-style-type: none"> • first-aid plan for foreseeable incidents • redundant electrical insulation 	<ul style="list-style-type: none"> • with untrained persons or children, even slight harm can have disproportionate consequences
Tier 3	<ul style="list-style-type: none"> • mechanical and thermal harm types possible • severity moderate to serious 	<ul style="list-style-type: none"> • documented harm-type matrix with reversibility analysis 	<ul style="list-style-type: none"> • long-term consequences of moderate-to-serious harm (scarring, restricted mobility)

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
	<ul style="list-style-type: none"> partially irreversible consequences plausible (laceration) 	<ul style="list-style-type: none"> protective measures per harm type first-aid resources on site 	<ul style="list-style-type: none"> are residual risks for the operator to cover
Tier 4	<ul style="list-style-type: none"> harm types mechanical, thermal, occasionally biological (in healthcare) severity serious; irreversible consequences plausible 	<ul style="list-style-type: none"> comprehensive harm-type matrix with probability of occurrence emergency plans with harm-type-specific response documented insurance cover 	<ul style="list-style-type: none"> irreversible harm despite protective measures remains the primary residual risk biological contamination in hospitals requires its own hygiene protocols
Tier 5	<ul style="list-style-type: none"> compound harm types plausible (mechanical with thermal or biological) severity very serious to fatal only irreversible consequences 	<ul style="list-style-type: none"> formal risk assessment with certified expert review emergency escalation chain to medical services comprehensive insurance including cyber cover 	<ul style="list-style-type: none"> fatal consequences despite protective measures are a residual risk to be mitigated through organisational measures (duty of supervision, emergency presence) compound hazards mean cumulative residual risks that have to be assessed individually

6.4 Expectations matrix Axis 4 — Degree of autonomy

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
Tier 1	<ul style="list-style-type: none"> degree of autonomy A0 (teleoperated) or A1 (assisted) operator carries every safety-relevant decision clear operator interface with emergency stop 	<ul style="list-style-type: none"> trained operator with certification remote-control link with latency guarantees 	<ul style="list-style-type: none"> loss of remote-control link: safe-state behaviour required (see Axis 5)
Tier 2	<ul style="list-style-type: none"> degree of autonomy A1 or A2 routine decisions autonomous, safety decisions escalated clear escalation interface to the operator 	<ul style="list-style-type: none"> logged escalation history operator response-time requirements 	<ul style="list-style-type: none"> where safety-relevant escalations are rare, operator response time becomes critical — the tier may be breached if operator presence is not guaranteed
Tier 3	<ul style="list-style-type: none"> degree of autonomy A2 to A3 defined catalogue of autonomously taken safety decisions human oversight available 	<ul style="list-style-type: none"> logging of all safety-relevant autonomous decisions response-time guarantees intervention interface kept simple 	<ul style="list-style-type: none"> unexpected decision situations outside the catalogue may produce unpredictable behaviour cognitive overload of oversight staff when

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
	at any time		escalations are frequent
Tier 4	<ul style="list-style-type: none"> degree of autonomy A3 with high frequency of autonomous safety decisions continuous human oversight plausible intervention available without delay 	<ul style="list-style-type: none"> oversight dashboard with confidence visualisation automated warning levels training programme for oversight staff 	<ul style="list-style-type: none"> with high decision frequency, sustained oversight is in practice unsustainable — the residual risk is human attention fatigue
Tier 5	<ul style="list-style-type: none"> degree of autonomy A4 (fully autonomous without human ultimate control at the moment of operation) comprehensive risk assessment of the autonomous decision logic clear take-over interface for emergencies 	<ul style="list-style-type: none"> tiered escalation to human ultimate responsibility auditable decision history regular training and validation runs 	<ul style="list-style-type: none"> A4 systems in human contact are highly contested, both regulatorily and ethically — buyers should approach deployment with particular caution and may wish to avoid Tier-5 deployment in vulnerable contexts MR Annex III and AI Act high-risk obligations apply in full

6.5 Expectations matrix Axis 5 — Safety architecture (three pillars)

The expectations matrix for Axis 5 is structured around three pillars: for each tier, Pillar 1 (classical safety architecture), Pillar 2 (AI safety architecture), and Pillar 3 (lifecycle update infrastructure) are assessed separately. The aggregation of the three pillars to the axis tier follows the minimum-floor rule — Pillar 3 only becomes aggregation-relevant from Tier 3 (see Chapter 8).

T	Pillar	E-Min	E-Plus	R
T1	1 — Classical safety	<ul style="list-style-type: none"> ISO 13849 PL a or basic protective measures emergency stop in place 	<ul style="list-style-type: none"> redundant emergency stop 	<ul style="list-style-type: none"> a hardware fault before an update may force a standstill
T1	2 — AI safety	<ul style="list-style-type: none"> AI without safety function, or no AI present 	<ul style="list-style-type: none"> where AI is present: simple plausibility check 	<ul style="list-style-type: none"> —
T1	3 — Lifecycle	<ul style="list-style-type: none"> not aggregation-relevant (see Chapter 8.2) 	<ul style="list-style-type: none"> — 	<ul style="list-style-type: none"> Pillar 3 is not aggregation-relevant at this tier (Chapter 8.2); CRA conformity remains binding regardless of tier
T2	1 — Classical	<ul style="list-style-type: none"> PL b 	<ul style="list-style-type: none"> two-channel 	<ul style="list-style-type: none"> single-point failures in

T	Pillar	E-Min	E-Plus	R
	safety	<ul style="list-style-type: none"> simple emergency stop with defined response time 	<ul style="list-style-type: none"> emergency stop self-diagnosis of safety functions 	<ul style="list-style-type: none"> non-redundant paths remain a residual risk
T2	2 — AI safety	<ul style="list-style-type: none"> where AI is present: simple plausibility check of decisions 	<ul style="list-style-type: none"> documented training and test data 	<ul style="list-style-type: none"> AI decisions outside the training distribution (distribution shift) are not covered
T2	3 — Lifecycle	<ul style="list-style-type: none"> not aggregation-relevant (see Chapter 8.2) 	<ul style="list-style-type: none"> where in place: regular updates communicated 	<ul style="list-style-type: none"> Pillar 3 is not aggregation-relevant at this tier (Chapter 8.2); CRA conformity remains binding regardless of tier
T3	1 — Classical safety	<ul style="list-style-type: none"> PL c or higher Safe Limited Speed, Safe Brake Control defined safe-state concept 	<ul style="list-style-type: none"> Cat-3 architecture to ISO 13849-1 Safe Stop 1/2 with monitored response 	<ul style="list-style-type: none"> with high-energy tools (rotating blades), minimum detection time is critical — see Axis 6
T3	2 — AI safety	<ul style="list-style-type: none"> safety-relevant AI decisions have plausibility checks ISO/IEC TR 5469-aligned risk assessment of AI components 	<ul style="list-style-type: none"> documented training and validation data confidence-based escalation 	<ul style="list-style-type: none"> non-deterministic AI behaviour in edge cases remains a residual risk explainability is limited
T3	3 — Lifecycle	<ul style="list-style-type: none"> aggregation-relevant from this tier upwards patchability ensured by architecture cryptographic signing of updates 	<ul style="list-style-type: none"> rollback mechanism atomic updates with safety-function test 	<ul style="list-style-type: none"> substantial-modification detection is still imprecise — interpretation evolving in practice
T4	1 — Classical safety	<ul style="list-style-type: none"> PL d, Cat-3 architecture redundant safety functions comprehensive safe-state concept 	<ul style="list-style-type: none"> Cat-4 architecture with diversity Active Damping and Safe Guided Falling for humanoid systems (e.g. via Active Short Circuit) 	<ul style="list-style-type: none"> for humanoid form factors, complete fall control is not fully achievable to the current state of the art classical Safe Torque Off (STO) does not produce a safe state for humanoid form factors — without active stability control, STO leads to an uncontrolled fall. STO must therefore be combined with Active Damping (e.g. via Active Short Circuit at the joint level) and Safe Guided Falling; the combination, not STO alone, is the safe-state mechanism battery-state-induced

T	Pillar	E-Min	E-Plus	R
				fall risk: low-battery conditions can produce uncontrolled fall if no battery-aware safe-state engagement (pre-emptive Safe Guided Falling on critical battery level) is implemented
T4	2 — AI safety	<ul style="list-style-type: none"> AI Act high-risk obligations met (risk management, data quality, transparency, oversight) MR Annex III applicable and met documented Behavioural Safety functions with manufacturer self-declaration; risk-assessed against ISO/IEC TR 5469 (and upcoming ISO/IEC TS 22440 once published); third-party assessment encouraged where available 	<ul style="list-style-type: none"> formal ISO/IEC TR 5469 risk assessment live monitoring with confidence values mechanism for detecting distribution shift 	<ul style="list-style-type: none"> AI behaviour in unfamiliar situations remains structurally uncertain explainability of VLAs is a current research topic
T4	3 — Lifecycle	<ul style="list-style-type: none"> CRA-aligned update infrastructure substantial-modification detection implemented minimum update path over expected lifetime (at least 5 years) 	<ul style="list-style-type: none"> redundant update paths emergency patches communicated within 72 hours 	<ul style="list-style-type: none"> update availability depends empirically on supplier longevity — check the secondary dimension R1/R2/R3 support-period gap risk: where the economic life of a robotics deployment exceeds the manufacturer's CRA support period (Article 13 minimum 5 years), the operator carries a residual cybersecurity-support gap. Contractual support-period extension should be considered as a procurement-side measure
T5	1 — Classical safety	<ul style="list-style-type: none"> PL e, Cat-4 architecture with diversity comprehensive safe-state concept redundant standstill paths 	<ul style="list-style-type: none"> multi-channel safety functions with independent power supply continuous self-diagnosis 	<ul style="list-style-type: none"> even PL e architectures carry a bounded probability of dangerous failure, not zero for humanoid form factors, classical Safe Torque Off (STO) is not a safe-state mechanism — STO must be combined

T	Pillar	E-Min	E-Plus	R
				with Active Damping and Safe Guided Falling to demonstrate a safe stop <ul style="list-style-type: none"> • battery-state-induced fall risk: low-battery conditions can produce uncontrolled fall if no battery-aware safe-state engagement (pre-emptive Safe Guided Falling on critical battery level) is implemented
T5	2 — AI safety	<ul style="list-style-type: none"> • full conformity with the AI Act high-risk regime • MR Annex III met • Behavioural Safety documented, risk-assessed, and independently reviewed where available • comprehensive explainability 	<ul style="list-style-type: none"> • formal verification of critical AI components, to the extent the state of the art allows • redundant AI models with cross-check 	<ul style="list-style-type: none"> • complete verification of neural networks is not achievable — a structural residual risk
T5	3 — Lifecycle	<ul style="list-style-type: none"> • CRA conformity with emergency update routes • substantial-modification detection documented and independently assessed where a recognised assessment route is available • regular audit updates documented 	<ul style="list-style-type: none"> • redundant OTA paths • pre-planned escalation on update failure • continuous compliance monitoring 	<ul style="list-style-type: none"> • supplier insolvency or market exit may put update delivery at risk — secondary dimension R1/R2/R3 is decisive • support-period gap risk: where the economic life of a robotics deployment exceeds the manufacturer's CRA support period (Article 13 minimum 5 years), the operator carries a residual cybersecurity-support gap. Contractual support-period extension should be considered as a procurement-side measure

6.6 Expectations matrix Axis 6 — Exposed persons and exposure frequency

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
Tier 1	<ul style="list-style-type: none"> • only trained staff exposed • individual persons only • exposure frequency: brief encounter, no prolonged co-presence 	<ul style="list-style-type: none"> • documented training history • attendance check 	<ul style="list-style-type: none"> • check for reclassification if the audience changes

Tier	E-Min — Minimum measures	E-Plus — Robustness uplift	R — Residual risks & measures
Tier 2	<ul style="list-style-type: none"> ▪ adults exposed ▪ rights of householder applicable for private persons ▪ no vulnerable groups ▪ exposure frequency: brief to moderate co-presence (minutes to a few hours) 	<ul style="list-style-type: none"> ▪ conduct rules documented ▪ access by children and animals excluded 	<ul style="list-style-type: none"> ▪ pet encounters (household dog, cat) cannot be ruled out in private settings — residual risk remains
Tier 3	<ul style="list-style-type: none"> ▪ mixed audiences possible ▪ children or older people excluded or protected organisationally ▪ exposure frequency: moderate co-presence, plannable and bounded 	<ul style="list-style-type: none"> ▪ exclusion times for sensitive groups ▪ audible or visual approach warning 	<ul style="list-style-type: none"> ▪ unannounced third-party presence cannot be fully ruled out — organisational duty of supervision
Tier 4	<ul style="list-style-type: none"> ▪ vulnerable groups plausibly or regularly exposed (children, older people, patients) ▪ exposure frequency: prolonged or frequent co-presence (several hours daily, or regularly recurring) ▪ documented supervision or protection plan 	<ul style="list-style-type: none"> ▪ adaptive behaviour by audience type (e.g. slower in areas with children) ▪ redundant person detection 	<ul style="list-style-type: none"> ▪ spontaneous reactions of vulnerable persons (startle, fall) cannot be fully prevented ▪ psychological consequences are a co-existing residual risk
Tier 5	<ul style="list-style-type: none"> ▪ vulnerable groups directly and continuously exposed (care, child supervision) ▪ exposure frequency: continuous co-presence or direct human contact over several hours daily ▪ comprehensive protection plan with supervising staff ▪ emergency presence guaranteed 	<ul style="list-style-type: none"> ▪ person-specific adaptation (e.g. patient profiles) ▪ continuous attendance monitoring of supervising staff 	<ul style="list-style-type: none"> ▪ continuous direct contact with vulnerable groups is regulatorily and ethically particularly sensitive — buyers should consider deployment only after explicit prior assessment by qualified experts

7 Secondary dimensions

The tier scale of 1 to 5 is the primary assessment dimension in HU-RSC. Two secondary dimensions complement it, parallel in structure to HU-RAC but reinterpreted in terms of safety.

7.1 Business criticality A/B/C

Business criticality measures the economic and reputational weight of a safety incident for the operator. Three classes guide the assessment:

- **A — critical** — a safety incident is existence-threatening (operational shutdown, serious personal injury with criminal proceedings, reputational loss leading to market exit). Typical in healthcare, regulated industry, and public-sector clients.
- **B — material** — a safety incident has material consequences (operational shutdown lasting weeks, significant damages claims, local reputational loss). Typical in SME industry.
- **C — limited** — a safety incident has limited consequences (short operational outage, insured damages claims). Typical in consumer settings and B2C services.

Business criticality is nominally independent of the HU-RSC tier, but in practice the two often correlate: healthcare applications (typically Tier 4 or 5) are mostly A-critical, and consumer applications (typically Tier 2 or 3) are mostly C-critical.

7.2 Supplier resilience R1/R2/R3

Supplier resilience assesses the likelihood that the technical safety capability demonstrated on Axes 1 to 6 will persist over typical robotics lifecycle horizons of five to ten years. Three classes guide the assessment:

- **R1 — high resilience** — an established supplier with a demonstrable market position, a documented record of update delivery, and clear lifecycle commitments. Creditworthiness verified (via Creditreform, CRIF, North Data, or equivalent). Emergency update path defined.
- **R2 — medium resilience** — a supplier with a working business model, but without a continuous lifecycle track record or with elevated credit risk. Update delivery is likely but not assured. Secondary sources (e.g. open-source components) may mitigate the risk.
- **R3 — low resilience** — a young company, a pre-revenue start-up, an unstable market position, or known credit problems. Update delivery over the expected lifetime is not plausibly assured. Vendor lock-in carries high residual risk.

Difference from HU-RAC:

In HU-RAC, R1/R2/R3 measures supplier resilience in the sense of data sovereignty (protection against acquisition, cloud lock-in, data ownership). In HU-RSC, R1/R2/R3 measures supplier stability in the sense of lifecycle persistence and compliance capacity (creditworthiness, market position, ability to deliver safety-relevant updates and spare parts over the lifetime). The two scales are structurally parallel but not identical in content

— buyers should not treat one as a copy of the other.

Note on vertical integration depth. The supplier's vertical integration depth is an additional resilience signal — both positive (lifecycle delivery capacity for in-house developed sub-systems) and negative (concentration risk, limited sub-system substitution). An industry overview of OEM strategy distribution is provided by FEV Consulting (Q3 2025): 100 % in-house full development at all 11 surveyed leading humanoid OEMs, but strongly differentiated outsourcing rates across vision sensors, compute, actuators, and battery system.

7.3 Relationship with Axis 5 Pillar 3

Axis 5 Pillar 3 (lifecycle update infrastructure) and the secondary dimension R1/R2/R3 (supplier resilience) are complementary. Axis 5 Pillar 3 measures the technical update capability as a property of the system — can the system in principle be patched safely? R1/R2/R3 measures the empirical likelihood that the supplier will actually deliver those updates.

A patchable system is of no use without a supplier that delivers updates. By the same token, a financially stable supplier is of no use if the system itself is not architecturally patchable. Both dimensions are necessary.

8 Aggregation and reclassification rules

8.1 Aggregation across the six axes

The tier of the system as a whole follows from the tiers of the six axes under the weakest-link principle: the worst-performing axis (the highest individual tier) sets the overall tier.

This rule is methodologically conservative and follows the ISO 12100 logic that the highest relevant risk drives the protective-measures requirement. It also protects buyers from the temptation to offset weak axis tiers against strong ones — which would be unsound from a safety standpoint.

Defence-in-depth as architectural principle. HU-RSC operationalises the principle that safety in human-robot interaction at Tier 3 and above depends on the layered combination of measures — environmental separation, sensor-based detection, software-level limits, and organisational procedures. No single measure is sufficient where vendor-certified safety functions are not yet available. The weakest-link aggregation rule above is the formalised expression of this principle: the buyer cannot offset weak measures on one axis with strong measures on another.

8.2 Aggregation within Axis 5 — the minimum-floor rule

The three-pillar Axis 5 is not aggregated by a plain application of the weakest-link principle. A modified minimum-floor rule applies instead:

- **Pillar 3 not aggregation-relevant for Tier 1/2** — for Tier 1 and Tier 2, Pillar 3 (lifecycle update infrastructure) does not enter the aggregation
- **Pillar 3 aggregation-relevant from Tier 3** — from Tier 3 upwards, Pillar 3 enters the aggregation in full
- **Pillars 1 and 2** — Pillars 1 and 2 are aggregated across all tiers under the weakest-link principle

The rationale is two-stranded. From the safety standpoint: at Tier 1 and Tier 2 — kinetically harmless or only mildly demanding applications without safety-critical AI — Pillars 1 and 2 (the initial safety architecture) sit on a low requirement level in any case, so a weak Pillar 3 would not improve the aggregated safety judgement, only shift it, while the underlying risk profile stays low. From the methodological standpoint: the minimum-floor rule prevents a weak lifecycle update infrastructure from pushing an otherwise harmless system onto an axis tier whose regulatory and operational weight it cannot carry.

Important clarification — no exemption from the CRA:

The minimum-floor rule expressly does not mean that Tier-1 and Tier-2 systems are exempt from Cyber Resilience Act obligations or from update-delivery duties. The CRA applies across product classes and is independent of the HU-RSC tier. At Tier 1 and Tier 2, Pillar 3 simply does not take effect within the HU-RSC tier aggregation; as a regulatory and operational requirement it remains relevant for any robotics system with digital elements. The expectations matrix for Axis 5 Pillar 3 carries this point in the R column for Tiers 1 and 2.

This minimum-floor rule will be reviewed in operation, with attention to its time-shifted effect. Pillar 3, unlike Pillars 1 and 2, does not act at the moment of deployment but over the system's lifetime. The aggregation rule may be refined in a later version if practice points to perverse incentives.

8.3 Reclassification on substantial modification

Robotics systems are dynamic. Software updates can change behaviour, tool changes can shift kinetic profiles, and changes in deployment context can shift Axis 1 or Axis 6. HU-RSC therefore calls for a reclassification check on any of the following triggers:

- **Software trigger** — a software update with an effect on safety functions (see Axis 5 Pillar 3 — substantial-modification detection).
- **AI trigger** — an AI model update where the AI is high-risk under the AI Act 2024/1689 (see MR Annex III — self-evolving systems).
- **Hardware trigger** — a tool or attachment modification that changes the form factor or kinetic injury potential (Axis 2).
- **Context trigger** — a change of deployment environment or audience (Axis 1 or Axis 6).
- **Regulatory trigger** — regulatory changes (e.g. new ISO standards taking effect, revised interpretation of the Machinery Regulation).

How the term 'substantial modification' under the Machinery Regulation 2023/1230 is to be interpreted is still an open question of detail. HU-RSC takes a conservative line and treats every trigger as a reason to reclassify; the legally binding judgement remains with the party placing the product on the market and with the competent authorities.

9 Minimum coupling HU-RSC × HU-RAC

9.1 The asymmetric minimum coupling

HU-RSC and HU-RAC are nominally independent models, but they carry a hard minimum coupling in one direction.

HU-RSC ≥ 4 requires HU-RAC ≥ 3

asymmetric — the reverse direction does not apply

9.2 Rationale for the coupling

Safety functions from HU-RSC Tier 4 upwards generate safety-relevant telemetry: emergency stops, collision events, AI confidence values, loss of stability. The Machinery Regulation 2023/1230 requires explicitly, in Annex III, that safety-relevant data and software are protected against corruption. Where this telemetry is not adequately integrity-protected, the safety argument for the Tier-4 system is broken — the safety function no longer delivers trustworthy evidence.

The concrete anchor for the minimum coupling HU-RAC ≥ 3 lies in the HU-RAC architecture itself. HU-RAC v0.3.1 Axis 5 (cybersecurity) introduces integrity protection of safety-relevant data and a secured update infrastructure as an E-Min requirement only at Tier 3 — Tiers 1 and 2 address data protection, authentication, and confidentiality, but not the integrity protection of telemetry and control data against corruption. HU-RSC Tier-4 systems, however, demand precisely that integrity protection (MR Annex III). HU-RAC 3 is therefore not an arbitrary choice; it is the lowest HU-RAC tier at which the Machinery Regulation's minimum requirement on data integrity is structurally addressed.

A high HU-RSC class without a sufficient HU-RAC class therefore cannot be operated at high integrity — the coupling is forced by regulation, not by aesthetic preference. Existing safety bus implementations — for example FSoE (Fail-Safe over EtherCAT) — provide Safe Process Data with integrity protection at the protocol level. From HU-RSC Tier 4 upwards, the safety telemetry path must demonstrably implement equivalent integrity protection; this is the technical anchor for the minimum coupling to HU-RAC Tier 3.

9.3 Asymmetry

The reverse direction does not apply: a high HU-RAC class does not force a high HU-RSC class. A data-sensitive but stationary, kinetically harmless robot is conceivable — for instance a telepresence system with video recording in a solicitor's office. The coupling is asymmetric because data sensitivity without kinetic risk is a realistic scenario, whereas kinetic risk without data sensitivity is not.

Empirical anchor for HU-RAC \leftrightarrow HU-RSC. AttackVLA (Li et al., 2025) shows in parallel that backdoor attacks on VLA models are feasible in both simulation and real-world robotic settings, with an average targeted success rate of 58.4 % and up to 100 % on specific tasks.

This makes the HU-RAC ↔ HU-RSC interface empirically tangible: a compromised AI controller translates immediately into a physical hazard. The cybersecurity axis (HU-RAC) is therefore not analytically cleanly separable from HU-RSC but is operationally entangled with it.

9.4 Operational consequences

For procurement practice, the minimum coupling means:

- **At Tier 4 or 5** — a robotics system with HU-RSC Tier 4 or 5 must come from a supplier whose data and update architecture is fit for HU-RAC Tier 3.
- **In SMEs outside healthcare** — HU-RAC ≥ 3 is recommended for Tier-4 and Tier-5 deployments, but may be relaxed under stricter operational supervision arrangements.
- **In healthcare** — HU-RAC ≥ 3 is non-negotiable — combining a Tier-5 system with a weak HU-RAC class creates a compliance risk that cannot reasonably be accepted.

10 Application in procurement practice

HU-RSC is designed as a procurement-stage assessment framework. This chapter sketches how the model is used in a procurement process and walks the reader through a worked example. Detailed operationalisation — questionnaires, scoring sheets, supplier-specific scorecards — is not held here, but in the internal Requirements Catalogue v2.x (Document 01) and the external Requirements Profile v2.x (Document 03).

10.1 The three-step logic

- **Step 1 — set the target** — the buyer identifies the level of safety the use case requires. What must the system at minimum deliver? This judgement is independent of the particular product. Output: target HU-RSC tier.
- **Step 2 — assess the product** — the buyer assesses the offered product against the six axes and the expectations matrices. Does it meet the E-Min requirements of the target tier? Output: actual HU-RSC tier of the product.
- **Step 3 — compare and decide** — if $\text{actual} \geq \text{target}$, procurement can proceed; the R column becomes a checklist of supplementary measures. If $\text{actual} < \text{target}$, look for an alternative product or adapt the use case.

10.2 Connection to HU-RAC and the Requirements Catalogue

A complete assessment uses both models together — HU-RAC for data and resilience, HU-RSC for safety. The minimum coupling (Chapter 9) sets the floor. Supplier-specific assessments (NEURA, Agile Robots, PAL, Hexagon, Yarbo, Nexaro and others) are held in the Requirements Catalogue v2.x, which is published separately and updated independently of this discussion paper.

Verification of conformity claims. Manufacturer claims of ‘PLe’, ‘SIL3’, ‘AI Act compliant’, ‘CRA ready’ or equivalent must be substantiated by a conformity certificate from a notified or competent body, identifying the standard version, the scope of the safety function, the certificate validity date, and the notified body’s identification number. Whitepaper or marketing-page claims alone do not constitute evidence for HU-RSC E-Min compliance at Tier 4 or higher. Buyers should request and archive certificate copies as part of the procurement file.

10.3 Worked example — humanoid robot in an SME workshop

To illustrate the three-step logic, this section walks through a fictional but realistic procurement case. A metalworking SME with thirty-five employees wants to pilot a humanoid robot for workpiece feeding and tool loading at a CNC machine. The robot is to move freely within a defined area of the workshop, carry workpieces from storage to the machine, and work unsupervised among trained employees.

Step 1 — set the target tier

The six axes are placed in the use-case context:

- **Axis 1 (deployment environment)** — semi-controlled workshop environment with employees present — Tier 3 as target.

- **Axis 2 (form factor)** — humanoid form factor with fall potential, kinetic energy in the 80–200 J range — Tier 4 as target.
- **Axis 3 (types of harm)** — mechanical severity plausibly serious, irreversible consequences possible — Tier 4 as target.
- **Axis 4 (degree of autonomy)** — partly autonomous with human oversight (A3), several safety-relevant autonomous decisions per hour — Tier 4 as target.
- **Axis 5 (safety architecture)** — the safety architecture must carry humanoid form factor and AI-driven decisions — Tier 4 as target.
- **Axis 6 (exposed persons)** — trained adult employees, no vulnerable groups, several hours of daily co-presence — Tier 3 as target.

Aggregated under the weakest-link principle: target HU-RSC tier = 4. The minimum coupling (Chapter 9) therefore also requires HU-RAC ≥ 3 .

Step 2 — assess the offered product

The supplier offers a humanoid robot with the following stated features: PL d on the critical safety functions; Active Damping (via Active Short Circuit) implemented but not third-party assessed; AI Act high-risk self-declaration in place but no external certification; documented update infrastructure with OTA but no explicit substantial-modification detection; the supplier is a pre-revenue start-up that has been on the market for eighteen months.

Assessment against the expectations matrices:

Axis-by-axis assessment of the offered product: Axis 1 (deployment environment) — SME workshop with employees, organisational protective measures, defined exclusion times — meets Tier-3 E-Min, considered against Tier-4 target through additional measures. Axis 2 (form factor / kinetic injury potential) — humanoid form factor, walking speed approximately 1 m/s, fall potential present — meets Tier-4 typical configuration but not Tier-5. Axis 3 (types of harm) — mechanical with plausibly serious severity, irreversible consequences possible — Tier 4 target met. Axis 4 (degree of autonomy) — partly autonomous (A3) with several safety-relevant autonomous decisions per hour — Tier 4 target met. Axis 6 (exposed persons) — trained adult employees, no vulnerable groups, several hours daily co-presence — Tier 3 in compliance with Tier 4 minimum operator-side measures. Axis 5 detailed assessment (three pillars):

- **Pillar 1** — Axis 5 Pillar 1 meets Tier-4 E-Min (PL d) — fits.
- **Pillar 2** — Axis 5 Pillar 2 does not fully meet Tier-4 E-Min — the AI Act high-risk obligations without external certification are borderline, and MR Annex III conformity is unclear.
- **Pillar 3** — Axis 5 Pillar 3 does not meet Tier-4 E-Min — the lack of substantial-modification detection is an E-Min gap at Tier 4.
- **Secondary dimension R class** — the supplier is a pre-revenue start-up without a documented lifecycle track record — R3 (low resilience).

HU-RAC parallel assessment. The supplier states that the backend is hosted within the EU on infrastructure operated by a European subsidiary of a US parent. Update signing is in place; tenant separation is documented in the supplier's technical and organisational measures. Under HU-RAC v0.3.1, this supplier reaches Tier 3 on the data sovereignty axis only conditionally — the US parent ownership triggers an Article 48 GDPR examination under the CLOUD Act. Without that examination, HU-

RAC Tier 3 is not demonstrable, and the minimum coupling (HU-RSC ≥ 4 requires HU-RAC ≥ 3) is not satisfied. This blocks Tier 4 even before the safety-side gaps are addressed.

Actual HU-RSC tier of the product: Tier 3. Axis 5 carries only Tier-3 conformity, because Pillars 2 and 3 do not meet the Tier-4 E-Min — the weakest pillar sets the axis tier, and the weakest axis tier sets the overall tier.

Step 3 — compare and decide

Target = Tier 4, actual = Tier 3. The product does not meet the target. Added to that is the secondary dimension R3, which is problematic for safety-relevant Tier-4 deployments because update delivery over the lifetime is not plausibly assured.

Procurement decisions then come in three shapes. First, decline the product and look for an alternative supplier that meets Pillar 2 and Pillar 3 at Tier 4. Second, adapt the use case to a Tier-3 level — for instance through additional separation of robot and employees by exclusion times or defined work phases. Third, contractual risk transfer aimed specifically at the R3 class: an update-delivery commitment over the planned five-to-seven-year lifetime, a creditworthiness guarantee or source-code escrow arrangement for the insolvency case, and documented spare-parts availability. The choice between these options is an operational risk-and-return judgement, not a model output — HU-RSC delivers the diagnosis, not the prescription.

Note: this example is illustrative and does not substitute for a concrete ISO 12100 risk assessment. Real procurement decisions need detailed examination by qualified experts and, where appropriate, legal advice.

11 Final remarks

HU-RSC is a model for procurement practice in German-speaking SMEs, not a universal classification system. It does not compete with established standards; it complements them with a procurement-side perspective that is under-represented in the engineering-centred world of standards. Comments, corrections, and practical experience from real applications are explicitly welcomed and will feed into future versions.

12 Annex — mapping tables for ISO classes

The table below maps HU-RSC tiers onto the relevant ISO safety classes. The mapping is a guide, not a conversion. ISO classes assess individual safety functions, whereas HU-RSC assesses systems as a whole — the two levels of granularity cannot be offset against one another. The table is included to give a buyer familiar with ISO standards a quick orientation.

HU-RSC tier	ISO 13849-1 (PL)	IEC 61508 (SIL)	ISO 26262 (ASIL)*	ISO 25785-1**
Tier 1	PL a typical	SIL 1 or no claim	QM or ASIL A (not directly applicable to robotics)	low class (provisional)
Tier 2	PL b typical	SIL 1	ASIL A–B (provisional guide)	low to medium class (provisional)
Tier 3	PL c typical	SIL 1–2	ASIL B (provisional guide)	medium class (provisional)
Tier 4	PL d typical	SIL 2–3	ASIL C–D (provisional guide)	high class (provisional)
Tier 5	PL e typical	SIL 3	ASIL D (provisional guide)	very high class (provisional)

* ASIL under ISO 26262 originates in the automotive sector and is not directly applicable to robotics. The column is included as a coarse risk-level guide.

** ISO 25785-1 is, as of April 2026, still a Working Draft (classification for dynamically stable mobile robots). When the final standard is published (expected in 2027), the mapping table will be updated in a subsequent version of this discussion paper. Until then the provisional guidance applies, with reservations.

12.1 Notes on using the mapping table

- **PL/SIL as indication, not as proof** — a supplier whose product is certified to PL d offers a strong indication of fitness for Tier 4 — but not a proof. The other axes (deployment environment, exposed persons, degree of autonomy, and so on) must be assessed separately. Reading the mapping table: ISO 13849-1 PL and IEC 61508 SIL assess individual safety functions, while HU-RSC assesses the system as a whole. Multiple safety functions with different PL or SIL ratings can coexist in one system; HU-RSC takes the overall residual risk as the tier driver, not the highest-rated safety function. The mapping is therefore an orientation guide, not a conversion.
- **PL and SIL are not synonyms** — ISO 13849-1 PL assesses safety-related parts of control systems; IEC 61508 SIL assesses safety integrity. Some components are rated in PL, others in SIL. Both are usable, but they are not the same thing. A note on the new ISO 10218-1:2025 class distinction: the 2025 edition introduces two classes

of industrial robots — Class 1 covers very low-power robots that do not pose a significant hazard and may be assessed with reduced control requirements (typically PL b); Class 2 covers the conventional industrial robot scope. For HU-RSC mapping, many Tier-1 and Tier-2 candidate systems will fall in Class 1; Tier-3 and above typically require Class 2 assessment.

- **Joule heuristic versus ISO/TS 15066 Annex A** — the joule thresholds used in Axis 2 are a procurement-stage heuristic. The engineering-precise assessment of human-robot collisions does not work through kinetic energy as a scalar but through the force and pressure thresholds for 29 specific body regions set out in ISO/TS 15066:2016 Annex A (Tables 2 and 3), now also reproduced in ISO 10218-2:2025. A joule-only assessment, with no reference to impact area or impact direction, is incomplete from the safety-engineering standpoint. In a procurement context, the joule heuristic remains useful as a quick first-cut categorisation; the ISO/TS 15066 assessment belongs with the ISO 12100 risk assessment. A further calibration limit: the biomechanical thresholds in ISO/TS 15066:2016 (now ISO 10218-2:2025) are calibrated to AIS 1 (Abbreviated Injury Scale, level 1 — minor injuries, bruising). Tier-4 and Tier-5 scenarios in HU-RSC, which describe AIS 3 to AIS 5 severities, lie outside the direct validity range of the biomechanical table.
- **Mapping updates** — when the final standards (ISO 25785-1 for humanoid and dynamically stable mobile robots, ISO/IEC TR 5469 for AI functional safety) are published, the mapping will be updated in a subsequent version of this paper. The model itself is not affected. A further known limitation: the transient and quasi-static contact distinction in ISO/TS 15066:2016 contains acknowledged interpretive ambiguities (Kirschner et al. 2022). The integration into ISO 10218-2:2025 partially clarifies this, but documented risk assessments should explicitly state which interpretation has been applied.

Empirical reference point. *Fraunhofer IPA KMU humanoid measurements (September 2025) on a Unitree G1 humanoid robot (35 kg, 1.32 m height, walking speed approximately 1 m/s, measured with a GTE CBSF 75 N/mm pressure sensor) reach approximately 530 N forward-walking contact force at the abdomen and 510 N at the chest. Both exceed the ISO 10218-2:2025 pain-onset thresholds (abdomen 220 N, hand and fingers 280 N) by roughly a factor of 2.4. A 35-kg humanoid at normal walking speed thus places systematically in HU-RSC Tier 4 on Axis 2 and triggers the minimum coupling to HU-RAC Tier 3 (Chapter 9). Source: Fraunhofer IPA, Leitfaden Sicherheit Humanoider Roboter, 25 October 2025.*

13 Annex — Glossary

This annex collects working definitions of technical terms used in this paper. Where a term originates from a specific vendor or technical specification, the source is noted. Definitions are intended for procurement orientation and do not substitute for the authoritative terms in the underlying standards or vendor documentation.

Active Damping. Vendor term (Synapticon, 2026): the controlled electromagnetic braking effect produced by Active Short Circuit at the joint level. Not a harmonised standard term. Used in HU-RSC as the actuator-level component of a safe-state mechanism for humanoid form factors.

Active Short Circuit (ASC). Vendor term (Synapticon, 2026): an actuator function that electronically short-circuits the motor windings upon fault detection, producing controlled electromagnetic braking. Building block for Safe Guided Falling.

AI Behavioural Safety. Vendor term (Synapticon, 2026): safety architecture using multimodal AI foundation models, pre-trained on internet data and post-trained for robot-specific contexts, to semantically understand danger beyond geometry. Not yet a harmonised standard term. UK spelling 'Behavioural' used throughout this paper; equivalent to US 'Behavioral Safety'.

Backdrivability. Mechanical property of an actuator whereby external force can move the output, dissipating impact energy through controlled compliance. Relevant for humanoid joints under fall or stumble scenarios.

Defence-in-depth. Architectural principle: safety at Tier 3 and above relies on the layered combination of measures (environmental separation, sensor-based detection, software limits, organisational procedures) rather than on any single measure.

FSoE (Fail-Safe over EtherCAT). Safety-rated industrial bus protocol providing Safe Process Data (Safe Position, Safe Velocity etc.) with protocol-level integrity protection. Example anchor for the minimum coupling HU-RSC × HU-RAC at Tier 4 and above.

Humanoid robot. A robotic system with anthropomorphic morphology (upper body and two arms; two legs or other suitable locomotion), with environment-sensing capability, intended to perform tasks designed for humans.

Humanoid robot system. The complete deployable unit comprising at least one humanoid robot, end-effectors or tools, workpieces, peripheral components, and charging infrastructure. Safety devices and measures are necessary system-level elements.

Inverted-pendulum problem. The dynamic instability of a humanoid robot when power is cut: classical Safe Torque Off does not produce a safe state but an uncontrolled fall. Mitigation requires Active Damping and Safe Guided Falling routines.

Power and Force Limiting (PFL). Collaborative operation mode under ISO 10218-2:2025 (formerly ISO/TS 15066) in which physical contact between robot and operator is permitted under controlled force and pressure thresholds.

Safe Guided Falling (SGF). Vendor term (Synapticon, 2026): a controlled collapse of a humanoid robot upon fault detection, using Active Damping rather than Safe Torque Off, to reduce impact energy and protect both hardware and nearby persons. Not yet a harmonised standard term.

Safe Limited Position. Certified safety function (ISO 13849-1) limiting movement to a defined position window.

Safe Limited Speed (SLS). Certified safety function (ISO 13849-1) limiting velocity below a defined threshold; monitored at the drive level.

Safe Limited Torque (SLT). Certified safety function limiting torque output below a defined threshold.

Safe Motion Monitoring (SMM). Family of certified safety functions monitoring motion parameters (position, velocity, acceleration) against defined limits.

Safe State. A state in which the safety hazard is reduced to an acceptable level. For most fixed-base robots this is a stop; for humanoid robots it requires combined Active Damping and Safe Guided Falling.

Safe Torque Off (STO). Certified safety function (ISO 13849-1, IEC 61800-5-2) that removes torque-producing energy from the motor. Sufficient as safe state for fixed-base robots; not sufficient for humanoid form factors (see Safe State).

Speed and Separation Monitoring (SSM). Collaborative operation mode under ISO 10218-2:2025: continuous monitoring of distance between robot and operator with speed reduction or stop on approach.

Substantial modification (MR 2023/1230). A modification of machinery already placed on the market that may trigger a fresh conformity assessment, including software updates that change safety-relevant or functional behaviour. The detailed interpretation is still under development by the Commission.

Transient vs. quasi-static contact. ISO/TS 15066:2016 (now ISO 10218-2:2025) distinguishes transient (unconstrained, recoiling) and quasi-static (constrained, clamping)

contact scenarios, with separate force and pressure thresholds. Definitions contain acknowledged interpretive ambiguities (Kirschner et al. 2022); risk assessments should declare which interpretation applies.

14 Annex — References

- Fraunhofer IPA, Leitfaden Sicherheit Humanoider Roboter, Stuttgart, 25 October 2025.
- Fraunhofer IPA, Leitfaden Wirtschaftlichkeit Humanoider Roboter, Stuttgart, 2025.
- Synapticon GmbH, Powering the Humanoid Revolution — High-Dynamic, Safety-Rated Actuation for Unstructured Environments (White Paper), 23 January 2026.
- Klindt, T. (Noerr), Humanoide Roboter: Diese Regeln gelten in Europa, ingenieur.de, 4 May 2026.
- Regulation (EU) 2023/1230 of the European Parliament and of the Council of 14 June 2023 on machinery (Machinery Regulation), OJ L 165, 29.6.2023, p. 1. Corrigendum: OJ L 169, 4.7.2023, p. 51 (corrected application date 20 January 2027).
- European Commission, Proposal for a Digital Omnibus on AI of 19 November 2025, together with Council and Parliament negotiating positions adopted in spring 2026 (status at time of writing: not yet formally adopted).
- Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L of 12.7.2024.
- Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements (Cyber Resilience Act), OJ L of 20.11.2024.
- Directive (EU) 2024/2853 of the European Parliament and of the Council of 23 October 2024 on liability for defective products (Product Liability Directive), OJ L of 18.11.2024.
- Directive 2014/53/EU on radio equipment (Radio Equipment Directive), OJ L 153, 22.5.2014, p. 62.
- Kirschner, R. J., Mansfeld, N., Abdolshah, S., et al., ISO/TS 15066: How Different Interpretations Affect Risk Assessment, arXiv:2203.02706, 2022.
- Fischer, C., Neuhold, M., Steiner, M., et al., Collision Tests in Human-Robot Collaboration, IEEE Access 11, 118395–118413, 2023.
- FEV Consulting GmbH, Humanoid robots — Potential or hype for automotive suppliers? (Insights), Q3 2025, 14 pages.
- Behrens, R., et al., A Statistical Model to Determine Biomechanical Limits for Physically Safe Interactions With Collaborative Robots, Frontiers in Robotics and AI, 2022.
- Behrens, R., et al., A Revised Framework for Managing the Complexity of Contact Hazards in Collaborative Robotics, 2021 IEEE International Conference on Intelligence and Safety for Robotics (ISR), 2021.
- Fischer, C., et al., Personalized Safety: Considering the Worker's Anthropometry in Safety Evaluation of Human-Robot Collaboration, 2024.
- Li, J., et al., AttackVLA: Benchmarking Adversarial and Backdoor Attacks on Vision-Language-Action Models, arXiv preprint, 2025.
- Morandini, S., et al., Collaborative Robots Adapting Their Behavior Based on Workers' Psychological States: A Systematic Scoping Review, Human Behavior and Emerging Technologies, 2025.
- Rustler, L., et al., Adaptive Electronic Skin Sensitivity for Safe Human-Robot Interaction, 2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids), 2024.
- Svarný, P., et al., 3D Collision-Force-Map for Safe Human-Robot Collaboration, 2021 IEEE International Conference on Robotics and Automation (ICRA), 2020.
- Zhang, B., et al., SafeVLA: Towards Safety Alignment of Vision-Language-Action Model via Safe Reinforcement Learning, arXiv preprint, 2025.