

Easycore — Instructions for Use

Conditions of use, supervision and output interpretation — Edition 2026 · v2026.06 · Public

PURPOSE OF THIS DOCUMENT

These instructions describe the intended conditions of use of Easycore deployments, the supervision they require and how outputs must be interpreted, in the spirit of the transparency obligations of the AI Act. Operating outside this envelope voids the intended use.

1 Intended conditions of use

- Easycore is deployed for professional users after a scoping phase that defines use cases, connected systems, risk levels and named approvers.
- Each deployment requires at least one trained supervisor with access to the supervision cockpit.
- The security envelope must remain intact: gates configured, budgets set, audit trail enabled.

2 Understanding outputs

Risk level	Perimeter	Execution mode
GREEN	Internal drafts, no outbound action.	Auto-executed, journaled.
ORANGE	Output to third parties or moderate financial impact.	Human approval gate.
RED	Regulatory submission, signature, payment, sensitive data.	Named gate + reinforced audit.

- **Drafts, not decisions.** Agent outputs are prepared work products; their release is governed by the risk level above.
- **Trust bands.** Each agent displays a Behavioral Trust Score recalculated every six hours by an independent service. A degraded band is a signal to increase supervision; a critical band suspends autonomous operation.
- **Escalation.** Agents are designed to escalate uncertainty to humans rather than guess on material matters.

3 Required competence

Supervisors must understand: the configured perimeter of each agent; the meaning of risk levels and gates; how to pause an agent and revoke an action awaiting a gate; how to read the audit trail. EasyLab AI provides supervisor onboarding at deployment.

4 Maintenance, updates and known limitations

- Security-relevant platform changes are communicated to clients; material model changes are notified and reflected in monitoring baselines (LLM-agnostic architecture).
- Language-model reasoning can produce plausible but incorrect content; gates and validation rules exist precisely for this reason.
- Agents do not browse or act outside their configured tools and channels; performance depends on the quality and relevance of the data made available by the deployer.

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