Acceptable Use Policy

POLICY | APPLIED ETHICS | USER FACING

Rules governing user interaction with CyberPlug AI features and the conditions of permitted use

This Acceptable Use Policy applies to all users of the CyberPlug platform, including tenant security engineers, asset owners, compliance officers, and any third parties granted access to CyberPlug systems. Acceptance of this policy is required as a condition of platform access.

1. Permitted Uses

- ✓ Security scanning of owned or explicitly authorised systems
- ✓ Receiving and reviewing AI-generated vulnerability summaries and remediation recommendations
- ✓ Overriding, dismissing, or escalating AI findings with a documented justification
- ✓ Accessing compliance-mapped audit reports and exporting them for regulatory purposes
- ✓ Configuring automated remediation workflows within approved policy boundaries
- ✓ Submitting feedback on AI findings through in-platform feedback forms or the support portal
- ✓ Requesting human review of any AI-generated finding at any time

2. Prohibited Uses

- ✗ Scanning systems, APIs, or infrastructure without ownership or explicit written authorisation
- ✗ Relying solely on AI output for Critical security decisions without the mandatory human review step
- ✗ Attempting to reverse-engineer model weights, thresholds, or confidence scoring logic
- ✗ Using platform outputs to infer personal data about individuals not related to system security
- ✗ Sharing tenant-specific vulnerability data with third parties without authorisation
- ✗ Circumventing, disabling, or interfering with the Human Interactions Log or audit trail
- ✗ Using CyberPlug to perform offensive security operations (penetration testing) without agreed scope of work

3. Human Review Obligations

Users must not treat AI-generated findings as final determinations. The following situations always require qualified human review before action:

Trigger	Required Reviewer	Documentation Required
Critical vulnerability (CVSS 9.0+)	Security Engineer + Asset Owner	Override or acceptance logged with rationale
AI confidence rated Low	Security Engineer	Finding reviewed and decision documented
Deployment block triggered	Asset Owner + Compliance Officer	Dual authorisation in Human Interactions Log
Risk acceptance selected	Asset Owner (CISO countersigns for Critical)	Business justification + ARC review
Payment or authentication system scan	Security Engineer + Compliance Officer	Mandatory sign-off in compliance report

4. Violations and Enforcement

VIOLATIONS REPORTED TO:	compliance@CyberPlug.cc | In-platform incident report function
CONSEQUENCES:	Access suspension; tenant notification; regulatory disclosure where required by law
APPEALS:	Any enforcement decision may be appealed in writing to ai-governance@CyberPlug.cc within 14 days
REVIEW CYCLE:	This policy is reviewed quarterly by the ARC and updated as required
EFFECTIVE DATE:	April 2026 | Version 2.0