

Easycore — System Card

Identity card of the Easycore agentic intelligence platform — Edition 2026 · v2026.06 · Public

PURPOSE OF THIS DOCUMENT

This System Card identifies the Easycore platform, its intended purpose, architecture, capabilities, limitations, data handling, models and monitoring. It is the reference document for IT departments and regulators, published without NDA, as handed over in due diligence.

1 System identification

Element	Value
System	Easycore, agentic intelligence platform
Provider	EasyLab AI, Luxembourg (RCS B276290)
Classification	General-purpose AI system, limited-to-minimal risk under Regulation (EU) 2024/1689. Risk additionally assessed per deployment context; deployments operate as deployer applications within the meaning of Article 26.
Intended purpose	Autonomous AI agents for B2B work augmentation under human supervision: preparation of professional work products, never autonomous decisions with legal effect.
Excluded uses	No credit scoring, no biometrics, no employment-decision automation, no Annex III use case.

2 Architecture overview

Each client deployment is an independent, isolated instance combining: an autonomous agent driven by an external reasoning runtime through an agent control protocol (ACP); a hardened gateway acting as control plane across communication channels; ephemeral, tool-restricted sub-agents for any processing of external content; and tool servers exposing memory, knowledge-graph and orchestration functions.

- **LLM-agnostic by design.** Easycore is not a model maker. The reasoning model is an interchangeable supplier behind the ACP abstraction, switchable to European-hosted models without re-architecting the platform.
- **Memory.** Proprietary tri-layer memory (working, episodic, distilled), stored locally per client. Read memory involves no third-party model.

- **Isolation.** Strict per-client isolation of data, memory and runtime at infrastructure level (dedicated volumes, networks, collections). The gateway binds to loopback; remote access only through encrypted tunnels; public exposure is prohibited in production.

3 Capabilities and limitations

3.1 Capabilities

- Autonomous task execution and scheduling, 24/7, within a defined perimeter.
- Persistent memory with extraction, distillation and consolidation cycles.
- Preparation of function work products: compliance templates, NAV assembly support, reporting packs, KYC file preparation.
- Inter-agent coordination under the orchestration layer: tickets, goals, budgets, approval gates, audit trail.

3.2 Limitations

- No outbound action without passing the applicable approval gate; sensitive actions require a named human approver.
- No self-modification of the security configuration; the envelope is locked and runtime-enforced.
- Outputs are professional drafts requiring human validation per the deployment risk level.
- Operation bounded by per-agent budgets: tokens, calls, spend, human-validation time.

4 Data and privacy

EU hosting with a sovereign provider; client data and memory remain in Europe; data minimisation by design. A single encrypted outbound flow exists: the call to the reasoning model, documented to the client and switchable to EU-hosted models. Strong encryption (AES-256-GCM pattern) is used for sensitive data at rest, proven on production systems handling health data.

5 Monitoring and behavioral trust

Every six hours, an independent scoring service evaluates each agent across **nine reliability dimensions**: reliability, gate compliance, output quality, security, memory consistency, budget compliance, appropriate escalation, transparency, stability. **The agent never scores itself.** Scores, confidence bands and history are visible to the client in the supervision cockpit and exportable for audit.

Living document. This card reflects the state of the platform at the date of edition and is updated as the platform evolves. A detailed technical dossier is available to clients and authorities in due diligence.

Easycore by EasyLab AI — Public governance documentation, edition 2026 (v2026.06). Published for transparency and due-diligence purposes under Regulation (EU) 2024/1689. This document describes compliance by design and does not constitute legal advice. It reflects the state of the platform at the date of edition. Due-diligence inquiries: contact@easylab.ai — easycore.ai/governance