Code of Ethics

ETHICS CHARTER | INTERNAL & PUBLIC

Company-wide standards for responsible AI development and deployment — CyberPlug DevSecOps Security AI

> *CyberPlug's Code of Ethics is a binding commitment by all staff, contractors, and AI systems to operate with integrity, fairness, and respect for the rights of all people affected by our technology. This charter is enforced by the Algorithmic Risk Committee (ARC) and reviewed annually.*

1. Core Ethical Commitments

- ✓ We will not deploy AI systems that cannot be explained to the people they affect
- ✓ We will not use human oversight as a scapegoat for system failures — blame must be systemic, not individual
- ✓ We will always disclose when AI is involved in decisions affecting our clients and their systems
- ✓ We will rotate human oversight roles annually to prevent automation bias entrenchment
- ✓ We will publish an annual Human Oversight Transparency Statement — publicly, without redaction of key metrics
- ✓ We will not deploy a model we cannot halt — the CISO holds unconditional authority to stop the AAA System
- ✓ We will never train models on tenant data without explicit consent and documented anonymisation
- ✓ We will treat algorithmic errors as system design failures, not user failures

2. ARC Ethical Mandate

The Algorithmic Risk Committee holds standing authority to halt any AI feature that conflicts with fundamental rights, regardless of commercial impact. ARC decisions on ethical grounds cannot be overruled by business units.

> *The ARC operates under institutionalised distrust principles (Laux, 2023): trust in our AI systems is earned through structural safeguards, not assumed. Every governance mechanism is designed on the assumption that the AI — and the humans overseeing it — may be wrong.*

3. Operator Ethics Standards

Standard	Requirement	Enforcement
Fitting Intentions	Operators must prioritise the rights and safety of affected parties over performance metrics or throughput targets	ARC oversight; anonymous reporting channel
Independence	Operators must feel safe to override AI recommendations. No implicit or explicit pressure to defer to the system	Culture policy; ARC monitoring of override rates
Expertise	Operators must hold the qualifications appropriate to the systems they oversee. Under-qualified operators must not be assigned Critical oversight roles	Role-based training requirements; LO-1 to LO-6 certification
Disclosure	Operators must immediately disclose any conflict of interest, organisational pressure, or knowledge of system bias to the ARC	Whistleblower protection policy

4. Endorsed Frameworks

EDPS TECHDISPATCH 2025:	Human Oversight of Automated Decision-Making — operationalised in our Human Interactions Policy
FORHUMANITY AUDIT CRITERIA:	Human Oversight & Interaction — fully implemented; ARC-certified
LAUX (2023) MODEL:	Institutionalised Distrust framework — applied to ARC governance structure
GHANA AI STRATEGY:	Pillars 4 (Data Sovereignty) and 6 (Ethical Governance) — regional deployment compliant
ARISE FRAMEWORK:	Accountability, Risk, Integrity, Security, Explainability — embedded in all governance layers

5. Charter Certification

This Code of Ethics is approved by the CyberPlug Algorithmic Risk Committee. All staff complete ethics training (LO-6) annually. Violations are reported to the ARC and may result in access revocation or regulatory referral. Version 2.0 | April 2026. Next review: April 2027.