

IMPULSE PAPER

Technical discussion paper – not a draft standard

HU-RAC

A Conceptual Framework for Classifying Smart Robotic Systems by Data Sovereignty and Data Security

A contribution to the public debate on a possible future standardisation and regulatory framework

Questions of technical and functional safety are deliberately excluded and will be addressed separately in an announced second impulse paper.

Version 0.3.1

As of: April 2026

Author: Humanide GmbH

*This document is neither a draft standard nor a standard.
It does not constitute a binding requirement and does not create legal or performance obligations.*

It serves the sole purpose of technical discussion.

Legal Notice

This impulse paper is a non-binding technical paper intended for discussion purposes. It is neither a draft standard within the meaning of the DIN or ISO procedure, nor a coordinated consortium position, nor a standard, a technical rule or a recommendation within the meaning of any product liability or safety-case regime. The classifications, categories, matrices and formulations contained herein are intended exclusively as suggestions for further discussion among interested parties.

The author excludes any liability for decisions made on the basis of the representations contained in this document. Use of this document or parts of it in tenders, contracts, regulatory proceedings or legally binding contexts is not intended. Where such use occurs, it takes place at the sole responsibility of the using party.

Statements referring to the General Data Protection Regulation (GDPR), the AI Act, the Machinery Regulation, the Cyber Resilience Act and other legal acts are placed in professional context but have not been reviewed by counsel and do not replace individual legal advice.

The designation HU-RAC is used as a working term within this paper. The RARI data referenced in this paper stem from the keynote by N. Radford (Persona AI) at the Humanoids Summit 2025 (11 December 2025, Mountain View) and are cited here as a secondary source. Trademark and copyright rights in third-party designations and representations remain with their respective holders.

Use and Distribution

The distribution and citation of this impulse paper is expressly welcomed, provided that the authorship of Humanide GmbH is named and the character as a non-binding discussion document remains recognisable. Excerpts may be used if the source is attributed.

Executive Summary in Three Paragraphs

A lawn-mowing robot in the garden and a humanoid care companion in the bedroom of a person with dementia should not be evaluated by the same logic when it comes to data sovereignty and data security. Today they often are – sometimes implicitly, sometimes explicitly. This paper proposes classifying smart robotic systems along six context dimensions into five categories (HU-RAC 1 to HU-RAC 5) and assigning each category a graduated expectation profile on questions of data protection, key sovereignty, tenant isolation, AI reference and operational continuity. The approach follows the principle of the most stringent condition: the highest classification on a single axis determines the overall category.

HU-RAC is not intended to replace existing legal acts, but to act as an overlay: a domain-specific organising and translation grid that makes the GDPR, the AI Act, the Machinery Regulation, the Cyber Resilience Act and the relevant ISO standards applicable in context. The regulatory frameworks exist; what is missing is their coherent, robotics-specific operationalisation.

Three recent events show that the gap is not hypothetical. In February 2026, security researchers discovered an authorisation flaw in the backend of a leading manufacturer of autonomous household robots, allowing third-party access to thousands of devices. In February 2025, a market-dominant US technology group withdrew its zero-knowledge cloud protection feature for British users after the British security authorities issued a corresponding order. And in 2024, a European manufacturer of networked blood-pressure devices filed for insolvency. The shutdown of the cloud infrastructure took with it not only the apps without which the devices were inoperable, but also the users' historical health data. For humanoid service robots in sensitive environments, these are not footnotes – they touch the core of the trust question.

The framework integrates a second perspective: the commercial one. The RARI framework of Persona AI (Radford, Humanoids Summit 2025) rates industries by their adoption readiness and shows that precisely industries with complex regulatory environments (Healthcare, Home) exhibit the lowest commercial readiness. A clear classification system lowers compliance friction and raises adoption readiness exactly where the market is currently stalling. In this reading, standardisation becomes a lever for market opening.

One important caveat up front: this paper deliberately does not address questions of technical and functional safety. The emergency-stop problem in dynamically stable humanoid systems, power-loss behaviour, power-and-force-limiting thresholds in non-industrial contexts – all of that deserves its own impulse paper, which Humanide GmbH is preparing separately. Chapter 2 makes this demarcation explicit.

1 Why This Paper

Smart robotic systems are permeating the world of work and everyday life at a speed for which the existing regulatory and standards landscape was not built. The Machinery Regulation treats lawn-mowing robots and humanoid care companions largely under the same roof. The AI Act distinguishes high-risk categories but provides no fine-grained differentiation that would do justice to the heterogeneity of robotics. The EU Product Liability Directive remains equally coarse in its risk logic. The result is a twofold misalignment: consumer products tend to be over-regulated, while high-risk deployments tend to be under-regulated. Both are detrimental to the market.

Three events from the past 18 months have shown how concrete this gap becomes as soon as service robots enter sensitive areas.

1.1 Three Episodes Worth Thinking About

A Backend Incident at a Consumer Robotics Manufacturer, February 2026

In February 2026, security researchers discovered an authorisation flaw in the MQTT backend of a leading manufacturer of autonomous household robots. Approximately 6,700 devices were compromised. Attackers with valid user credentials could, by manipulating device references, access cameras, microphones and floor plans of apartments belonging to other users. Transport encryption was sound. The vulnerability lay not in cryptography but in authorisation logic: a multi-tenant backend with a shared pool and token-based assignment enabled cross-access. Lesson: Transport encryption does not protect against flaws in tenant isolation.

Withdrawal of a Zero-Knowledge Cloud Feature in the United Kingdom, February 2025

A market-dominant US technology group had since 2023 offered a zero-knowledge architecture for cloud storage services: the cryptographic keys lay with the user, and the provider could not technically access plaintext data. In February 2025, the group withdrew this feature for British users after British security authorities issued an order under the applicable legal framework for intelligence powers. Lesson: Only what the provider has never technically possessed can it, even under compulsion, not hand over to state authorities. Comparable legal bases exist in Germany (G10 Act, TKÜV) and France, and are being discussed at the EU level.

The Insolvency of a Cloud-Dependent Medical-Device Manufacturer, 2024

A European manufacturer of networked blood-pressure devices filed for insolvency in 2024, shut down its cloud infrastructure – and with it the apps without which the devices were inoperable, as well as the historical measurement data of its users, which were not available locally. It is a pattern that has repeated itself among several smart-home and connected-health providers. For a healthcare or care setting with a humanoid assistance system, such an outage would be existential. Lesson: Cloud dependency without architectural and contractual continuity protection creates a risk that is difficult to reconcile with the care and medical context.

1.2 The Guiding Question

How can one arrive at a graduated assessment that does not burden a pool-cleaning robot with the regulatory weight of a surgical-assistance system, and conversely does not grant a care companion the regulatory leniency afforded to a vacuum-cleaning robot? And do so in a way that does not appear arbitrary, but can be thought through along clear context dimensions that fit into the existing regulatory landscape? This paper offers a proposal – limited to questions of data sovereignty and data security, as an overlay on existing law.

Why this twofold limitation – thematically and in relation to existing regulatory frameworks – is necessary and sensible is made explicit in the following chapters.

2 What This Paper Addresses – and What It Does Not

2.1 The Scope

This impulse paper addresses two connected subject areas: **data sovereignty** (control over the data generated during the operation of smart robotic systems, including storage location, key management, tenant isolation and protection against state access) and **data security** (protection of these data against unauthorised access, loss of integrity and interruption of availability). Within this scope, it also addresses the directly adjacent fields: regulatory classification under the GDPR and the AI Act, the protection of data subjects in the sense of fundamental data-protection rights, and operational continuity as a precondition for data not to be lost in the event of provider insolvency.

2.2 Outside the Scope: Technical and Functional Safety

The paper deliberately does not address questions of technical and functional safety. These deserve their own conceptual framework with their own professional culture, their own normative references and their own circle of participants (trade associations, test laboratories, mechanical engineering bodies, aviation authorities). Mixing both fields in a single paper would do justice to neither.

Concretely, the following questions are excluded from the scope of this paper:

- Emergency-stop logic and shutdown concepts for dynamically stable and flight-capable systems
- Behaviour during power loss in bipedal humanoids and quadrupeds
- Power-and-force-limiting thresholds in non-industrial contexts
- Dynamic stability, fall behaviour and crash risks
- Aerospace intersection issues (UAS, aviation approvals) in conjunction with the Machinery Regulation
- Certifiability of learning AI components as a safety function under ISO 13849-1
- Approval and type-examination issues by notified bodies

2.3 A Concrete Example of the Need for Clarification: The Emergency-Stop Conflict

Why these questions cannot simply be appended to a data-sovereignty paper is made especially clear by a single example: the emergency-stop problem in humanoid robots.

Existing machine-safety standardisation – in particular ISO 13850 – starts from a simple foundational assumption: the energetically disconnected system is the safe state. For stationary industrial robots, for wheel-based cobots and for most classical machines, this is correct. For a bipedal humanoid robot, it is not. Cutting power to the actuators of a walking humanoid via emergency stop provokes an uncontrolled fall – with damage to people in the vicinity, to the device itself and possibly to sensitive objects. For a flight-capable system, the problem is even more drastic: a multicopter falls from the sky on emergency stop. An aircraft has, for good reason, no emergency-stop button in the cockpit.

The obvious alternative would be to enclose humanoid robots in the EU the way classical industrial robots are enclosed: physically separated, in a protective cage, with safety light barriers. That, however, is functionally absurd. A care companion in a cage has no purpose. Neither does a household assistance robot. Neither does a cooking robot in a restaurant or a delivery robot in a public space. The very justification of these systems depends on their ability to move uncaged in human environments.

Between "emergency stop with collapse" and "enclosure in a cage" lies the actual solution space – and it is largely open from a standardisation point of view. Conceivable approaches range from graceful degradation (controlled sitting or bracing) through adaptive shutdown logic (depending on the motion state) to redundant fallback systems (analogous to aviation practice). ISO 25785-1 for dynamically stable systems is at the working-draft stage and will at least partially address these questions, but its publication date is open.

2.4 Urgency: Why a Second Impulse Paper Must Follow

The Machinery Regulation (EU) 2023/1230 becomes binding on 20 January 2027 – without a transition period. Humanoid and dynamically stable systems are clearly covered by it. If no viable answer to the emergency-stop question and related problems has been developed by then, providers and operators in the EU will essentially be left with three unattractive paths: the risky implementation of classical emergency-stop logic with known side effects, the de facto enclosure of the systems in useless cages, or withdrawal from the European market. None of these paths serves the goal of making humanoid service robotics a success in Europe.

Humanide GmbH therefore announces its intention to prepare **a second impulse paper** on questions of technical and functional safety of smart robotic systems. The planned focus is on:

- Shutdown logic for dynamically stable and flight-capable systems (graceful degradation, controlled sitting, fallback redundancy)
- Demarcation from the classical emergency-stop concept under ISO 13850
- Discussion of the certifiability of learning safety functions
- Power-and-force-limiting thresholds in care and household contexts
- Interfaces with aviation approval for UAS and flight-capable robots

The second paper will be developed independently, with its own circle of participants and its own discussion architecture. The two impulse papers will, however, cross-reference each other – because in practical application, data sovereignty and technical safety belong together, even when they must be treated separately in problem analysis.

Until then, this paper refers in the relevant chapters to the respective applicable or in-development machine-safety standards without formulating its own expectations in that field. Where the matrices in Chapter 5 border on machine-safety aspects, this is expressly to be understood as an open boundary point.

3 Four Guiding Ideas of the Framework

3.1 Idea 1 – Not Form Factor, but the Interplay of Axes Decides

A humanoid-form robot is not per se more sensitive than a stationary industrial robot. A lawn-mowing robot with integrated camera and cloud connection can, in the data dimension, carry more sensitivity than a bipedally walking service robot without person recognition. The model therefore proposes six observation axes, each of which is evaluated individually. The overall classification follows the principle of the most stringent condition: the highest value on any axis determines the category. This prevents a single high-risk dimension from being averaged out by low values on other axes.

3.2 Idea 2 – Five Levels Are Enough, More Would Harm

A model that is too finely granular will not be applied in practice. Five categories (HU-RAC 1 to 5) appear to be a compromise that differentiates adequately between consumer trivialities and state-sensitive applications without being operationally overwhelming. The reference points are established tier models from other fields: the Common Criteria Evaluation Assurance Levels, the Automotive SPICE Levels, and the Performance Levels of ISO 13849-1. Five levels are applicable and accepted in these contexts.

3.3 Idea 3 – Regulatory and Commercial Perspectives Are Mirrors

The present work was inspired by the encounter with the RARI framework (Robot Adoption Readiness Index), presented by Nicolaus Radford (Persona AI) at the Humanoids Summit 2025 in Mountain View. RARI rates eight application industries on commercial adoption readiness. The ranking reaches from Heavy Industry (score 100) and Automotive (96.9) through Warehousing and Construction down to Home (near 0). Overlaying this ranking onto a classification by data sovereignty and data security reveals a surprising pattern: industries with high commercial readiness lie in a narrow middle regulatory corridor. Industries with low commercial readiness (Healthcare, Home) are spread broadly – from very low to very high. The opacity of data-protection and data-sovereignty requirements is part of the adoption friction. A clear classification model can lower this friction – and thereby unlock market potential that is currently blocked.

3.4 Idea 4 – HU-RAC Is Overlay, Not Replacement

A dense framework of applicable legal acts and standards already exists: the GDPR with Art. 32 (security), Art. 25 (privacy by design), Art. 28 (data processing agreement), Art. 35 (DPIA); the AI Act with high-risk obligations on documentation, logging, transparency and human oversight; the Machinery Regulation; the EU Product Liability Directive; the Cyber Resilience Act; ISO 10218, ISO 13482, EN ISO 12100. The gap therefore does not lie in missing rules, but in the missing coherence between these regulatory frameworks and in the missing operationalisation for concrete smart-robotics deployment scenarios.

HU-RAC explicitly positions itself not as a replacement for these legal acts and standards, but as an overlay: a domain-specific organising and translation grid that makes existing frameworks applicable in context. The model orchestrates rather than duplicates. Annex E contains a mapping scheme that makes this orchestration comprehensible in four steps: AI Act classification, HU-RAC classification, GDPR alignment, and derivation of additional HU-RAC-specific expectations. This positioning also has consequences for the expectation matrix in Chapter 5: many items marked as essential are already legally mandatory, others go beyond market standard. Chapter 5 makes this distinction visible through the split into E-Min and E-Plus.

HU-RAC thereby stands in a defined field of related classification approaches but differs, in the combination of its focus, from any known system: autonomy-grading systems such as SAE J3016 (vehicles), LASR (surgical robotics), and the Chinese standards T/CIE 298-2025 and the MIIT standard system published in 2025/2026 address capability levels and decision autonomy, not data sovereignty. The CPST taxonomy by Ning and Ding (2026) addresses the ontological question of what kind of entity an autonomous system is, and derives governance levels from it up to questions of legal personhood. The Google Robotics Privacy Framework (2024) remains principle-based (Identify, Assess, Enforce) without a graduated classification system. HU-RAC is the only model that combines the robotics deployment context with a data-sovereignty focus in an overlay logic on European law. Annex D.4 contains an overview of related classification approaches.

4 The Framework: HU-RAC in Five Categories

4.1 The Six Observation Axes

Axis 1 – Deployment Environment

The environment in which the system operates determines its structuredness, public exposure and legal framework. The range spans from private outdoor use to commercial-enclosed and publicly intensive areas, and up to care, education, childcare and militarily sensitive contexts.

Axis 2 – Form Factor and Mobility

Stationary, wheel-based mobile, flight-capable, bipedal or dynamically stable, aquatic. Each form factor triggers different machine-safety aspects and different normative references.

Axis 3 – Data Categories

From pure machine telemetry through person-identifiable environmental data and biometric data up to intimate long-term observation (everyday life, conversations, care situations). The data category determines the data-protection gravity.

Axis 4 – AI Risk Class

Classified by the AI Act: from trivial functions through limited risk and high risk up to high risk with vulnerable target group or state sensitivity.

Axis 5 – Physical Hazard Potential (as Context Information)

Kinematic energy, reach, interaction capability with humans. From power-and-force-limiting class through cobot level up to weapons-capable or dual-use-sensitive. This axis serves exclusively as context information for classification into an HU-RAC category; machine-safety-related expectations are deliberately not derived from it in this paper (see Chapter 2).

Axis 6 – Affected Persons

Who is legally and ethically touched? Only the operator, employees, uninvolved third parties, vulnerable groups (patients, children, people with dementia), or legally specially protected circles (bearers of professional secrets).

4.2 The Aggregation Logic

The HU-RAC framework proposes determining the overall category through the highest value on any of the six axes. A system that would rank low on five axes and reach the highest level on one axis would be thought of as a Tier-5 system. This rule follows the principle of the most stringent condition, analogous to the hazardous-goods classification logic and to the risk-oriented classification practice of the Machinery Regulation. It prevents a single critical dimension from being masked by averaging.

Exception 1 – Technically Enclosed Temporary High-Risk Modes

A temporally or functionally limited high-risk mode (for example, a diagnostic mode with one-time biometric capture) does not automatically lead to the elevation of the overall category to Tier 4 or 5, provided that the following conditions are cumulatively met:

- The mode can only be activated with the explicit consent of the affected person or by authorised specialist personnel.
- Strong technical safeguards are active (local processing without cloud connectivity, air-gap from general telemetry, immediate deletion after purpose fulfilment).
- Every activation is logged and auditable.

- The functional limitation is documented in the technical documentation and in the declaration of conformity.

If these conditions are met, the classification follows the regular operation. Otherwise the basic rule applies.

Exception 2 – Research Deployments

For pure research deployments within the meaning of Art. 2(6) AI Act and under appropriate safeguards (ethics committee, limited circle of participants, anonymisation, Art. 89 GDPR), a classification regime one tier lower may be applied for the duration of the research, provided this is explicitly justified in the DPIA.

Second Aggregation Rule – Multiple Mid-Level Load

To prevent under-classification of multi-faceted systems, an additional rule applies: if three or more of the six axes lie in the mid-range, at least HU-RAC 3 is reached, even if no single axis reaches the highest level. This ensures that complex systems with broad mid-level load are not classified as consumer applications.

4.3 The Five Classification Categories

HU-RAC 1 – Baseline

Smart robotic systems without person reference in productive operation, low hazard, no AI high-risk classification, no vulnerable target groups. Typical reference examples: lawn-mowing robots in private gardens, pool-cleaning robots, autonomous lawn irrigation, simple outdoor agricultural robots on delimited areas.

HU-RAC 2 – Enhanced

Systems with environmental sensors that can technically capture persons but are not primarily used for biometric identification. In private or commercial-enclosed spaces. Typical reference examples: vacuum-cleaning robots, window-cleaning robots, simple indoor drones, household assistance robots without companion function, warehouse logistics systems in enclosed areas.

HU-RAC 3 – Professional

Commercial-industrial deployment with ERP or process integration, or commercial environments with public contact without vulnerable target groups. Typical reference examples: cobots and industrial robots with human-robot collaboration, welding robots in heavy industry, logistics humanoids, process-integrated AI bridges, cleaning robots in public spaces, cooking robots in restaurants, autonomous delivery robots in last-mile delivery.

HU-RAC 4 – Sensitive

Deployment in care, health, education, childcare, or with biometric data capture and/or vulnerable target groups. Typical reference examples: companion and care robots in senior residences and home-care situations, surgical-assistance systems, therapy robots, robots in childcare facilities, guard robots with person recognition, humanoid home assistants in care contexts.

HU-RAC 5 – Critical / Sovereign

Military or dual-use applications, deployment in critical infrastructures, high-risk deployments with state-related or legally specially protected sensitivity (justice, law enforcement, intelligence-related environments). This category is deliberately provided in the model, even though its specific configuration is essentially governed by specialised frameworks (BSI regulations, NATO STANAGs, national VS rules). The framework here only maps the civilian outline on which these specialised frameworks can build.

5 Expectation Matrix for Discussion

The following section translates the five categories into a differentiated expectation profile – limited to the fields of data sovereignty, data security, regulatory classification, AI reference and data-subject protection. The matrices are intended as illustrative offers of what a graduated regime could look like – not as a list of binding obligations.

Expectations from the area of technical and functional safety (emergency-stop logic, power-and-force-limiting, dynamic stability, power-loss behaviour) are deliberately not represented. They will be treated independently in the announced second impulse paper. Where standards from the machine-safety area are occasionally cited in the following matrices (for example concerning labelling or transparency), this is done only to the extent that they are also relevant to data protection or information law.

Legend

E-Min	Minimum essential – in this category predominantly already expected under applicable law (GDPR, AI Act, Machinery Regulation, CRA) or established market standard.
E-Plus	Essential for a conservative fundamental-rights and resilience baseline – goes beyond the current market standard and is considered aspirationally necessary in this category.
R	Recommended – regularly recommended, but not uniformly seen as indispensable.
O	Optional – considered in the professional discussion, but without a uniform expectation.
–	Not regularly raised in the discussion for this category.

The split into E-Min and E-Plus allows a realistic self-assessment: start-ups and SMEs can begin with E-Min and declare E-Plus as a development roadmap without falling out of the HU-RAC category. The matrices are cumulative in intent: a higher category encompasses the underlying considerations and adds further ones.

5.1 Data Sovereignty and Data Protection

Aspect	HU-RAC 1	HU-RAC 2	HU-RAC 3	HU-RAC 4	HU-RAC 5
Data processing in EU/CH	R	E-Min	E-Min	E-Min	E-Min
Independence from US hyperscalers	–	R	E-Plus	E-Plus	E-Min
Complete DPA under Art. 28 GDPR	R	E-Min	E-Min	E-Min	E-Min
Art. 48 GDPR clause	–	E-Plus	E-Min	E-Min	E-Min
Disclosure of sub-processors	–	E-Min	E-Min	E-Min	E-Min
Consent requirement for changes	–	–	E-Plus	E-Min	E-Min
Warrant canary / transparency obligation	–	–	R	E-Plus	E-Plus
Art. 9 GDPR concept for biometrics	–	–	R	E-Min	E-Min
DPIA with data-flow documentation	–	–	R	E-Min	E-Min
Sovereign cloud (C5-certified)	–	–	O	R	E-Min

5.2 AI Reference under the AI Act

Aspect	HU-RAC 1	HU-RAC 2	HU-RAC 3	HU-RAC 4	HU-RAC 5
Documented AI classification	–	E-Min	E-Min	E-Min	E-Min
Technical documentation Art. 11	–	–	E-Min	E-Min	E-Min
Recording obligations Art. 12	–	–	E-Min	E-Min	E-Min
Transparency obligation Art. 13	–	–	E-Min	E-Min	E-Min

Aspect	HU-RAC 1	HU-RAC 2	HU-RAC 3	HU-RAC 4	HU-RAC 5
Human oversight Art. 14	–	–	E-Min	E-Min	E-Min
Explainability for safety-critical decisions	–	–	R	E-Plus	E-Plus
ISO/IEC 42001 (AI management system)	–	–	O	R	E-Plus
AI readiness attestation by third party	–	–	O	E-Plus	E-Plus
Opt-out for continuous field learning	–	E-Min	E-Min	E-Min	E-Min
No implicit learning transfer between customers	–	R	E-Min	E-Min	E-Min

5.3 Cybersecurity

Aspect	HU-RAC 1	HU-RAC 2	HU-RAC 3	HU-RAC 4	HU-RAC 5
CRA readiness (Reg. 2024/2847)	E-Min	E-Min	E-Min	E-Min	E-Min
Software Bill of Materials (SBOM)	–	R	E-Min	E-Min	E-Min
Vulnerability disclosure process	–	R	E-Min	E-Min	E-Min
CVE monitoring and response	–	R	E-Min	E-Min	E-Min
Signed updates with rollback	E-Min	E-Min	E-Min	E-Min	E-Min
Corruption protection per MR	R	E-Min	E-Min	E-Min	E-Min
Incident reporting with SLAs	–	R	E-Min	E-Min	E-Min
Bug bounty programme	–	–	R	E-Plus	E-Plus
ISO/IEC 27001	–	O	R	E-Min	E-Min
BSI Grundschutz certification	–	–	O	E-Plus	E-Min

5.4 Key Sovereignty and Zero Knowledge

Aspect	HU-RAC 1	HU-RAC 2	HU-RAC 3	HU-RAC 4	HU-RAC 5
Transport encryption (TLS 1.3)	E-Min	E-Min	E-Min	E-Min	E-Min
Data encryption at rest	R	E-Min	E-Min	E-Min	E-Min
Bring Your Own Key (BYOK)	–	O	E-Plus	E-Min	E-Min
Hold Your Own Key (HYOK)	–	–	R	E-Plus	E-Min
Hardware Root of Trust (TPM/SE)	–	–	R	E-Plus	E-Min
Customer-side HSM	–	–	O	E-Plus	E-Min
Zero-knowledge architecture	–	–	O	E-Plus	E-Plus
M-of-N recovery (Shamir Secret Sharing)	–	–	–	E-Plus	E-Plus
Independent cryptography audit	–	–	–	E-Plus	E-Plus
BSI TR-03145 (Secure CA)	–	–	O	R	E-Plus

5.5 Tenant Isolation and Backend Security

Aspect	HU-RAC 1	HU-RAC 2	HU-RAC 3	HU-RAC 4	HU-RAC 5
Single-tenant deployment	–	R	E-Plus	E-Min	E-Min
Zero-trust authorisation at API level	–	E-Min	E-Min	E-Min	E-Min
External pen-tests of authorisation logic	–	–	R	E-Min	E-Min
Red-team exercises tenant isolation	–	–	O	E-Plus	E-Min

Aspect	HU-RAC 1	HU-RAC 2	HU-RAC 3	HU-RAC 4	HU-RAC 5
Separated network segments (OT-VLAN)	–	R	E-Min	E-Min	E-Min
Session recording for remote support	–	–	R	E-Min	E-Min
Customer-initiated remote support	–	–	E-Min	E-Min	E-Min
Public disclosure process	–	–	R	E-Plus	E-Plus

5.6 Data-Subject Protection, Transparency and Ethics

Aspect	HU-RAC 1	HU-RAC 2	HU-RAC 3	HU-RAC 4	HU-RAC 5
Transparent labelling as a robot	R	E-Min	E-Min	E-Min	E-Min
Opt-out of data transmission	R	E-Min	E-Min	E-Min	E-Min
Deactivatable telemetry	–	E-Min	E-Min	E-Min	E-Min
No continuous biometric surveillance	–	E-Min	E-Min	E-Min	E-Min
Child-protection concept for target group	–	–	–	E-Min	E-Min
Ethics committee involvement (care)	–	–	–	E-Min	E-Min
Documented dignity-preservation principles	–	–	–	E-Plus	E-Plus

Rows on emergency stop, normal-stop function and force-protection shutdown are deliberately not listed here. They belong to the field of technical and functional safety and will be addressed in the announced second impulse paper.

5.7 Standards Connection and Future Viability

Aspect	HU-RAC 1	HU-RAC 2	HU-RAC 3	HU-RAC 4	HU-RAC 5
Implementation of new standards within 12 months	–	R	E-Min	E-Min	E-Min
Proactive notification on updates	–	R	E-Min	E-Min	E-Min
Open interfaces (ROS2 or similar)	–	O	R	E-Plus	E-Plus
Participation in ISO TC 299 WG 12	–	–	O	R	R
Participation in DIN NA 060-30	–	–	O	R	R
Participation in CEN/CENELEC JTC 21	–	–	O	R	R
Participation in ETSI TC CYBER	–	–	O	R	R
Annual transparency report	–	–	R	E-Plus	E-Plus

6 Two Additional Observation Levels

6.1 Business Criticality

The HU-RAC category alone represents the infringement and regulatory risks. It does not yet represent how critical the failure of the system is for the operator. An HU-RAC-3 cobot in a paced production line and an HU-RAC-3 research robot are regulatorily equivalent, but operationally extremely different. A failure in the first case means production standstill, in the second only a postponed test series.

The framework therefore envisages an orthogonal business-criticality classification with three levels:

Level A	Process- or safety-critical. Failure leads to production standstill, patient endangerment or supply-chain interruption. Reference examples: ERP-integrated automation bridge, surgical assistance, industrial cobot in paced production, care companion in unattended living situation, critical-infrastructure inspection robot.
Level B	Operationally important, but bridgeable. Reference examples: warehouse logistics with manual fallback option, commercial cleaning robots, restaurant cooking robots, agricultural robots.
Level C	Not critical. Failure tolerable. Reference examples: research robots, demo devices, single-workstation automation, private household robots.

Operationalisation of the Levels

To make classification less subjective and thereby more comparable, the following orientation criteria are proposed. Level A applies if at least one of the following conditions is met:

- The system contributes more than 20% to the operator's core operational performance or revenue.
- An outage of more than four hours leads to patient endangerment, interruption of medically necessary care, or a critical-infrastructure outage.
- Operations are subject to statutory availability obligations (e.g. under the Critical Infrastructure Umbrella Act, German Social Code, medical-device law) with defined service levels.
- Recovery after an outage requires more than 72 hours or external service providers.

Level B applies if the system is operationally important but an outage remains bridgeable within 24 hours through manual or alternative procedures and no patient endangerment arises. Level C applies if an outage does not immediately impair operations and no third parties are affected.

Minimum Standards for Degraded Mode

The combination of HU-RAC category and business-criticality level provides a comprehensible framework for minimum standards for continued operation in the event of manufacturer-cloud failure:

HU-RAC × Level	Minimum Cloud Autonomy	Further Requirements
HU-RAC 4 × A	24 months	Basic functions local; data export possible at any time; source and data escrow mandatory
HU-RAC 4 × B	12 months	Basic functions local; data export; source escrow recommended

HU-RAC 3 × A	12 months	Core operation local; data export; exit clauses contractual
HU-RAC 3 × B	6 months	Core operation local; data export
HU-RAC 2 × C	Recommended	Easy data export possible; degraded mode optional

The threshold values are intended as discussable orientation and can be refined in concrete standardisation procedures.

6.2 Operational Continuity and Provider Resilience

Insolvencies of cloud-dependent providers in adjacent market segments (see Chapter 1.1) suggest that provider resilience and continuity of operations deserve independent observation dimensions. The framework proposes three interlocking mechanisms whose intensity scales with the HU-RAC category and the business-criticality level.

Architectural Continuity

- From HU-RAC 2 onwards, the ability to continue operating without a cloud (degraded mode) appears sensible.
- From HU-RAC 3 onwards, open interfaces (ROS2 or equivalent) and data portability in standard formats appear appropriate.
- From HU-RAC 4 onwards, an on-premises operating option that functions at least 24 months without the manufacturer's cloud appears worth discussing.
- For HU-RAC 5, air-gap capability would be expected.

Contractual Continuity

A simple data-export claim from HU-RAC 1 onwards, source-code escrow from HU-RAC 3 onwards (especially for Level A), data escrow and cloud-continuity guarantee from HU-RAC 4 onwards. The concrete contractual drafting lies outside the impulse idea and would have to be shaped by counsel.

Provider Resilience Rating

Three classes derived from financial stability, ownership structure and operational resilience: R1 (established large enterprise), R2 (stable mid-sized enterprise), R3 (start-up with elevated insolvency or change-of-control risk). For HU-RAC 4 and 5, R3 would only be justifiable with fully developed continuity mechanisms.

7 The Role of Operators

The framework has so far primarily addressed manufacturers and standardisation actors. In practical application, however, it is often the operators – hospitals, care homes, industrial enterprises, municipal institutions – who decide on the success or failure of a robotics deployment. They select, procure, operate, document and share liability. A classification system that omits this perspective remains incomplete.

7.1 Distribution of Responsibilities Along the HU-RAC Logic

The distribution of roles between manufacturer and operator generally follows existing law. The framework fits into this distribution without legal changes:

- Manufacturers are typically controllers within the meaning of Art. 4 GDPR for product development, for the technical documentation under AI Act and Machinery Regulation, and for the declaration of conformity. In the HU-RAC logic, they provide the classification basis: they pre-classify their product, substantiate the classification comprehensibly and name which expectations are met in which category.
- Integrators (system integrators, implementation partners) carry responsibility for embedding into the respective IT landscape, for interfaces to existing systems, and for technical configuration. In the HU-RAC logic, they ensure that the classification remains valid after integration – for example that tenant isolation is not undermined by configuration errors.
- Operators are controllers within the meaning of Art. 4 GDPR for the deployment context. They conclude data-processing agreements under Art. 28 GDPR, carry out DPIAs under Art. 35 GDPR, document technical and organisational measures, and bear the operator obligations under the Machinery Regulation. In the HU-RAC logic, they check whether the manufacturer's classification matches the concrete deployment context, and raise it where the context deviates.

7.2 HU-RAC in Procurement and Tendering

For practical anchoring in procurement processes, the framework could be applied at several points:

- Performance specifications in tenders could name an HU-RAC target category and refer to the corresponding expectation matrix. Bidders would be required to submit their self-classification plus evidence.
- In public-procurement proceedings, HU-RAC matrix criteria could be mapped as eligibility or award criteria, particularly in the health and social sector and in critical-infrastructure-adjacent municipal procurement.
- For contracts with longer duration (leasing, rental, service models), exit clauses, data-portability guarantees and escalation mechanisms could be designed with reference to the HU-RAC category.
- Operator policies and data-protection management systems could use HU-RAC classifications as an input to internal risk assessment and to the design of technical and organisational measures.

7.3 A Brief Note on Practical Feasibility

For smaller operators – municipal bodies, mid-sized care facilities, clinics without a central IT department – the large number of requirements in Chapter 5 is potentially overwhelming. The E-Min/E-Plus split addresses this problem in part: an operator can rely on the E-Min columns as minimum selection criteria and treat E-Plus as a differentiating factor. The

associations of the respective sectors (care associations, hospital associations, municipal umbrella organisations) could support uptake by developing sector-specific procurement guidelines on an HU-RAC basis.

8 Impact of the Framework: Market Opening Rather Than Compliance Burden

A classification system acts not only as a burden but also as a bridge. A provider developing a care robot today has to work against a diffuse cloud of regulatory uncertainty. What data-protection measures are appropriate in care situations? What cryptographic architecture is expected? What level of explainability for AI-assisted decisions is indicated? The answers to these questions are today conversations declined by purchasers, cancelled pilot projects, failed funding rounds. Each of these answers costs the market time and capital.

A clear classification model lowers this friction. The provider knows in which category its product operates. The procurer knows which expectation profile is appropriate. The investor can estimate the regulatory capital and time costs. Market supervision has a reference framework that makes plausibility verifiable.

The comparison with the RARI framework underscores this argument. Industries with a broad regulatory corridor – Healthcare, Home – are at the same time those with the lowest commercial adoption readiness. This is no coincidence. Regulatory uncertainty is part of the adoption friction. A framework that comprehensibly closes the corridor thus simultaneously acts as a market-opening instrument.

Put differently: economically, such a system is an investment in market development, not only a compliance cost item. From the author's perspective, this twofold effect would be a worthwhile leitmotif for further discussion and, where appropriate, for a later transition into a standardisation procedure.

9 Where the Framework Could Be Connected

This impulse paper is not a standard and does not wish to become one – unless the resonance from professional circles makes the step plausible. For that case, the following connection possibilities are conceivable:

- An impulse exchange with DIN NA 060 (Robotics) on the suitability of the framework as a starting point for a DIN SPEC under the PAS procedure.
- A substantive discussion with ISO TC 299 WG 12, which deals with humanoid and dynamically stable systems.
- A scientific contribution in cooperation with Fraunhofer IAO or IPA that empirically substantiates or refutes the economic impact hypothesis (market opening through classification).
- A dialogue track with the German Federal Office for Information Security (BSI) regarding the cryptography and zero-knowledge aspects in HU-RAC 4 and 5.
- A resonance loop with Persona AI and other industry partners already working on adoption-readiness maturity models.
- Inclusion of civil-society voices (data-protection authorities, Stiftung Datenschutz, Bitkom, care associations, patient representatives).
- Publication of the announced second impulse paper on technical and functional safety, which will address the gaps identified in Chapter 2.

The author does not regard this paper as a final word, but as an invitation to discussion. Feedback, criticism and proposals for extensions are expressly welcome. Contributions on the open technical-safety questions sketched in Chapter 2 are particularly welcome and will inform the preparation of the second paper.

Annex A – The Five Commercial Dimensions (RARI Reference)

This annex supplements the regulatory and risk-oriented framework with the commercial-operational perspective. The reference is the RARI framework (Robot Adoption Readiness Index), presented by Nicolaus Radford, CEO of Persona AI, at the Humanoids Summit 2025 in Mountain View, California, on 11 December 2025.

A.1 The Five Evaluation Dimensions

Labor Scarcity	Skills shortage in the deployment segment. High scarcity increases willingness to pay and pressure to automate.
Labor Cost	Wage level in the deployment segment. High wages justify higher capital costs for humanoid replacement labour.
Task Ease	Simplicity and clarity of the tasks to be automated. Structured, repeatable tasks are technically more accessible.
Env. Predictability	Structuredness and predictability of the work environment. Industrial and logistics environments are typically structured; domestic environments are not.
Dev. Ease	Development effort for industry-specific adaptation. Low development friction accelerates market entry.

A.2 The Eight RARI Reference Industries with Normalised Scores

Industry	RARI Score (0–100)	Adoption Readiness
Heavy Industry	100.0	very high
Automotive	96.9	very high
Warehousing	72.5	high
Construction	58.5	medium
Healthcare	~ 40	low-medium
Retail/QSR	39.7	low-medium
Agriculture	39.7	low-medium
Home	near 0	very low

The scores are cited here as a secondary source to make the commercial perspective transparent. They are not part of any normative determination of this impulse paper. Authorship and interpretive authority rest with Persona AI.

Annex B – Twelve Application Domains at a Glance

The range of smart robotics cannot be reduced to a few application fields. The following overview extends the eight RARI industries by four further fields (Logistics/Last-Mile, Hospitality, Public Sector/Security, Defence/Dual-Use) in order to represent the diversity adequately. For each domain, a typical HU-RAC range, a qualitative adoption-readiness assessment, and references to relevant additional legal layers are indicated.

Domain	HU-RAC Range	Adoption Readiness	Additional Legal Layers
Heavy Industry	3	very high	ATEX Directive 2014/34/EU; Industrial Safety Ordinance; professional-association rules
Automotive	3	very high	ISO 26262; TISAX; UN-R155/156 for vehicle integration
Warehousing	2–3	high	GDP; ISO 13485 in pharma/medical logistics; IFS/BRC for food
Construction	3	medium	Construction Products Regulation (EU) 305/2011; construction-site rules
Healthcare	2–5	low-medium	MDR (EU) 2017/745; national medical-device law; IEC 60601 family; social-security law
Retail/QSR	2–3	low-medium	Food hygiene regulations; price indication; youth-protection law where applicable
Agriculture	1–3	low-medium	Reg. (EU) 167/2013; fertiliser ordinance; plant-protection law
Home	1–4	very low	Household exception Art. 2(2)(c) GDPR; product-safety law
Logistics / Last-Mile	2–3	high	Road traffic law (public space); EU 2018/1139 for drones; transport law
Hospitality	2–3	medium	Hospitality law; food hygiene; accessibility law; GDPR for guest data
Public Sector / Security / Emergency	4–5	medium	Police laws; federal criminal-police act; BSI Act; Critical Infrastructure Umbrella Act
Defence / Dual-Use	5	not applicable	Dual-Use Reg. 2021/821; foreign trade law; BSI classified regulations; NATO STANAGs

Annex C – Eight Illustrative Examples

The following examples illustrate how the framework could be applied in concrete cases. They are deliberately intended as illustration and do not replace case-by-case analysis. Examples C.6 to C.8 deliberately address edge cases in which the robustness of the model under complex conditions becomes visible.

C.1 Lawn-Mowing Robot in Private Garden

Deployment environment private-outdoor (low), form factor wheel-based (low), data categories telemetry (low), AI risk trivial (low), physical hazard potential PFL class (low), affected persons only the operator (low). Highest value low. Illustrative classification: HU-RAC 1, Level C. Domain: Home.

C.2 Industrial Cobot with ERP Integration

Deployment environment industrial-manufacturing (medium), form factor stationary (medium), data categories process data with person reference (medium), AI risk high risk under the AI Act (high), physical hazard potential cobot class (medium), affected persons employees (medium). Highest value high on the AI axis. Illustrative classification: HU-RAC 3, Level A. Domain: Heavy Industry or Automotive.

C.3 Care Companion in Senior Residence

Deployment environment sensitive-care (very high), form factor mobile or dynamically stable (high), data categories biometric and intimate long-term observation (very high), AI risk high risk with vulnerable target group (very high), physical hazard potential medium to high, affected persons vulnerable individuals (very high). Highest value very high on several axes. Illustrative classification: HU-RAC 4, typically Level A, R1 or R2 indicated. Domain: Healthcare with MDR review depending on intended purpose.

C.4 Autonomous Delivery Robot in Last-Mile Delivery

Deployment environment commercial with public contact (high), form factor wheel-based mobile (medium), data categories environmental sensors with person capture (medium), AI risk limited to high risk (medium-high), physical hazard potential medium, affected persons uninvolved third parties (high). Highest value high. Illustrative classification: HU-RAC 3, Level B. Domain: Logistics/Last-Mile.

C.5 Inspection Robot in Critical Infrastructure

Example substation, grid control room or waterworks: deployment environment critical-infrastructure sensitive (very high), form factor variable (medium), data categories state-security-relevant operational data (very high), AI risk high risk with state sensitivity (very high), physical hazard potential medium, affected persons operator-side (medium). Illustrative classification: HU-RAC 5, Level A, R1 recommended. Domain: Public Sector / Security / Emergency with reference to the BSI Critical Infrastructure Ordinance.

C.6 Research Robot at a University Hospital (Edge Case)

A humanoid research system is used at a university hospital to test a new care-interaction concept. Without contextual consideration, the classification would be HU-RAC 4 (vulnerable target group, sensitive deployment environment). The deployment, however, takes place under ethics-committee supervision, with a limited circle of participants, with explicit consent under Art. 9(2)(a) GDPR and under the research exceptions of Art. 89 GDPR. Under the

second exception rule in Chapter 4.2, the classification may be reduced by one level for the duration of the research. Illustrative classification: HU-RAC 3 for the research phase, with documented DPIA, Level B, R2 or R1. On transition to regular operation, HU-RAC 4 applies again. Lesson: the model remains applicable in research contexts without disproportionately restricting research freedom.

C.7 Consumer Humanoid with Optional Subscription Service (Edge Case)

A humanoid assistance robot for private households is marketed as a consumer product (basic function offline, Home domain, household exception under Art. 2(2)(c) GDPR). Optionally, a cloud subscription with extended functions (routine learning, speech recognition, remote access for relatives) can be added. Without the subscription, the classification is HU-RAC 2 (humanoid form factor, sensors with person capture, but the household exception applies). With the subscription, the household exception falls away because the provider processes the data as controller – the classification rises to HU-RAC 3 to 4, depending on the extent of biometric processing and the target group (children, elderly people in the household). In addition, the platform risk arises in the event of provider insolvency: without the subscription the device remains functional, with the subscription it loses its core functions. Illustrative classification: HU-RAC 2 in the baseline state, HU-RAC 3-4 with active subscription; Level B to A depending on household context. Lesson: dynamic classifications are conceivable, but they must be clearly identified contractually and technically.

C.8 Hospital Logistics Robot in a Mixed Environment (Edge Case)

An autonomous delivery robot distributes medication, samples and laundry between wards. It moves both in non-public service corridors and in semi-public areas (elevators, hallways with patient traffic). The data categories include transport logs with patient reference (sample assignment), but no direct biometrics. Deployment environment mixture of sensitive and public-facing (high), form factor wheel-based (medium), data categories indirectly patient-related (medium-high), AI risk limited to high risk (medium), physical hazard potential medium, affected persons mixed (medium-high). Under the principle of the most stringent condition: HU-RAC 3 for regular operation. Under the second aggregation rule (several mid-level axes cumulatively), at least HU-RAC 3 is reached – particularly plausible here because of the mixed environment. Illustrative classification: HU-RAC 3, Level A because of medication-logistics criticality. Lesson: mixed environments are not adequately characterised by a single highest-value classification; the supplementary rule in Chapter 4.2 applies and lifts the classification where individual axes alone cannot.

Annex D – Sources and Reading Recommendations

D.1 Legal Acts

- Regulation (EU) 2016/679 – General Data Protection Regulation
- Regulation (EU) 2023/1230 – Machinery Regulation
- Regulation (EU) 2024/1689 – AI Act
- Regulation (EU) 2024/2847 – Cyber Resilience Act
- Directive (EU) 2024/2853 – Product Liability Directive
- Regulation (EU) 2021/821 – Dual-Use Regulation

D.2 Technical Standards and Regulatory Frameworks

- ISO 10218-1/-2:2025 – Safety requirements for robots
- ISO/DIS 25785-1 – Safety of dynamically stable mobile robots
- ISO 13482 – Safety requirements for personal care robots
- EN ISO 12100 – Risk assessment
- ISO 13849-1 – Performance Levels
- ISO/IEC 42001 – AI management system
- BSI TR-03145 and BSI C5

D.3 Professional Literature and Industry Reports

- Radford, N. (2025): Humanoid Adoption Readiness – The RARI Framework. Keynote, Humanoids Summit 2025, Mountain View, 11 Dec 2025. Persona AI Inc.
- Morgan Stanley Research (2025): Humanoid Robotics – The Future of Labor (A. Jonas et al.).
- IDTechEx (2025): Humanoid Robots 2025–2035: Technologies, Markets and Opportunities.
- Bain & Company (2025): Humanoid Robots – From Demos to Deployment. Technology Report 2025.
- Yole Group (2025): Humanoid Robots 2025 – The Race to Useful Intelligence.

D.4 Related Classification Approaches

The following approaches address adjacent classification and governance questions in robotics. None of them, however, combines the combination constitutive for HU-RAC of robotics deployment context, data-sovereignty focus and overlay logic on European law. The field is therefore open for a European data-sovereignty classification system.

Ontological Governance Classification

- Ning, Huansheng; Ding, Jianguo (2026): Beyond Tools and Persons: Who Are They? Classifying Robots and AI Agents for Proportional Governance. arXiv:2604.05568, submitted 7 April 2026. The model uses Cyber-Physical-Social-Thinking (CPST) space theory with four dimensions (computational, embodied, relational, cognitive) and a three-tier taxonomy (Confined Actors, Socially-Aware Interactors, CPST-Integrated Agents). Governance focus: product liability up to qualified legal personhood. Structurally related to HU-RAC through multi-dimensional aggregation; substantively complementary through an ontological rather than contextual focus.

Autonomy and Intelligence Grading (Robotics and Vehicles)

- SAE International (2018, continuously updated): J3016 – Taxonomy and Definitions for Terms Related to On-Road Motor Vehicle Automated Driving Systems. Six levels (Level 0–5) of driving automation. Methodological reference point for HU-RAC's tier logic; not substantively concerned with data sovereignty.
- Lee, A. et al. (2024): Levels of Autonomy in Surgical Robotics (LASR). Systematic review of FDA-cleared surgical robots, defining five levels from Robot Assistance to Full Autonomy. The framework builds on earlier autonomy taxonomies for surgical robotics by Yang et al., Haidegger et al. (2019), Fosch-Villaronga et al. and Attanasio et al.
- SemiAnalysis (2025): Robotics Levels of Autonomy – A five-level classification model for general robotics, oriented to commercial maturity and capabilities along the axes Agency and Dexterity.
- Beijing Humanoid Robot Innovation Center (May 2025): T/CIE 298-2025 – Humanoid Robot Intelligence Grading Standard. Four-Dimension, Five-Level framework with 22 primary indicators. Addresses intelligence grading, not data sovereignty.
- Ministry of Industry and Information Technology (MIIT), China (February 2026): Humanoid Robot and Embodied Intelligence Standard System (2026 Edition). National standards framework with six pillars. Worldwide first state-sponsored humanoid robotics standards framework; focus on technical capability, autonomy and system design.
- NLJIC/Shanghai (2024): Industry standards with four technical levels (L1–L4) and five intelligence levels (G1–G5) for Chinese humanoids.

Robotics Privacy Frameworks

- Google / Relford, E. (2024): Robotics Privacy Framework. Principle-oriented approach with three focal areas: Novel Interactions, Autonomous Engagement, Transparency. No graduated classification; orientation framework rather than classification system.
- Rueben, M. et al. (2016): Privacy in Human-Robot Interaction – Survey. Classifies territorial areas into four categories (individual, shared, public, jurisdiction). Spatial-territorial focus, not data-sovereignty focus.
- Fosch-Villaronga, E. (2021): Toward Privacy-Sensitive Human-Robot Interaction. Analysis of the privacy policies of eight leading consumer-robotics manufacturers; not a classification model but evidence for the heterogeneity of current manufacturer practice.

Ethics and Assurance Standards

- IEEE 7000 family (since 2021): among others IEEE 7005 (Transparent Employer Data Governance), IEEE 7007 (Ontologically Driven Robotics), IEEE P7018 (Security/Trustworthiness Generative AI). Thematic-ethical standards without a graduated classification logic.
- BS 8611 (2016): Guide to the Ethical Design and Application of Robots and Robotic Systems. Principle-driven British standard on ethical aspects (privacy, bias, explainability, worker impact).

Data Classification Frameworks (Adjacent, Non-Robotics-Specific)

- Microsoft Service Assurance (ongoing): Data Classification Framework. Three- to five-tier data classification (Public, Internal, Confidential, Highly Confidential) with associated controls. Methodological reference point for HU-RAC's tier logic and E/R/O logic; substantively concerned with enterprise data, not robotic systems.

- ISO 27701 (2019): Privacy Information Management System. International standard for data-protection management systems; forms the basis for many privacy-governance frameworks.

D.5 Relevant Incidents and Cases

- Backend authorisation flaw at a leading manufacturer of autonomous household robots (February 2026): third-party access to approximately 6,700 devices through a flaw in the tenant isolation of the MQTT backend.
- Withdrawal of a zero-knowledge cloud protection feature by a market-dominant US technology group in the United Kingdom (February 2025) following an order by the British security authorities.
- Insolvency of a European manufacturer of networked blood-pressure devices (2024): loss of cloud infrastructure and user data as a pattern of cloud-dependent outages; comparable cases at several smart-home and connected-health providers.

Annex E – Mapping Scheme: Orchestrating Existing Frameworks

The HU-RAC framework understands itself as an overlay on existing law (see Idea 4 in Chapter 3.4). The following four-step process shows how HU-RAC can be embedded in a classification decision without competing with the GDPR, the AI Act or the Machinery Regulation. The scheme is intended as a workflow for manufacturers, integrators and operators.

E.1 Four-Step Process

Step	Framework	What to Check
1	AI Act (Reg. 2024/1689)	Does the system fall into a high-risk category under Annex III or as a safety component? Do the obligations under Art. 9 ff. apply (risk management, data quality, documentation, logging, transparency, human oversight, accuracy)?
2	HU-RAC classification	Classification into HU-RAC 1–5 along the six axes (Chapter 4.1), applying the aggregation logic (Chapter 4.2). Determination of the business-criticality level A/B/C (Chapter 6.1) and the provider-resilience class R1/R2/R3 (Chapter 6.2).
3	GDPR alignment	Are personal data processed? If so: Art. 6 legal basis, Art. 9 special categories, Art. 28 DPA, Art. 35 DPIA obligation (regularly required from HU-RAC 3 onwards), Art. 44 ff. third-country transfer, Art. 48 orders by third countries.
4	HU-RAC-specific controls	Derivation of additional expectations from Chapter 5: key sovereignty, tenant isolation, data-subject protection, cybersecurity, standards connection. Split into E-Min (obligation or market standard) and E-Plus (aspirational, conservative). In addition: sectoral law (MDR, medical-device operator ordinance, Critical Infrastructure Umbrella Act, police law, social-security law) and the Machinery Regulation (MR 2023/1230).

E.2 Handling Discrepancies

Three constellations are to be distinguished:

- AI Act high risk without corresponding HU-RAC high category: occurs rarely in practice but is conceivable. The AI Act with its specific obligations remains applicable; HU-RAC sits alongside and adds data-sovereignty-related expectations.
- HU-RAC 4 or 5 without AI Act high risk: typical case for systems with non-AI-based core functions in sensitive deployment environments. HU-RAC here leads to additional expectations going beyond the AI Act on key sovereignty, tenant isolation and data-subject protection.
- Both high: the regular case for humanoid care or medical applications with adaptive AI components. The obligations of both frameworks accumulate; HU-RAC orients implementation context-specifically.

E.3 Example Application: Surgical-Assistance System

Step 1 AI Act: high risk (Annex III, medical intended purpose, safety component under MDR). All AI Act obligations apply. Step 2 HU-RAC: data categories health data under Art. 9 GDPR (very high), affected persons patients with endangerment potential (very high), AI high risk (very high). HU-RAC 4, Level A, R1 or R2. Step 3 GDPR: Art. 9 applicable, special legal basis (typically Art. 9(2)(h)), DPIA mandatory. Step 4 HU-RAC specific: additional E-

Min controls from Chapter 5.1 (EU/CH data processing, DPA, DPIA documentation) and 5.4 (Hardware Root of Trust, HYOK options, customer-side HSM recommended). E-Plus controls such as zero knowledge and independent cryptography audit as roadmap elements. Sectoral law: MDR with CE conformity and, where applicable, notified body, medical-device operator ordinance, social-security law.

Contact and Feedback

This impulse paper is intended as an invitation to dialogue. Feedback, criticism, proposals for extensions and pointers to related work are welcome. The author will review the suggestions and revise the paper as needed. A second version is conceivable once sufficient resonance has been received.

Humanide GmbH

– *Contact person and contact details to be inserted here* –

Citation: Humanide GmbH (2026): HU-RAC – A Conceptual Framework for Classifying Smart Robotic Systems by Data Sovereignty and Data Security. Impulse Paper, Version 0.3.1, April 2026.