

Easycore — Acceptable Use Policy

Permitted use, prohibited use, detection and enforcement — Edition 2026 · v2026.06 · Public

PURPOSE OF THIS DOCUMENT

This policy defines what Easycore agents may be used for, what is strictly prohibited, and how violations are detected and handled. It applies to every deployment, agent instance and user with access to a cockpit, configuration or channel, and is incorporated by reference into deployment agreements.

1 Permitted use

- Professional work augmentation for the deploying organisation: preparation, assembly, reconciliation, drafting, monitoring and reporting under human supervision.
- Operation strictly within the configured perimeter: defined channels, tools, data sources and budgets.
- Use of outputs as drafts, subject to the validation rules of the applicable risk level (GREEN / ORANGE / RED).

2 Prohibited use

- Any Annex III high-risk use under the AI Act: credit scoring of natural persons, biometric identification or categorisation, employment selection decisions, or any automated decision producing legal effects on individuals.
- Bypassing, disabling or socially engineering approval gates, budgets or the security envelope.
- Processing data the deployer is not lawfully entitled to process; introducing third-party data in violation of confidentiality or professional secrecy obligations.
- Generating illegal content, harassment, deceiving third parties about interacting with an AI system, or any use contrary to applicable law.
- Resale, sublicensing or exposure of the agent to untrusted external parties without EasyLab AI's written agreement.

3 Detection and enforcement

Mechanism	Role
Audit trail	All agent activity is journaled in an immutable, timestamped, exportable log.
Behavioral Trust Score	Independent scoring every six hours surfaces drift and anomalous behaviour.

Mechanism	Role
Graduated measures	Alert, gate tightening, agent suspension, termination of the deployment for serious breach.

4 Reporting

Suspected violations or vulnerabilities should be reported to **contact@easylab.ai**. Reports are acknowledged and investigated; clients are informed of incidents affecting their instance in line with the incident-management procedure and applicable law.

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